



Founders AI

founders-ai.dev

PRACTICAL AI FIELD GUIDE

AI Decision System

How to Compare Options Without Letting AI Decide for You

A practical decision guide for founders, business owners, creators, consultants, and professionals who want AI to compare options, surface trade-offs, and improve judgment without taking over the decision.

- Name the decision
- Generate alternatives
- Set criteria
- Score with evidence
- Red-team advice
- Write the brief

AI is very good at helping you think through a decision. It can summarize options, list trade-offs, build comparison tables, find missing risks, challenge your assumptions, and turn messy thinking into a decision brief.

But AI should not quietly become the decision-maker.

This guide gives you a practical system for using AI as a decision assistant while keeping human judgment in charge. Use it for founder decisions, business decisions, work decisions, creator decisions, tool comparisons, offer design, project prioritization, and everyday choices where you want a clearer process.

The rule:

Let AI improve the decision. Do not let AI own the decision.

What this guide helps you do

Use this when you need to:

- Compare tools, vendors, offers, strategies, content ideas, or projects.
- Choose between multiple “good enough” options.
- Stop asking AI for one confident answer too early.
- Make trade-offs visible instead of vague.
- Avoid letting AI’s first recommendation anchor your thinking.
- Document why you chose something.
- Know when a decision needs human review, expert review, or no AI at all.

This is not a legal, medical, tax, financial, hiring, safety, or compliance decision substitute. For high-stakes decisions, use AI to organize your thinking and then get qualified human review.

The core idea

Most people use AI for decisions like this:

Which option should I choose?

That is too much authority to give away at the start.

Use this instead:

Help me make a better decision. Do not decide for me yet. First frame the decision, list alternatives, define criteria, surface trade-offs, identify risks, and tell me what evidence is missing.

The difference is control. You are not asking AI to be the judge. You are asking AI to build the courtroom.

The AI Decision System

Use this loop:

1. Name the decision.
2. Define the owner and stakes.
3. Frame the real question.
4. Generate alternatives.
5. Set decision criteria.
6. Score options with evidence.
7. Map trade-offs.
8. Red-team the recommendation.
9. Decide what AI cannot decide.
10. Write the decision brief.
11. Create a review loop.

This system is slower than asking “what should I do?” It is much faster than making a confident bad decision.

Step 1: Name the decision

A decision is not a topic. It is a choice between actions.

Weak:

Help me think about content.

Stronger:

I need to decide whether my next 30 days of content should focus on AI prompting, AI workflows, or AI for business owners.

Strongest:

I need to choose one primary YouTube content lane for the next 30 days. The options are AI prompting, AI workflows, AI research, or AI for business owners. The goal is audience growth, replay value, and eventual product demand.

The clearer the decision, the less likely AI is to wander into generic advice.

Decision naming prompt

Help me turn this messy situation into a clear decision.

Situation:

[explain what is going on]

Return:

1. The decision I actually need to make
2. The options currently on the table
3. The decision owner
4. The deadline
5. What happens if I do nothing
6. What would make this decision successful

Step 2: Define the owner and stakes

Before comparing options, define who owns the decision and what is at risk.

Question	Your answer
Who owns the decision?	
Who is affected by it?	
How reversible is it?	
What could go wrong?	
What is the upside if right?	
What is the cost if wrong?	
Does this need expert review?	

NIST's AI RMF emphasizes governance, accountability, context, and risk management across AI use, including clear roles, responsibilities, monitoring, and documentation ([NIST AI Resource Center](#)). For everyday use, that means someone must own the decision, not just the prompt.

Step 3: Frame the real question

Bad framing creates bad decisions. If you ask the wrong question, AI may give you a polished answer that solves the wrong problem.

Ask:

- Is this a yes/no decision?
- Is this a choice between options?
- Is this a prioritization problem?
- Is this a timing decision?
- Is this a risk decision?
- Is this actually a values decision?

Frame-check prompt

```
Review this decision frame:  
[decision]
```

```
Tell me:
```

1. Is this framed as the right kind of decision?
2. What hidden assumptions are inside the frame?
3. What better framing could change the answer?
4. What options are being excluded too early?
5. What would a skeptical operator ask before deciding?

Strategic Decisions Group's decision-quality framework starts with setting the right frame, then considering alternatives, gathering information, clarifying values and trade-offs, reasoning correctly, and committing to action ([Strategic Decisions Group](#)). Start there before asking for recommendations.

Step 4: Generate real alternatives

AI will often compare the options you give it. That is useful, but incomplete.

Before scoring anything, ask it to create more options.

Alternative-generation prompt

For this decision:
[decision]

My current options are:
[options]

Generate additional alternatives in these categories:

1. The simplest option
2. The cheapest option
3. The fastest option
4. The highest-upside option
5. The lowest-risk option
6. A hybrid option
7. The option I am probably overlooking
8. The option that preserves future flexibility

Do not recommend yet. Just expand the option set.

Saveable tool: Option list

Option	What it means	Why it could work	Why it could fail
A			
B			
C			
Hybrid			
Do nothing			

Always include “do nothing.” It is often the hidden option everything else should beat.

Step 5: Set decision criteria

Do not let AI choose the criteria by itself. Criteria are where values hide.

For a founder, the criteria might be:

- Revenue potential
- Speed to launch
- Cost
- Difficulty
- Audience fit
- Strategic value
- Risk
- Reversibility
- Energy required
- Proof generated

For an everyday learner, the criteria might be:

- Ease of use
- Time required
- Cost
- Confidence gained
- Practical benefit
- Maintenance burden
- Risk of getting overwhelmed

Criteria prompt

Help me define decision criteria for this choice:
[decision]

Use my context:
[context]

Return:

1. 6-10 criteria
2. Why each criterion matters
3. How to score each one from 1-5
4. Which criteria are must-have vs nice-to-have
5. Which criteria might conflict with each other

Step 6: Score options with evidence

A score without evidence is just a vibe in spreadsheet form.

Use this table:

Option	Criterion	Score	Evidence	Confidence
		1-5	What supports this score?	High / medium / low

Scoring prompt

Score these options against the criteria below.

Decision:
[decision]

Options:
[options]

Criteria:
[criteria]

Rules:

- Use a 1-5 score
- Explain the evidence behind each score
- Mark confidence as high, medium, or low
- Do not average scores until after the evidence is shown
- Flag any score that is mostly assumption

This keeps the comparison honest. The evidence matters more than the number.

Step 7: Map trade-offs

There is rarely a perfect option. If there were, you probably would not need a decision system.

Ask AI to make the trade-offs explicit.

Trade-off prompt

Compare the top options by trade-off, not just score.

For each option, show:

1. What I gain
2. What I give up
3. What risk I accept
4. What future option it preserves
5. What future option it closes
6. Who would prefer this option and why

Trade-off table

Option	Gain	Give up	Main risk	Preserves	Closes

This is where the real decision usually appears.

Step 8: Red-team the recommendation

Once AI recommends something, do not accept it immediately. Challenge it.

Automation bias is the tendency to over-rely on automated recommendations. The EU AI Act's human-oversight language explicitly tells overseers to remain aware of over-reliance on AI outputs, correctly interpret them, and retain the ability to disregard, override, or reverse them ([EU Artificial Intelligence Act](#)). Studies in medical decision support show why this matters: experienced users can still be influenced by incorrect AI suggestions ([Radiology](#), [JAMIA](#)).

Red-team prompt

Red-team your own recommendation.

Assume your top recommendation is wrong.

Tell me:

1. The strongest case against it
2. What evidence could disprove it
3. What assumption it depends on most
4. What kind of person or business should not choose it
5. What second-best option may actually be better
6. What early warning signs would show this decision is failing

Do not skip this step. This is where AI stops being a cheerleader and starts being useful.

Step 9: Decide what AI cannot decide

Some parts of a decision are not analytical. They are personal, strategic, ethical, legal, or values-based.

AI can help clarify:

- The options
- The evidence
- The trade-offs
- The risks
- The likely consequences
- The questions to ask

AI should not own:

- Your values
- Your risk tolerance
- Your ethics
- Your legal responsibility
- Your customer promises
- Your family priorities
- Your final accountability

Human-judgment prompt

For this decision, separate what AI can help with from what only I should decide.

Return two columns:

1. AI can assist with this
2. Human judgment must own this

Then list any parts that require expert review before action.

If the AI cannot clearly separate those columns, slow down.

Step 10: Use override rules

Before you decide, define when you would ignore the AI recommendation.

Override rules protect you from trusting a polished answer just because it sounds confident.

Override rule examples

- If the recommendation depends on a source I have not verified, I will not act yet.
- If the decision affects legal, tax, medical, hiring, safety, compliance, or major money outcomes, I will get expert review.
- If the AI recommends the fastest option but ignores long-term trust, I will pause.
- If the score is high but confidence is low, I will gather more evidence.
- If the recommendation conflicts with direct customer evidence, I will investigate before acting.
- If the decision is hard to reverse, I will require stronger evidence.

Override prompt

Create override rules for this decision.

Decision:

[decision]

Recommendation:

[AI recommendation]

Return:

1. Conditions where I should ignore or delay the recommendation
2. Conditions where I should seek expert review
3. Conditions where I should run a small test instead of committing
4. Conditions where the recommendation is safe enough to act on

Step 11: Write the decision brief

The output of the process should not be a giant chat thread. It should be a one-page decision brief.

Use this:

Decision:

Owner:

Deadline:

Options considered:

Recommendation:

Why this option:

Best evidence:

- 1.
- 2.
- 3.

Main trade-offs:

Risks:

Assumptions:

Override rules:

What would change the decision:

Next action:

Decision-brief prompt

Turn this analysis into a one-page decision brief.

Use:

- Decision
- Owner
- Options considered
- Recommendation
- Evidence
- Trade-offs
- Risks
- Assumptions
- Override rules
- What would change the decision
- Next action

Keep it practical. Do not hide uncertainty.

Step 12: Create a review loop

A decision system gets better when you review outcomes.

After you act, capture:

NIST's AI RMF emphasizes that risk management should continue as knowledge, methods, contexts, risks, benefits, and impacts evolve over time ([NIST AI Resource Center](#)). The same idea applies to personal and business decisions. The decision is not finished when you choose. It is finished when you learn from the result.

Review prompt

Review this past decision.

Original decision brief:
[paste brief]

Outcome:
[what happened]

Tell me:

1. What assumptions were right
2. What assumptions were wrong
3. What evidence mattered most
4. What warning signs I missed
5. What I should change in my decision process next time

The 20-minute AI decision workflow

Use this when you need clarity fast.

Minute 0-3: Name the decision

Turn this situation into a clear decision with owner, deadline, stakes, and options.

Minute 3-7: Expand options

Generate alternatives, including simplest, lowest-risk, highest-upside, hybrid, and do-nothing options.

Minute 7-11: Define criteria

Create decision criteria and scoring rules based on my context.

Minute 11-15: Score with evidence

Score each option, explain evidence, and mark confidence.

Minute 15-18: Red-team

Assume the recommendation is wrong. Make the strongest case against it.

Minute 18-20: Brief

Create the one-page decision brief with override rules and next action.

The “AI is deciding for me” checklist

Slow down if:

- You asked “what should I do?” before defining criteria.
- You only compared the options you already had.
- The AI gave a recommendation before showing trade-offs.
- You accepted a score without evidence.
- You did not define what would change your mind.
- You did not create override rules.
- You feel relieved because the AI gave you permission to do what you already wanted.
- The recommendation affects money, safety, legal exposure, customers, employees, or reputation.
- You cannot explain the decision without referring back to the AI.

If three or more are true, restart the decision process.

Best everyday use cases

Choosing AI tools

Use the system to compare price, setup time, risk, learning curve, integration needs, and actual use cases.

Picking content topics

Use the system to compare audience demand, replay value, personal credibility, product alignment, and filming difficulty.

Prioritizing projects

Use the system to compare revenue impact, speed, bottleneck removal, strategic value, and energy required.

Designing an offer

Use the system to compare customer pain, proof, delivery effort, price sensitivity, and trust gap.

Evaluating opportunities

Use the system to compare upside, downside, reversibility, required resources, and what has to be true for the opportunity to work.

One-page version

Copy this into your notes.

AI Decision System

1. Name the decision.
What choice am I actually making?
2. Define owner and stakes.
Who owns it, who is affected, and what happens if wrong?
3. Frame the real question.
Am I solving the right problem?
4. Generate alternatives.
What options am I missing?
5. Set criteria.
What matters and how will I score it?
6. Score with evidence.
What supports each score?
7. Map trade-offs.
What do I gain, give up, preserve, or close?
8. Red-team the recommendation.
Assume the top answer is wrong.
9. Separate AI help from human judgment.
What can AI assist with, and what must I own?
10. Create override rules.
When should I ignore, delay, test, or escalate?
11. Write the decision brief.
Recommendation, evidence, risks, assumptions, next action.
12. Review the outcome.
What happened, what was wrong, and what changes next time?

Final reminder

AI is not dangerous because it gives advice. It is dangerous when the advice sounds good enough that you stop thinking.

Use AI to expand your options. Use AI to clarify criteria. Use AI to challenge assumptions. Use AI to write the brief.

Then make the decision like an owner.