

The AI Contractor Playbook: Win Bids Without Guessing

By Joe Giler

Table of Contents

1. Chapter 1: What AI Can and Cannot Estimate (Read This Before Chapter 2)
2. Chapter 2: Same-Day Quote Turnaround: The Two-Hour Workflow
3. Chapter 3: Building Your Real Cost Baseline From Past Jobs
4. Chapter 4: Handling Plans and PDFs Without Uploading Client Data You Shouldn't
5. Chapter 5: Turning a Messy Site Walkthrough Into a Structured Scope
6. Chapter 6: Follow-Up Sequences for Bids That Went Quiet
7. Chapter 7: Prompting for a Line-Item Takeoff You Can Actually Check
8. Chapter 8: Post-Job Debrief: Feeding Actuals Back Into Your Estimates
9. Chapter 9: Catching the Three Errors AI Makes on Material Counts
10. Chapter 10: When to Walk Away From a Bid (And the Signals AI Surfaces)
11. Chapter 11: Pricing Labor Hours When Your Crew Is the Variable
12. Chapter 12: Your 30-Day Rollout Without Blowing Up Current Jobs
13. Chapter 13: Writing Exclusions and Assumptions That Prevent Change-Order Fights
14. Chapter 8: Generating a Bid Document a Homeowner Understands in 90 Seconds
15. Conclusion
16. References and Further Reading

Chapter 1: What AI Can and Cannot Estimate (Read This Before Chapter 2)

If you skipped straight to the workflow chapters, come back here. This chapter is the one that keeps you from putting your name on a number that loses you \$18,000 on a kitchen remodel.

I've watched contractors get burned two ways with AI. The first is refusing to use it at all, then losing bids to the guy who turns a quote around in three hours while you're still "getting to it Sunday night." The second — worse — is trusting it with the part of estimating it is genuinely bad at, and finding out on framing day.

This chapter draws the line. It tells you exactly where a language model earns its keep on an estimate, exactly where it produces confident garbage, and how to tell the difference in real time without being an AI expert.

The one sentence that explains everything AI gets wrong about your numbers

A language model does not calculate. It predicts what a plausible answer looks like.

Read that again, because every failure mode in this book descends from it. When you ask ChatGPT, Claude, or Gemini "what does 2x6 pressure-treated cost per linear foot in Ohio," it is not looking anything up. It is generating the text that most resembles a correct answer based on patterns in its training data. Sometimes that text happens to be right. Sometimes it is a 2021 price presented in 2026 language. The model cannot tell you which, because it does not know the difference — it has no internal flag for "this figure is real" versus "this figure is shaped like a real one."

That is not a bug you can prompt your way out of. It is what the technology is. Every technique in this book is built around that constraint instead of pretending it away.

Here's the practical translation. When a model produces a *structure* — a list of scope items, a sequence of trades, a set of questions to ask the homeowner, a paragraph

explaining an exclusion — it is doing the thing it is actually good at: reproducing well-worn patterns of professional language. When it produces a *quantity* or a *price*, it is guessing with a straight face.

So: **let it write, let it organize, let it check. Never let it price.**

What AI is genuinely excellent at on a bid

I want to be specific here, because "AI helps with estimating" is the kind of vague claim that gets contractors into trouble. These are the tasks where I'd say the model outperforms most humans working alone under time pressure.

1. Scope extraction from messy inputs. A homeowner sends you four texts, a voicemail you transcribed, and a Pinterest board. An architect sends 22 pages where the good information is on sheets A2.1, A5.3, and one note in the general conditions. Paste all of it in and ask for a structured scope list, and you get a clean itemized breakdown in a minute. This is genuinely the highest-value use and it is not close. Most bad bids I've seen were not arithmetic errors. They were *missed scope* — the model is very good at not missing things that are stated in the text you gave it.

2. Finding what's missing. Ask it: "Based on this scope for a bathroom gut renovation, list every item a residential GC typically includes that is NOT mentioned here." You get back things like dump fees, permit runs, floor protection, temporary toilet arrangements, waterproofing membrane at the shower pan, blocking for grab bars, the electrician's second trip for trim-out. Every one of those is a real line you might have forgotten at 11 p.m. This is pattern completion, which is exactly what the technology does well.

3. Turning your numbers into a client-ready document. You have the pricing. You always had the pricing — it's in your head, your spreadsheet, or your supplier's quote. What eats your evening is formatting it into something that doesn't look like it was typed on a phone. A model will take your rough figures and produce a clean scope-of-work, inclusions, exclusions, assumptions, and payment schedule in the tone you specify. Ten minutes instead of ninety.

4. Writing the exclusions and assumptions section. This is underrated and it directly protects your money. Most disputes I've heard about trace back to something

that was neither included nor explicitly excluded. AI is excellent at generating a thorough exclusions list because exclusions are boilerplate patterns — and thoroughness, not creativity, is what makes them work.

5. Question generation before the site visit. Feed it the scope and ask for the 15 questions whose answers most change the price. You'll walk in with better questions than you would have improvised, and you'll leave with fewer unknowns to pad against.

6. Sanity-checking your own math. Not doing the math — checking it. "Here is my line-item estimate. Flag anything where the unit, quantity, or relationship between lines looks internally inconsistent." It will catch that you priced 340 SF of tile but 190 SF of backer board. That's a real save.

7. Reading spec language you don't want to read. Commercial specs, HOA architectural guidelines, municipal requirements. "Summarize every requirement in this document that costs the contractor money or time." Excellent use. Verify anything you'll rely on, but as a first pass it's fast and thorough.

8. Rewriting the same bid for a different audience. The GC wants line items. The homeowner wants a narrative and a number. The bank wants a draw schedule. Same underlying data, three documents, five minutes each.

What AI cannot estimate — the hard list

Now the part that matters more. If you take one page out of this book, take this one.

1. Current material prices. Not even close. Model training data has a cutoff, material markets move weekly, and pricing is regional to the point of absurdity. Lumber has swung wildly enough over the past several years that any number a model gives you is a coin flip on which year it's remembering. Copper, steel, and anything petroleum-derived are the same story. Ask a model for a price and you will get a confident number with no date and no region attached. It is worthless for your bid.

Where the price actually comes from: your supplier's current quote, your distributor's portal, your last three invoices for the same item, or a live phone call.

That's it. Those are the only four sources. If a number in your estimate didn't come from one of them, it's a guess wearing a suit.

2. Your labor rates and your crew's actual productivity. A model can tell you that hanging drywall takes roughly some hours per sheet. It cannot tell you that your two-man crew does 38 sheets a day on open walls and 19 in a 1920s house with plaster keys and no straight studs. Your productivity data is the single most valuable proprietary asset you own, and no model has ever seen it. Feed yours in and it becomes useful. Ask it to invent them and you're bidding somebody else's crew.

3. Site conditions. Access. Parking. Stairs. The 14-foot ceiling nobody mentioned. Whether the driveway will hold a dumpster. Whether the crawlspace is 30 inches or 18. Whether the neighbor is going to call the city on day two. Whether the "flat" backyard has a two-foot fall across the patio footprint. Every one of these moves a number and none of them are in any file you can paste.

4. What's behind the wall. No model knows about the knob-and-tube, the abandoned cast iron stack, the sistered joist somebody added in 1974 with sixteen-penny nails and hope, or the fact that the "load-bearing" wall isn't. You handle this with allowances and contingency language, which we'll cover — not with prediction.

5. Local code interpretation as your inspector applies it. A model can summarize published code. It cannot tell you that your county's inspector wants the strapping every 4 feet regardless of what the table says, or that your city has a local amendment on egress that nobody put online. Ask AI for a code question and you get a plausible-sounding answer that is frequently correct for the model code and frequently wrong for your jurisdiction. **Building code compliance is a legal matter with real liability. Verify with your local building department or a licensed design professional — never with a chatbot.**

6. What your competition is bidding. There is no data source for this. Anything a model says about "typical bid ranges in your market" is invented. Your read on your market comes from bids you won, bids you lost, and what the GC's PM lets slip.

7. Your overhead and your break-even. Your true overhead rate is a function of your truck payments, your insurance, your shop rent, your admin hours, and your billable-hour volume. It is arithmetic on your own books. A model that guesses

"contractors typically run 12–18% overhead" is giving you a shape, not your number. Do this one on paper, once, properly, and then use it for a year.

8. Anything involving structural load, engineering, or life safety. Beam sizing. Footing design. Retaining walls over four feet. Truss modification. Anything holding up anything. This is engineer territory, full stop, and it is not an AI question. **If your project touches structure, hire a licensed structural engineer. Nothing in this book substitutes for that.**

9. Insurance, licensing, lien law, and contract enforceability in your state. Lien deadlines and notice requirements vary by state and are unforgiving — miss one and you lose the right to collect. Contract terms that hold up in one state are void in another. **Have a construction attorney licensed in your state review your contract template once. It's a few hundred to low four figures and it is the cheapest insurance in this business. Do not use AI-generated contract language as-is.**

10. Tax treatment of anything. Deductibility, entity structure, sales tax on materials versus labor, worker classification. All state-specific, all consequential. **Talk to a CPA who works with contractors. AI will give you an answer that sounds like a CPA and isn't one.**

The confidence problem, and why it's the actual danger

Here's what makes this genuinely hazardous rather than merely annoying: **the model's tone does not change based on whether it knows.**

When it tells you "framing lumber for a 12x16 deck typically runs \$1,400–\$1,900 in most markets," it uses the exact same measured, professional voice it uses when telling you "always verify joist span against your local code table." One of those statements is a fabrication with a plausible shape. The other is solid advice. Nothing in the writing distinguishes them.

Humans have tells. When a lumber guy isn't sure, he says "let me check" or "somewhere around, don't hold me to it." A model has no such reflex. It generates fluent text because generating fluent text is the entire mechanism.

So you have to supply the skepticism yourself, structurally, every time. That's what the next section is for.

The Verify Test: four questions before any number goes in your bid

Any figure that appears in a document you send a client has to pass all four. Write these on a card and tape it inside your truck's visor if you have to.

1. Where did this number come from, specifically? Not "AI gave it to me." Name the invoice, the quote, the portal, the phone call, or the past job. If you cannot name a source, the number does not go in the bid. It becomes an allowance instead.

2. How old is it? A number from a February invoice in August is different from a number from last week. For volatile materials — lumber, copper, anything with fuel in the supply chain — I treat anything over 30 days as needing confirmation. For stable items, 90 days is fine. Date-stamp your price file so you can answer this in two seconds.

3. Is it for my region and my quantity? Prices differ between metros and rural areas, and unit prices change with volume. A per-unit price from a job three times the size of this one is not your price.

4. What happens to me if it's 25% wrong? This is the triage question. If a line is 25% off and it costs you \$60, ignore it and move on. If a line is 25% off and it costs you \$9,000, that line gets a real quote before you send anything. Most estimators spend equal effort on every line. Spend it where the money is.

The 80/20 of where your estimate actually goes wrong

On most residential jobs, a small handful of lines carry the overwhelming majority of your risk. Find them, and you've found where to spend your verification time.

Here's the drill. Once your line items exist, sort them by dollar value, descending. Draw a line where the running total hits about 80% of the job. Everything above that line gets a real, sourced, dated number — a supplier quote or a sub's written price.

Everything below the line can carry a reasonable historical figure with a modest cushion.

On a typical \$60,000 kitchen, the lines above that cut are usually cabinets, countertops, appliances if you're supplying, and the labor block. That's four to six lines out of forty. Get those right and a 20% error on "cabinet hardware" cannot hurt you.

What surprises people: **labor is almost always the biggest single number and the one most often estimated by feel.** If you do one thing after reading this chapter, start tracking actual hours per task per job. Not for accounting — for estimating. Six months of that data will make you the sharpest bidder in your market, and it is the one input AI can never fake for you.

The three ways AI actively hurts contractors

Let me be blunt about the failure modes I've seen, because the marketing around this stuff never mentions them.

Failure one: the invented unit price that survives to the contract. Somebody asks a chatbot for a square-foot cost, drops it into the estimate, sends the bid, wins the job — and finds out at buyout that the real number is 40% higher. Now you're eating it or having an ugly conversation. This is the most common and most expensive AI mistake in this trade, and it comes entirely from skipping the Verify Test.

Failure two: the confident scope omission. You ask for a scope list, get a beautiful one, and assume it's complete. It's complete *relative to what you pasted in*. If the drawing note about the structural header was on a sheet you didn't include, the model didn't skip it — it never saw it. AI cannot tell you about the thing you didn't give it. It also won't say "hey, I notice there's no sheet A3 here." You have to check your own inputs.

Failure three: bid documents that sound like a corporation. Send a homeowner a bid full of "leverage," "holistic solution," and "we are committed to exceeding expectations" and you will lose to the guy who wrote three plain paragraphs. Homeowners can smell generated text now — they see it all day. The tone

fix is easy and I'll give you the exact instruction in Chapter 2, but you have to know the problem exists.

How to actually get useful output: feed it your data

The single biggest upgrade in your AI estimating is embarrassingly simple. **Stop asking the model what things cost. Start telling it what things cost.**

Build one file. Call it `my-rates.md` or `pricing.txt` or whatever. Plain text. It contains:

- **Your labor rates** — burdened cost per hour per crew type, and your billing rate. Burdened means wage plus taxes, workers' comp, insurance, and any benefits, not the number on their paycheck.
- **Your productivity figures** — hours per unit for the 20 tasks you do most. Sheets of drywall per day. LF of baseboard per hour. SF of tile per day, per setter, broken out for floor versus wall versus mosaic.
- **Your last-known material prices, with dates** — pulled from actual invoices. Twenty to forty lines covers most of what you buy.
- **Your subs' pricing patterns** — what your electrician charges for a standard bath rough-in, what your plumber's day rate is, who marks up fixtures and who doesn't.
- **Your standard markup and overhead rate** — with the arithmetic that got you there, so you can re-derive it next year.
- **Your standard exclusions and assumptions** — the boilerplate you want in every bid.

That file takes an afternoon to build from your existing invoices and job records. From then on, every AI estimating session starts with pasting it in. Now the model is doing arithmetic and formatting *on your real numbers* instead of hallucinating industry averages. The output goes from "interesting but unusable" to "this is my bid, minus verification on the big lines."

Update it monthly. Fifteen minutes. Put a recurring reminder on the first of the month. This one habit is worth more than every other technique in this book