Estimate the Unknown

Facilitator guide · solo or team

Your job: take one unknown the team has to put a number on, break it into parts, get an honest range on each, surface the two or three things really driving the uncertainty, and combine into one band the team can defend — with the conditions that move it to each end. Run it solo in 10–15 minutes, or with a team in 30–45 — on a worksheet or whiteboard, no kit.

The one line you must hold: the output is never a single number. If the room drifts toward "so, 17 days?", pull it back to the band and its conditions. **A range you can defend beats a number you can't.**

· WHEN DO I USE THIS?

Use this when

GOOD FIT

A real estimate you have to commit to with genuine unknowns in it — sizing a sprint, a release, a migration, a discovery spike — where a single confident number would be a guess.

WRONG TOOL

You don't know what question to ask yet › Question Finder. You're tracing why something broke › Root Cause Drill. You want to express confidence in a claim › How-Sure Line. Judging whether a past estimate was right › Calibration Trap.

The one test before you start: can you name the unknown as a single number you have to commit to? If yes, this tool sizes it. If you can't even frame the question, decompose the problem elsewhere first.

PART 1 SET UP — 2 minutes

1 · WHAT IS THE IDEA?

Your job and the one idea

You run this to one outcome: a band you can explain confidently, with its assumptions and drivers visible — not a single number. Solo, you're making your own estimate honest. With a team, you're making the group's uncertainty visible, which only happens when people estimate independently and then compare.

THE ONE IDEA

A whole unknown is unguessable, but its parts usually aren't. Decomposing turns one impossible guess into several reasonable ones — and the places where the parts are widest, or where people disagree, are exactly where the real uncertainty lives. Your job is to make that uncertainty visible and put a range on it, not to manufacture a single confident number.

© Witmer

2 · WHAT DO I NEED?

What you'll need

What	Why it's there
One real estimate	A number the team has to commit to — a sizing, a timeline — that can't be known precisely. Bring a live one, not a hypothetical.
The Estimate Worksheet	Columns for the parts, a low and high for each, the assumption behind it, and a confidence note. Landscape, so the ladder has room to grow. One sheet each for the independent pass; one shared sheet to combine.
Pens or markers	Solo: a pen. Team: one per person for the independent pass, plus a shared sheet or whiteboard to combine on.

Problem Solving · Work with Unclear Inputs · solo 10–15 min, team 30–45 min. The time is indicative and will depend on how many parts the estimate has and the size of the team.

Two words to keep straight — you'll use both: an **assumption** is something you're taking to be true ("the API is stable"); an **uncertainty driver** is something that could change the estimate ("we've never used this API"). Every wide range should have at least one of each written next to it.

PART 2 RUN IT

solo or team

THE SINGLE MOST IMPORTANT RULE

Never let people estimate together first. Independent estimates create the disagreement that reveals the uncertainty — that is the engine of the tool, and everything else depends on it.

Before you begin — 30-second check

- ✓ There's one concrete number to commit to — not a vague topic
- ✓ The people who'll own the estimate are the ones doing it
- ✓ Each person has a worksheet (team: for the independent pass)
- ✓ You're ready to hold the line: a band, never a single number

3 · HOW DO I RUN IT?

Run sheet — solo or team

One run sheet, both modes. The solo path is the same five steps without the compare. For a team, the independent-then-reveal rhythm is the whole point — don't let people estimate in a group discussion, or you lose the disagreement that surfaces the uncertainty.

© Witmer

1. Write the unknown	**2 min**	*Write the one number you have to commit to at the top: "how long to migrate billing?"* Team: appoint a facilitator first. One concrete number, not a topic.
2. Break it into parts	**5–8 min**	*List the parts — the pieces of work you can reason about separately — one per row.* Solo: write your own. Team: each person decomposes silently on their own sheet, then: "let's compare our breakdowns — where did we split things differently? Let's agree one combined list of parts everyone can work with."
3. Put a range on each part	**5–8 min**	*For each part, write a low and a high — honest best and worst realistic case. Never a single number.* Team: do this alone first, don't discuss yet. A wide range is information, not weakness.
4. Name what's driving the uncertainty	**5–8 min**	*For the widest ranges, write the two or three things making them uncertain.* Team: reveal everyone's ranges together now. Wherever they differ most, ask why — those differences point straight at the real uncertainty.
5. Combine and state the band	**5–8 min**	*Combine the lows into one lower bound and the highs into one upper bound. Then write the two conditions: the low-end and high-end conditions: what moves it to the low end, what pushes it to the high end.* Team: "now let's agree one combined range — not an average, a band we can defend. What would move it to the low end? What pushes it to the high end?"

How many parts? When to stop decomposing: break the work down until each part feels easy to estimate on its own. If a part still feels like one big unknown, split it again. If splitting it further wouldn't change how you'd estimate it, stop.

The signal it's working: the conversation stops being about "what's the number?" and becomes "where's the hidden guesswork?" When the team is arguing about the third-party API rather than the total, the uncertainty is visible — that's the point of the tool.

© Witmer

When to step in — and when to stay out · a principle, not a step

STEP IN

Someone writes a single number → ask for the low and the high.

The room estimates in a group discussion → reset to independent first.

A wide range gets no driver → "what's making this one uncertain?"

The team drifts to one figure → pull back to the band and its conditions.

STAY OUT

Estimates differ widely → good; that gap is the uncertainty. Don't average it away.

A range feels uncomfortably wide → leave it; honesty beats false precision.

The team debates a driver → let it; that's the real work.

Someone revises after the reveal → fine; comparing is meant to update them.

The clearest tell it's time to step in: the team has quietly converged on one tidy number and looks ready to move on. Step in: "what's the range, and what would move it to each end?" A single number that feels certain is the exact thing this tool is designed to prevent.

PART 3 AFTER

make it stick

4 · WHAT WILL PEOPLE SAY?

What people say — and your move

What they say	Your response
"Can't you just give me one number?"	"A single number would be a guess. The band is the honest answer — and I can tell you what moves it."
"This range is too wide to be useful."	"The width is the information — it's telling you where the real unknown is. Narrow it by resolving the driver, not by pretending."
"We're all just guessing anyway."	"Decomposed, we're not — each part is reasoned. And where we disagree tells us exactly what to go find out."
"Let's just average everyone's estimate."	"If two experienced engineers disagree by four days, the average isn't the truth — the disagreement is the thing we need to understand."
"Do we really need the assumptions written?"	"Yes — an estimate without its assumptions is just a number no one can defend later."

© Witmer

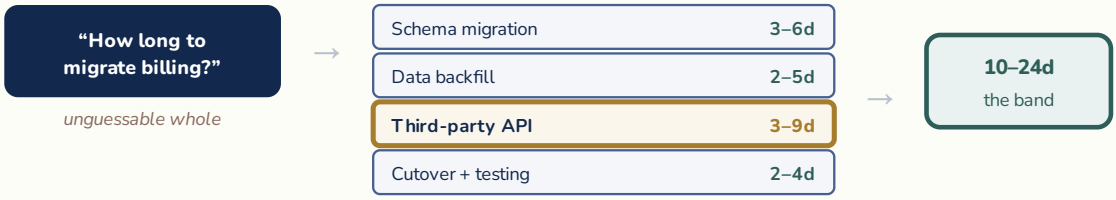
5 · WHY DOES THIS WORK?

Why it works

A whole unknown is unguessable, but its parts usually aren't — decomposition turns one impossible guess into several reasonable ones. And precision creates false confidence: people trust a single number more than a range, even when the range is the better answer. Naming the range and its drivers makes the uncertainty visible instead of hidden — and in a team, the places people disagree show where the hidden guesswork still lives. The tool's real job isn't better estimates; it's making that guesswork visible before it becomes a missed deadline.

ONE UNKNOWN → PARTS → RANGES → A BAND

you can't guess the whole — but you can reason about each part



← widest range = where the hidden guesswork lives

10 if the API is stable
↓ what moves it to the low end

24 if approvals slip
↑ what pushes it to the high end

"10–24 days — 10 if the API is stable, 24 if approvals slip."

A range you can defend beats a number you can't.

SAME LADDER, TWO MODES

SOLO · 10–15 MIN

1 · Write the unknown
2 · Break into parts, range each
3 · Name what's driving the width
4 · Combine into one band
5 · Write the low/high conditions

a defensible number, fast

TEAM · 30–45 MIN

Each person decomposes **alone** first
↓ no discussion yet
Reveal all ranges **together**
↓ where they differ = the uncertainty
Agree assumptions → one band

the differences are the point

Before you finish — the check

- ✓ The estimate is a band, not a single number
- ✓ Every wide part has its assumption and its uncertainty driver written
- ✓ The two conditions are stated: what moves it low, what pushes it high
- ✓ Team: the band was agreed after comparing independent estimates, not averaged on the spot

© Witmer

6 · HOW DO I MAKE THIS STICK?

Take it back to work

THE HABIT

Whenever someone asks for a number, give a band and name what moves it.

The worksheet trains the reflex: decompose before you estimate, always answer with a range, and write down what would change it. Teams that run this a few times stop asking “what’s the number?” and start asking “what’s driving it?” **A range you can defend beats a number you can’t.**

Appendix · Reference

You should not need this while you run the activity. Read it before, or use it if a run goes sideways.

7 · A WORKED RUN

Watch the disagreement surface the unknown

A team had to size a discovery spike everyone was quietly guessing at.

WORKED RUN — SIZING A DISCOVERY SPIKE

Decompose: the spike splits into research, prototype, review, write-up — four parts, not one fog.

Two estimators differ: on “prototype” one says 2 days, one says 6 — that gap surfaces the real unknown.

The driver: the prototype needs an API no one’s used — that’s the uncertainty, named and written down.

The band: “6–12 days — 6 if the API behaves, 12 if we have to build around it.” A number they can defend.

The disagreement wasn’t noise — it was the uncertainty, made visible.

What the facilitator did: when the two estimates for “prototype” came back 2 days apart from 6, they didn’t split the difference. They asked why — and the gap surfaced an unfamiliar API neither had flagged. That became the named driver, and the band (“6–12 days, 6 if the API behaves, 12 if we build around it”) carried the real risk on its face. The disagreement wasn’t noise; it was the uncertainty, made visible.

© Witmer

8 · WHAT IF SOMETHING GOES WRONG?

Common mistakes and edge cases

What happens	What to do
The team estimates in a group discussion	You lose the independent signal. Reset: everyone decomposes and ranges alone first, then reveal. The disagreement is what makes the estimate useful.
Ranges get averaged into one number	Averaging hides the uncertainty. Keep the band; narrow it only by resolving a driver, never by splitting the difference.
Every part gets the same wide range	Push for honesty per part — some parts are well understood. If everything's truly unknown, the decomposition isn't fine enough yet; break it down further.
A part can't be estimated at all	That's a real finding: it's a research task, not an estimate. Size a spike to remove the unknown, then estimate.
Someone insists on a single point	Give the band and the midpoint if they must have a headline — but keep the range and conditions attached. The point alone is the failure mode.
The estimate is genuinely tiny	If there's no real unknown, this tool is overkill — just estimate it. Use this when the unknowns are real.

9 · WHERE DOES THIS FIT?

Where this sits in the Gym

The real need	Use this
Put a defensible number on something with real unknowns	**This tool — Estimate the Unknown.** Decompose, range, surface the drivers, commit to a band.
You don't yet know what question to ask	**Question Finder** — find the real question before estimating it.
Trace why something actually went wrong	**Root Cause Drill** — get past symptoms to the cause.
Express how confident you are in a claim	**How-Sure Line** — put calibrated confidence on a statement.
Judge whether a past estimate was any good	**Calibration Trap** — check estimate quality after the fact.

© Witmer