Scott Tischler

FOUNDER & CHAIRMAN · AIRECOMMEND.AI

STRATEGY

The “How Long Does It Take?” Query: Winning the Timeline Question in AI Answers

By Scott Tischler · October 17, 2026 · 8 min read

“How long does this take?” is one of the first questions a buyer asks — and more and more they ask AI before they ask you. If your timelines aren’t written down somewhere a model can read them, the model won’t stay quiet. It will answer with an estimate it assembled from category averages, a competitor’s page, or a plausible-sounding guess. And once that number is in the buyer’s head, you inherit the expectation whether it’s accurate or not. Winning the timeline query means being the clear, honest source on how long your work actually takes — so the model quotes you instead of filling the gap.

I spend a lot of time looking at how AI engines answer real buyer questions about real businesses, and timeline questions come up constantly. People don’t just ask what something costs. They ask how long it takes to get an appointment, how long the project runs, how fast they’ll see results, when they can expect delivery. These are not idle questions — they’re often the deciding factor. A buyer who needs something done by a deadline is screening on timeline first and everything else second.

The old playbook was to keep timelines vague on purpose. “It depends” felt safer than committing to a number you might miss. That instinct made sense when the buyer had to call you to get any answer at all. It makes a lot less sense when the buyer’s first move is to open ChatGPT or Perplexity and ask how long a thing like yours usually takes — and the model hands back a confident answer that has nothing to do with how you actually operate.

A vague timeline doesn’t protect you — it gets replaced

Here’s the shift that trips people up. When you refuse to give a timeline, you don’t create caution in the buyer’s mind. You create a blank the model is happy to fill. Ask an engine how long a kitchen remodel takes, how long until a new patient can be seen, how fast a loan closes, or how long results take, and it will answer. It pulls from whatever it can find: industry rules of thumb, forum threads, a transparent competitor’s process page, or just a reasonable-sounding inference.

The problem is that the number it lands on becomes the anchor. If the model says “most providers take six to eight weeks” and you routinely turn work around in three, you’ve lost an advantage you actually have — because the buyer walked in already believing the category is slow. If the model guesses low and you’re genuinely more thorough and therefore slower, you look like you’re underdelivering against an expectation you never set. Either way, the vacuum didn’t help you. It handed your most important selling point — or your most important caveat — to a guess.

My read is that timeline is one of the most under-managed signals in AI visibility right now. Businesses have slowly gotten comfortable talking about what they do. Far fewer have made *how long it takes* legible in a way a model can lift cleanly. That gap is an opportunity, because the engines want a concrete answer and will reward the source that gives them one.

You can be honest about “it depends” without being useless

The reason people hide behind “it depends” is usually that it genuinely does depend. Fair. But “it depends” as your only answer is a non-answer, and a model can’t cite a shrug. The fix isn’t to invent a false precision you can’t honor. It’s to explain the dependency in a structured way the engine can actually use.

That means naming the ranges and the variables that move them. Think about the honest version of your timeline answer:

- **A typical range.** “Most projects run four to six weeks” is infinitely more extractable than “every project is different.” It gives the model a defensible number and gives the buyer a real frame.
- **The fast path and the slow path.** What does the quickest realistic version look like, and what pushes a job to the long end? “Simple cases can be done in two weeks; complex ones with permitting involved can take three months” tells the model how to answer a nuanced question nuancedly.
- **The variables that drive it.** Scope, approvals, materials, your current backlog, how quickly the client gets you what you need. When you name the drivers, the model can match them to the specific buyer’s situation instead of guessing.
- **The parts you control versus the parts you don’t.** Being honest that a step depends on a third party (a lender, a supplier, a licensing body) is not a weakness. It’s exactly the kind of grounded detail that makes an answer read as trustworthy rather than salesy.

Notice that none of this requires you to promise a date you can’t hit. It requires you to describe your reality clearly. The engines are not looking for the shortest timeline — they’re looking for the clearest, most reliable one. Clarity is the thing you’re competing on here, not speed.

Structure the timeline so the model can lift it

Making a timeline legible to AI is partly a content problem and partly a formatting problem. A model extracts best from text that states the answer plainly and close to the question it answers. Buried in a paragraph of marketing prose, your timeline might get skipped. Stated directly under a clear heading, it gets picked up.

A few things that consistently help:

- **Answer the literal question.** If buyers type “how long does X take,” have a line on your site that says, in plain words, how long X takes. Mirror the phrasing people actually use.
- **Put it in an FAQ.** A question-and-answer block is the most extractable format there is, because it maps a real query to a real answer. “How long until I can get an appointment?” followed by a concrete response is close to ideal for how engines chunk and retrieve content.
- **Describe the process as stages with rough durations.** “Consultation (week one), design (weeks two to three), build (weeks four to six)” gives the model both a total and a breakdown it can use for follow-up questions.
- **Keep it current.** A timeline that reflected your capacity two years ago is now misinformation about your own business. If your lead times have changed, the page should change with them. Stale operational facts quietly erode the model’s trust in you as a source.

You also want these details corroborated beyond your own site. When your stated process timeline matches what shows up in your reviews, your booking page, and the way you actually talk about the work, the engines gain confidence that the number is real and not just a marketing claim. Consistency across sources is a signal in its own right.

Timeline is also an expectation-setting tool — use it

There’s a second reason to win this query that goes beyond visibility. The timeline the model gives the buyer becomes the expectation they carry into the relationship. If you own that number, you get to set the expectation correctly. If you cede it, you’re managing against a number a machine invented.

This cuts both ways, and the honest version is better business. If your work genuinely takes longer because you do it properly, saying so up front filters out the buyer who needed it yesterday and would have been unhappy anyway. If you’re genuinely fast, claiming that number plainly is a real advantage that’s currently invisible if it lives only in your head. Either way, the buyer who shows up has already accepted your timeline — which is a far better conversation to start from than one where they arrive expecting something you never promised.

I'd treat the timeline query the same way I treat the pricing query: as a question the buyer is going to get answered one way or another, where your only real choice is whether you're the source. The businesses that will win these queries are the ones willing to say the honest number out loud, explain what moves it, and keep it current.

Check what AI says about your timeline right now

Before you change anything, find out what the gap actually is. Open a few engines and ask the timeline questions your buyers ask — how long your kind of work takes, how fast someone can get started, when they can expect results. Ask it about your category, and then ask it about you specifically. See whether the answer is close to your reality, whether it's anchored to a competitor, or whether the model is just guessing.

That five-minute exercise usually makes the priority obvious. If the model is quoting a number that undersells you or overpromises on your behalf, that's the expectation walking in your door every day. Fixing it is mostly a matter of writing down the honest answer where the model can read it — and that's a lot more within your control than most of what shapes an AI recommendation.

Key takeaways

- Timeline questions — how long, how fast, when — are among the first things buyers bring to AI, and the model will answer them with or without you.
- A vague "it depends" doesn't create caution in the buyer; it creates a blank the model fills with a category average, a competitor's number, or a guess.
- You don't need false precision to win the query: a typical range, a fast path and slow path, and the variables that move the number are honest and extractable.
- Format matters — state the timeline plainly, mirror the words buyers use, and put it in an FAQ or a staged process breakdown the model can lift cleanly.
- The number the model gives becomes the buyer's expectation, so owning it lets you set that expectation correctly instead of managing against a machine's invention.
- Audit it in five minutes: ask the engines your buyers' timeline questions about your category and about you, and see whether the answer matches your reality.

Frequently asked questions

Won’t committing to a timeline publicly just set me up to get blamed when a job runs long?

That’s a real worry, but a stated range with named variables protects you better than silence does. If you say “most projects run four to six weeks, longer when permitting or custom materials are involved,” you’ve set an honest expectation and explained what pushes it. The buyer who hired you already accepted that. The bigger risk is staying vague and letting the model quote a number you never agreed to — that’s the expectation you actually get blamed against. My read is that an honest range you control beats a guess you don’t, almost every time.

My timelines genuinely vary a lot. How do I make that legible without lying?

Describe the variation instead of hiding behind it. Name the typical range, then explain the fast path and the slow path and what determines which one a given buyer is on — scope, approvals, your current backlog, how quickly they get you what you need. A model can work with “two weeks for simple cases, up to three months for complex ones” far better than it can work with “every situation is different.” You’re not inventing precision; you’re explaining your real dependencies in a structured way.

Where should the timeline actually live on my site?

Put the direct answer somewhere a model reads easily and close to the question it answers. An FAQ block is ideal because it maps a real query to a real answer. A process page that breaks the work into stages with rough durations is also strong, since it gives the model both a total and a breakdown for follow-up questions. Mirror the phrasing buyers use, and keep it current — a timeline that reflects last year’s capacity is now misinformation about your own business.

About the author. Scott Tischler is the Founder & Chairman of AIrecommend.ai and a practitioner-authority on AI search and Answer Engine Optimization. With 20+ years in marketing technology — including American Express, MetLife, and UBS — and executive and professional study at Wharton, Harvard, and Oxford, he helps businesses become the ones AI recommends.

© 2026 Scott Tischler · scotttischler.com
scott@airecommend.ai · AIrecommend.ai