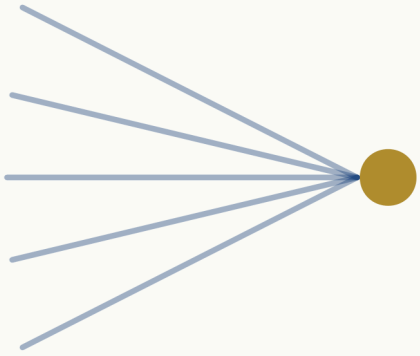
Expert Consilium

GUIDE · EDITION 5

A PRACTICAL OVERVIEW

One AI Answer, or Several Perspectives

When one AI answer is enough for a business decision, and when it is not



Multiple angles. One clear decision.

AUTHOR
Goran Granić

EDITION
November 2026

WEB
expertconsilium.com

01 / EXPERT CONSILIUM

The answer came back in thirty seconds

It is late. The supplier's offer is open on your laptop: a lower unit price, three years of exclusivity, a yearly minimum order. You paste the key terms into an AI chat and type one line: "Is this a good offer?"

Thirty seconds later the answer is on the screen. It is fluent, polite and confident. A lower price is a clear advantage. Exclusivity is common in the industry. The minimum looks manageable.

So why do you still not feel ready to sign?

This guide is about that feeling. It does not tell you to stop using AI. It shows you when one quick answer is enough, when it is not, and how to tell the difference before the contract is signed rather than a year after.

What you get from this guide

- A two-minute test that tells you how much checking a question deserves.
- One supplier offer worked through, with the numbers.
- The reason a single prompt tends to agree with you, backed by published research.
- Prompts and checks you can use tonight, in any AI tool.
- A one-page worksheet for the decision in front of you.

No technical knowledge is needed. If you have a real offer on your desk, keep it next to you as you read.

The examples in this guide are simplified educational scenarios, not client cases. This guide does not replace legal, tax, accounting or other regulated professional advice.

02 / EXPERT CONSILIUM

Why this question matters

In a small company there is no board to stop you and no legal department to read the fine print. The evening before you sign, you are usually alone with the papers.

It usually goes like this. You describe the offer as a good price from a reliable partner. The answer agrees, because that is the picture you gave it. The number that would change the picture, your sales over the last two years, sits in your own records, not in the chat. If the minimum is above what you sold, one slow year means paying for stock that waits in the warehouse until it expires. The tool was not wrong. The question was incomplete.

Everything in this guide works with any AI tool.

03 / EXPERT CONSILIUM

For plenty of work, one answer is fine

AI skeptics rarely admit it, but a single AI chat is genuinely useful. It drafts an email to a customer in German. It summarizes a forty-page tender document. It explains a term in a contract you have not seen before. It tightens the wording of an offer.

For these tasks one quick answer is enough, because a mistake is cheap and you will see it right away. You read the draft, fix what is wrong and move on.

It changes when the answer becomes the basis for a commitment. You sign a multi-year contract. You accept exclusivity. You agree to 90-day payment terms for a large retail chain. You change your price list. You hire the one person the business will depend on. Here a mistake is not cheap, and you usually find out only when the cost shows up.

The question to ask before you open the chat

Before you type anything, ask yourself one thing:

If this answer is wrong, who will notice, and when?

If the honest answer is "nobody, until we lose money," the question deserves more than one quick reply. It deserves a check of the assumptions, the missing information and the conditions under which the conclusion would be wrong.

Write it down: The commitment you are considering is _____. If you get it wrong, the cost is _____, and you would find out by _____.

04 / EXPERT CONSILIUM

Four kinds of decisions, four levels of checking

Not every decision deserves the same effort. Place yours in one of these four groups.

Type of decision	Typical example	What usually fits
Easy to undo, small cost	A draft text, a summary, an explanation of a term	One quick answer. Read it and correct it.
Can be undone with some effort	Trying a new tool, a small first order, a one-month pilot	A quick answer, then check the key numbers yourself and ask what is missing.
Hard to undo or expensive	A multi-year contract, exclusivity, a price change, a key hire, relying on one large customer	A structured check of assumptions, a serious case against, and open questions before you sign.
Legal, tax or regulatory effect	Contract wording, tax treatment, employment terms	Outside the scope of this guide. AI can help you prepare the questions, not answer them.

The two-minute test: do you need more than one answer?

Score each question from 0 to 2 and add up the points. This is a rough aid for your own judgment, not a measurement.

Question	0 points	1 point	2 points
Can you undo it?	Yes, easily	Partly	No, or only at a high cost
How much do you lose if you are wrong?	Little	A noticeable amount	Enough to hurt the business
How soon would you notice?	Right away	In a few weeks	Only when the loss shows up
Will anyone challenge it before you sign?	Yes, someone whose job that is	Maybe	No one
How solid is your information?	Documents and figures	Partly estimated	Mostly your impression

What your total suggests

- **0 to 3:** one quick answer is usually enough.
- **4 to 6:** check the key assumptions yourself with the prompt and the five checks later in this guide.
- **7 to 10:** the decision deserves a structured check before you commit.

Notice the fourth row. In a small company it is often the one that tips the score. Nobody's job is to say no to the owner.

Write it down: Decision: _____. Points: __ of 10. First check: _____.

06 / EXPERT CONSILIUM

One offer, worked through

Simplified educational scenario based on a typical situation.

A supplier offers you a lower unit price in exchange for three years of exclusivity and a yearly minimum order. You type the situation into an AI tool and ask: "Is this a good offer?"

What a short question usually gets

The answer tends to follow the direction the question already set. The lower price sounds favorable. Exclusivity is described as standard practice. The minimum sounds like a small formality. The answer is coherent and sure of itself. But it works only with what you gave it. It does not check whether your assumptions are right, because nobody asked it to.

What stays unchecked

- Whether your realistic yearly sales actually reach the minimum.
- What happens to the minimum if demand drops in year two.
- How you get out of the exclusive contract, and what it costs.
- Whether exclusivity stops you from negotiating with others when market prices fall.
- Whether the price is fixed in euros while you sell in dinars, and who carries the exchange rate.

A short prompt rarely raises these on its own. A better prompt can. So can the seven checks in the guide Before You Say YES.

Now with numbers, and with your sales

An illustrative calculation with round numbers, not data from a real offer. Today you pay EUR 50 per unit and buy as needed. The offer is EUR 46 per unit, 8% less, with a yearly minimum of 2,000 units. Last year you sold 1,700 units. The year before, 1,900. The AI tool knows none of this unless you paste it in.

Your yearly sales	Today (EUR 50, as needed)	Offer (EUR 46, min. 2,000)	Units left in stock
1,700 (like last year)	EUR 85,000	EUR 92,000	300
1,900 (like the year before)	EUR 95,000	EUR 92,000	100
2,000	EUR 100,000	EUR 92,000	0

Where the two cost the same. The offer costs at least 2,000 × 46 = EUR 92,000 a year. At today's price that buys 92,000 ÷ 50 = 1,840 units. If you sell fewer than 1,840, the lower price costs you more.

07 / EXPERT CONSILIUM

One offer, worked through

In a year like last year, the lower price costs you EUR 7,000 more and leaves 300 units in the warehouse. In a year like the one before, you save EUR 3,000 and hold 100 extra units. Whether this is a good offer depends on which of those two years comes next.

This is a break-even point for purchase spending in a simplified example. It is not a full profitability test.

That is the number the short question never saw. Paste your sales history into the prompt and ask the tool to build this table. The full method, with the other terms that change the result, is in check 2 and worked example 1 of Before You Say YES.

08 / EXPERT CONSILIUM

Why a single prompt tends to agree with you

The problem is not that AI models are bad. It is how one conversation works: one question, one answer, written to be helpful and consistent with what you already said.

The tendency to agree has a name

Researchers call it sycophancy: an answer that confirms what the user seems to expect, rather than what is accurate.

In March 2026, Stanford researchers published a study in *Science* that tested 11 large language models, including ChatGPT, Claude, Gemini and DeepSeek. On personal advice questions, the models endorsed the user's position on average 49% more often than people did. In a second part of the study, more than 2,400 participants rated the agreeable and the non-agreeable answers as equally objective. In other words, people could not tell when the AI was simply agreeing with them. [1]

That study looked at personal dilemmas, not supplier contracts. So read it as a warning, not as proof about business questions. But think about how you describe an offer you already like. You call it a good price. You mention the growth it could bring. The answer has every reason to follow you.

Nobody pushes back

In a single chat there is no second party to question the conclusion, ask where a figure came from or argue the other side. The answer stays unchallenged until you challenge it. And challenging your own preferred decision is exactly the hard part.

It is not the tool, it is the task

The same tool can argue against you, look for sources and work in several steps if you give it that job and enough information. What matters is not how many models you use. It is whether the process actually checks the questions that could change your mind.

09 / EXPERT CONSILIUM

Five patterns to watch for in a confident answer

These are patterns, not verdicts on any particular tool. A tool with a clear task and enough information can avoid several of them.

Pattern	What it looks like	What to ask
Follows your framing	Your question hints at a conclusion and the answer agrees	"Make the case against this as hard as you can."
Fills gaps with plausible details	Figures, dates or terms nobody gave it	"Which of these facts came from the description, and which did you assume?"
Treats the unknown as known	A firm conclusion without the missing data	"What would have to be true for this conclusion to be wrong?"
Stops at the first reasonable option	One path, no alternatives	"What are the alternatives, including waiting or a smaller first step?"
Sounds equally sure everywhere	Weak points written in the same firm tone	"Where are you unsure, and why?"

Each pattern is easy to miss because the answer reads well. The questions on the right are short on purpose. You can add them to any chat in a few seconds.

NIST's profile for generative AI lists false content presented as fact among the core risks of these systems. Data handling and answer quality need separate checks. [2]

A better prompt asks for the work, not the verdict

"Is this a good offer?" asks for a verdict. A better prompt asks for the work that should come before one.

A prompt you can copy and adapt

> Here is a supplier offer (pasted below). Do not say yet whether it is good. First, list the assumptions any conclusion would depend on. Then make the case against accepting it as hard as you can. List the information that is missing and the questions to ask the supplier. Do not state numbers you do not have: mark every figure as given in the offer or description, assumed, or unknown.

What a useful answer contains

- Assumptions named one by one, not buried in a paragraph.
- A case against that is specific to this offer, not general advice.
- Missing information written as questions you can actually send.
- No figures you did not supply, unless they are clearly marked as assumptions.

One small habit. The team behind the Stanford study also reports that simply telling a model to start its answer with "wait a minute" made it more critical. [1] Put that line at the start of your prompt. It takes a second, and it does not replace the checks below.

A better question improves one answer. It does not turn one answer into a checked one.

Five checks you can do yourself tonight

No special tool needed. These work in any AI chat, with a colleague, or on paper.

Check	How to do it
Ask twice	Put the same facts into two separate, fresh conversations. Compare where the answers differ.
Ask for the opposite	"Argue for the opposite conclusion as hard as you can."
Separate facts from guesses	Ask for three lists: stated in the documents, assumed, concluded.
Assume it failed	"It is a year from now and this deal went badly. What three reasons would you bet on?"
Check the three claims that carry it	Find the three claims the conclusion rests on. Check each one against a document or a person who knows.

Where these checks stop. They rely on the description you wrote. If an important document is missing, no number of questions replaces it.

Write down what would change your mind. Before you check, note the claim you rely on, how you will test it, and the result that would make you say no. This follows the logic of Strategyzer's Test Card: hypothesis, test, metric, threshold. [3]

Three situations where these checks earn their keep

These are typical situations for a small company in Serbia and the region, not stories from clients.

- **A supplier wants exclusivity.** Run "assume it failed" and ask specifically about year two demand and the exit clause. See also the guide Before You Say YES.
- **A large customer asks for a discount in exchange for volume.** Ask the AI to separate committed volume from promised volume. More on this in A Big Customer Asks for a Discount.
- **A retail chain wants 90-day payment terms.** Ask how much cash this ties up in your business and for how long. In Serbia a term over 60 days between companies needs instalments or a bank guarantee, so also ask what guarantee the chain will provide. See 90-Day Payment Terms.

Reading several answers without fooling yourself

Sometimes you end up with more than one answer: from two chats, from colleagues, or from separate AI perspectives. What matters is not how many you have. It is how you read them.

Look for	What it tells you
Where the answers agree	It may rest on facts you supplied, or on an assumption they all share. Find out which.
Where they differ	This is the part of the decision that is still open. Spend your checking time here.
A point only one answer raises	It may be noise, or the objection that matters. Check it before you set it aside.
The ranking of the answers	A ranking is not proof. Read the reasons, not the position.

Agreement is not proof. Five answers built on one unchecked assumption do not turn that assumption into a fact. When answers agree, ask what they all took for granted.

Keep the lone objection. One short objection can matter more than a long agreement among four. Write it down. Decide what evidence would settle it and who will get that evidence.

Disagreement is useful. It shows you where the decision is still open.

What it costs to skip this

A bad contract rarely ruins a business on its own. What happens is quieter. You accept a minimum you will struggle to reach. You find the exit clause only when you need it. You stop talking to other suppliers because the contract says you should. Each of these is a cost you could have seen on paper, the evening before, if someone had asked the right question.

When the decision still feels heavy

If you have done the checks above and the decision still feels heavy, give it a structured check before you commit. Three things are usually enough to start: separate answers that do not see each other first (two fresh chats, or colleagues who write their view before anyone speaks), one person or prompt whose job is to argue against, and a written summary that keeps the lone objection, not only the points of agreement.

Separate answers can still share a blind spot, because they all read the same description. If that description is missing something, they all miss it. More answers do not promise a more accurate result.

What a useful decision report contains

A report is not better because it is longer. It is useful when it shows what the recommendation rests on, where its limits are and what to do next. Use this as a checklist for any written analysis, from an AI tool, a colleague or an advisor.

Part	The question it answers
Decision and goal	What are we deciding, and what are we trying to achieve?
Facts and assumptions	What was supplied, what was assumed, what is still unknown?
Options	What are the alternatives, including waiting or a conditional yes?
Risks and trade-offs	What could change the choice?
Reasoned direction	Which step makes sense, and under what conditions?
Verification plan	Who gets which evidence, by when, and when do we look again?

Weak: "Take offer B, it is cheaper."

Stronger: "B is cheaper per unit. Whether it is cheaper overall depends on the minimum quantity and your realistic sales. Without those two numbers, an unconditional yes is not recommended."

A conditional recommendation ties the decision to what must be true. More on this in What a Good Decision Report Looks Like.

When it makes sense, and when it does not

A structured check does not make sense for every decision. If the question is easy to undo or cheap to get wrong, a quick AI answer or a talk with a colleague is enough. The same goes if someone in your company already has the job of challenging these decisions before you sign.

When the decision still feels heavy

It starts to make sense when three things are true at once: a mistake is expensive, it is hard to undo, and nobody in the company will push back before you sign. Even then, judge it by your specific situation, not by fear of what could go wrong.

Your decision on one page

Fill this in before you prompt an AI tool, or talk to your team. Keep it short, but make it define the decision. Mark the unknowns. Do not replace them with false certainty.

Decision and deadline:

Goal and constraints:

Facts you have, and where they come from:

Assumptions not yet confirmed:

Options, including waiting or a smaller first step:

What evidence would change your choice:

Next check: who, what, by when:

If you cannot say what would change your choice, ask yourself honestly whether you are looking for analysis or for confirmation.

Your next step

If there is an offer on your desk right now, start with the two-minute test and the five checks. They take about twenty minutes. If the score is still high after that, fill in the page above and run the structured check from *When the decision still feels heavy*.

Which question in your current decision is the one nobody has asked yet?

The final say is always yours.

Further reading

- Guide: Before You Say YES, seven checks before you accept a proposal.
- Guide: How AI Helps Businesses, from everyday tasks to reviewing important decisions.
- Blog: Why One AI Answer Is Not Enough for a Strategic Decision
- Blog: AI Sycophancy: When AI Tells You What You Want to Hear
- Blog: How to Make a Business Decision with AI

Sources

External sources support general statements about AI behavior and evaluation. Prompts, worksheets and scenarios in this guide are editorial suggestions.

1. Cheng et al., "Sycophantic AI decreases prosocial intentions and promotes dependence," *Science*, March 26, 2026. https://www.science.org/doi/10.1126/science.aec8352. Summary: Stanford Report, "AI overly affirms users asking for personal advice," March 26, 2026. https://news.stanford.edu/stories/2026/03/ai-advice-sycophantic-models-research

2. NIST, AI 600-1, *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*, July 2024. https://doi.org/10.6028/NIST.AI.600-1

3. Strategyzer, "Validate Your Ideas with the Test Card," 2015. https://www.strategyzer.com/library/validate-your-ideas-with-the-test-card

Sources checked on 21 November 2026. AI tool capabilities and terms change.

Informational material, not regulated professional advice. Examples are simplified educational scenarios. This guide makes no promise about the accuracy of any analysis or about any business outcome.