Risk Trade-off Lab

Facilitation guide · team, one facilitator

What this is: a short, facilitated lab you run when two risks trade against each other and you can't eliminate both. Instead of pretending you'll mitigate everything, the team chooses which residual risk it will consciously carry — accept, and be responsible for watching — sizes the exposure, and gives it a named owner. The mechanism is the Risk Acceptance Ledger — the capture sheet is the ledger. It takes about 30–45 minutes for one risk trade-off.

The one line to hold: your job isn't to make the risk go away — it's to stop the team pretending it has. A risk that's named, sized and owned is being managed; one that isn't is just waiting. **You can't eliminate every risk. You can choose which one you'll carry — the danger you'll live with, named and owned.**

· WHEN DO I USE THIS?

Use this when

GOOD FIT

Reducing one risk clearly increases another and you keep trying to avoid both; there's a residual risk everyone senses but no one has named or owns; a decision means carrying some exposure whatever you choose; an audit or review flagged a risk you can't fully close; or a risk landed as a surprise before and you want the next one accepted on purpose.

WRONG TOOL

Two *goals* collide and you must choose which governs › Strategic Trade-off Lab. You want to imagine how a plan could fail › It Already Failed. You're mapping how different futures could unfold › Scenario Split. The trade-off is a conflict *between people* › Trade-off Conflict Resolver.

The boundary that matters most: this is not the Strategic Trade-off Lab. That tool governs two *goals* — which do we optimise for? This one accepts a *risk* you can't remove — which danger do we consciously carry? Governing a goal and owning a residual risk are different decisions; if you're choosing what to optimise, you want the other lab.

DON'T CONFUSE THESE TWO LABS

STRATEGIC TRADE-OFF LAB

two **goals** collide

choose which **governs**

which do we optimise for?

Problem Solving · Make Decisions

RISK TRADE-OFF LAB

two **risks** can't both go to zero

choose which you'll **carry**

which danger do we live with?

Constraints · Manage Risks

© Witmer

PART 1 RUNNING THE LAB

team, one facilitator

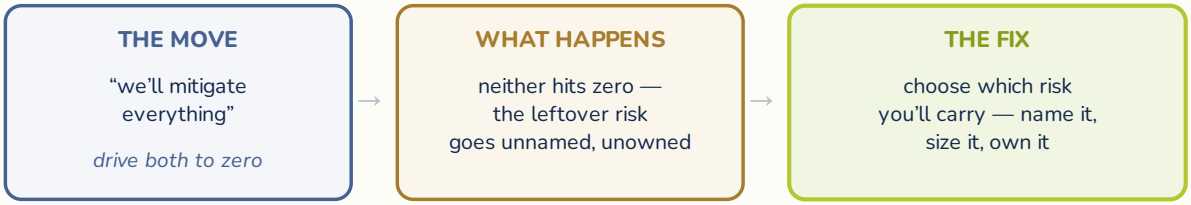
1 · WHAT IS THIS REALLY DOING?

The one idea

THE ONE IDEA

Teams treat “eliminate all risk” as the goal, so when two risks can’t both be driven to zero, the one left over gets hidden rather than chosen — it slips through unnamed, unsized and owned by no one. And an unowned residual risk doesn’t disappear; it waits, and lands later as a surprise. This lab makes the acceptance conscious: the team names the residual risk it will carry, sizes what it would cost, and gives it a single owner. Accepting a risk on purpose is what separates managing risk from being caught by surprise.

THE DRIVE-TO-ZERO TRAP



Unowned residual risk doesn’t disappear. It just waits.

© Witmer

2 · WHAT DO I NEED?

What you'll need

What	Why it's there
A real risk trade-off (core)	Two or more risks that genuinely trade against each other — reducing one increases another — where you can't drive them all to zero. If every risk can simply be fixed, you don't need this yet.
The Risk Acceptance Ledger (core)	The capture sheet is the ledger: it holds the competing risks, the one you accept, its exposure, and the named owner. Writing it down is what turns a vague worry into an owned decision.
A facilitator	Runs the ledger, makes sure the accepted risk is named and sized honestly, and holds the room to assigning a real owner — not "the team" in general. Needs the authority to challenge unrealistic "we'll mitigate everything" answers.
The team, and the person who'll own it	Whoever can genuinely carry the accepted risk needs to be in the room — ownership assigned to an absent person doesn't stick.

Constraints · Manage Risks · about 30–45 minutes for one risk trade-off, depending on how many risks are in play and how much the exposure is debated.

© Witmer

3 · HOW DO I RUN IT?

The run-sheet

Confirm it's a genuine trade-off first — you can't drive all the risks to zero. Then run the ledger. The whole lab turns on two disciplines: accept the risk *on purpose*, and give it *one named owner*.

Before you begin

- ✓ The risks genuinely compete — reducing one increases another
- ✓ You can't realistically eliminate both
- ✓ The likely owner is in the room
- ✓ Everyone understands this is about accepting a residual risk, not avoiding one

1. List the competing risks	**5 min**	*Write the risks that trade against each other on the ledger, in the team's real language.* Be specific — "a known Sev-1 ships" beats "security risk."
2. Confirm you can't eliminate them all	**3 min**	*Check it's a genuine trade-off: reducing one really does increase another, and you can't drive them all to zero.* If you can simply fix them, do that instead — you don't need the lab.
3. Choose the risk you'll accept	**7 min**	*Decide together which residual risk the team will consciously carry* — the one it makes most sense to accept given the alternatives. A decision the team reaches together, not a vote on preferences.
4. Explain why	**3 min**	*Before sizing it, ask the team to say why this is the risk they're prepared to carry.* Saying the reason out loud makes the acceptance deliberate rather than a quiet default.
5. Size the exposure	**6 min**	*State it plainly, as cost × likelihood: what it costs if the accepted risk lands, and how likely that is.* Out loud — a vague risk can't be managed, only feared. Size it honestly; making it sound smaller just recreates the surprise.
6. Assign a named owner	**5 min**	*Give the accepted risk to one specific person who says, out loud, "I'll carry this."* Choose someone who can genuinely monitor or act on it — not merely someone affected by it. One name, not "the team."

© Witmer

7. Set the watch trigger	**3 min**	*Agree what would make you revisit the decision: "what would tell us this risk is moving?"* A named trigger or review date tells the owner when to act, so acceptance isn't just a one-time shrug.
8. Record it on the ledger	**3 min**	*Capture the accepted risk, its exposure, its owner and its watch trigger where the team will see it.* The ledger is the record that the risk was chosen, sized, owned and watched — on purpose.

THE RISK ACCEPTANCE LEDGER



A risk that is named, sized and owned is being managed. One that isn't is just waiting.

FACILITATOR TIP — THE OWNER IS THE POINT

The moment the lab is most likely to go soft is the owner step. The team will want to leave it with "we'll all keep an eye on it" — which means no one will. Don't accept a collective owner. Ask for one name, and have that person say "I'll carry this" out loud, in the room. Ownership spoken aloud sticks in a way a name written quietly on a sheet does not. If no one will own it, that's important information: a risk no one will carry may be one the team hasn't really accepted.

PART 2 GETTING IT RIGHT

why it works & common mistakes

4 · WHY DOES THIS WORK?

Why it works

Teams treat "eliminate all risk" as the goal, so any risk that can't be removed gets hidden rather than chosen — and slips through unnamed. Naming it turns something waiting to go wrong into a decision the team owns; sizing it makes the cost real instead of vague; and giving it an owner means someone is actually watching it. Accepting a risk on purpose is what separates managing risk from being surprised by it.

© Witmer

5 · WHAT IF IT GOES WRONG?

Common mistakes

What happened	The fix
The risks weren't really a trade-off	Almost every failed session starts here. If both can simply be mitigated, do that — you don't need the lab. This is only for risks that genuinely can't both go to zero.
No one really owned it	"We'll all keep an eye on it" is no owner. Get one name, said out loud: "I'll carry this." Ownership by everyone is ownership by no one.
The owner couldn't actually influence the risk	An owner who can only watch it happen isn't an owner. Choose someone who can genuinely monitor or act on the risk, not merely someone affected by it.
The team tried to mitigate it away	If you could remove the risk you wouldn't be here. The job is to accept a residual risk consciously, not to make it disappear — pretending you removed it is how it comes back.
The exposure was understated	Naming a smaller cost to feel better recreates the surprise you came to avoid. Size it honestly: cost × likelihood, out loud.
It became "which goal do we want?"	That's the other lab. Here you're not choosing what to optimise — you're choosing which danger to carry. Keep it on the risks, not the goals.
It stayed a conversation, not a record	An accepted risk that isn't written on the ledger will be quietly forgotten. Record the risk, its exposure and its owner where the team will see it.

Before you close the lab

- ✓ The team chose which residual risk to carry — consciously, not by default
- ✓ The exposure was sized honestly — cost and likelihood, out loud
- ✓ One named person accepted ownership aloud — not "the team"
- ✓ A watch trigger or review date was agreed — the owner knows when to act
- ✓ The accepted risk, its exposure, its owner and its trigger are recorded on the ledger

© Witmer

6 · HOW DO I MAKE THIS STICK?

Take it back to work

THE TEST

Ask anyone on the team:

“Which risk are we carrying, why, and who owns it?”

If they can answer all three without hesitating — the risk, the reason, and the owner — it’s genuinely being managed. If they can’t name the owner, it isn’t — run the owner step again.

Appendix · Reference

You should not need this to start. Read it before, or if a run goes sideways.

7 · A WORKED RUN

Watch a risk get accepted

A team could harden a release against a known security flaw, or ship on the promised date — not both. They kept saying they’d do both, and the leftover risk sat unnamed.

Stage	What happened
List & confirm	Two risks: “a known Sev-1 ships” vs “we miss the launch commitment.” Reducing one increased the other — a genuine trade-off, neither removable.
Choose	The team decided to accept the schedule risk was worse for the business, and consciously carry the security risk — ship with the Sev-1 documented and contained.
Size	Exposure, out loud: “if exploited, ~2 days’ incident response and a customer-trust hit; likelihood low but not negligible in the first month.”
Own	The security lead said “I’ll carry this” — monitoring for exploit attempts, ready to hotfix, with a review at 30 days.
Record	Logged on the ledger: the accepted risk, its exposure, its owner, and the watch trigger — a 30-day reassessment. When it came up later, it was a managed decision — not a surprise.

© Witmer

8 · WHERE DOES THIS FIT?

Where this sits in the Gym

The real need	Use this
Consciously accept a risk you can't eliminate, and own it	**This tool — Risk Trade-off Lab.** Accept the residual risk, size the exposure, assign a named owner.
Choose which of two **goals** governs when they collide	**Strategic Trade-off Lab** — governs goals (which do we optimise?). This accepts a risk (which danger do we carry?). The sharpest boundary — different decision entirely.
Imagine how a plan could fail, before it does	**It Already Failed** — a premortem that surfaces failures. This chooses which surfaced risk to consciously accept.
Map how different futures could unfold	**Scenario Split** — branches scenarios. This trades two residual risks against each other and accepts one.
Resolve a trade-off that's a conflict between people	**Trade-off Conflict Resolver** — for interpersonal conflict. This is for accepting a risk, whoever agrees.

The move underneath: take a danger the team is carrying by accident and make it a choice — named, sized and owned — so no one is surprised by it later. That same move is one of the Gym's most steadying, wherever a team is quietly hoping a risk won't land instead of deciding who carries it.

© Witmer