



PRACTICAL AI FOR RUBBER MANUFACTURERS & DISTRIBUTORS

Six prompts. Real numbers. Decisions that changed.

What a rubber shop and distributors actually do with a \$20 AI subscription — the prompts, the outputs, and where the answer surprised me.

Josh Petrawski

President, Southern Michigan Rubber · ISD Board Member



Every prompt you'll see today has the same four parts.

1

ROLE

Who the AI should be.

`"You are a CFO-level analyst for a $40M rubber products manufacturer."`

2

CONTEXT

What happened, the numbers, the constraints, the audience.

`"Attached: 24 months of order history by customer and SKU."`

3

TASK

The single specific ask, in numbered steps.

`"(1) Compute reorder intervals. (2) Flag gaps >30% over norm. (3) Rank by revenue at risk."`

4

FORMAT

Exactly what the deliverable looks like.

`"A one-page table, then three talking points per account. Show your math."`

Two habits: the first prompt is a starting point, not the answer — the value is in the follow-ups. And for anything with numbers, **make it show its work.**



HOW TO WATCH THE NEXT 40 MINUTES

Six use cases. Same five beats every time.

1

The question

What I was actually trying to decide

2

The prompt

Verbatim — steal it

3

The input

What I fed it (usually one ERP export)

4

The output

What came back, lightly redacted

5

The decision

What changed because of it

- 1 Sales order review — who is quietly slipping
- 2 TAM — sized down to the molds we already own
- 3 Margin deep-dive and the price / cost / mix walk

- 4 Payment terms — NET120 and the early-pay discount
- 5 Capital decision — three quotes for one machine
- 6 Executive coaching — the 360 my team can't give me

THE QUESTION

Which customers have quietly stretched their reorder cadence?

THE PROMPT

ROLE: You are a sales-operations analyst.

CONTEXT: Attached is 24 months of order-line history: order date, customer, part, revenue.

TASK: (1) Compute each customer's historical average reorder interval by part family. (2) Flag every customer whose most recent gap exceeds their own norm by 30%+. (3) Rank flagged customers by trailing-12-month revenue at risk. (4) For the top 10, draft three talking points for the outside rep.

FORMAT: An at-risk table ranked by exposure, then one rep-ready brief per account.

THE OUTPUT — AT-RISK ACCOUNTS (REDACTED)

Customer	Avg cadence	Last gap	Over norm	TTM revenue
Parco	28 d	61 d	+118%	\$412K
SKF	35 d	52 d	+49%	\$288K
Seal Company	21 d	33 d	+57%	\$196K
Great Lakes Industrial	45 d	63 d	+40%	\$174K
Hydril	14 d	20 d	+43%	\$151K
Schlumberger	30 d	41 d	+37%	\$98K
Revenue at risk				\$1.32M

Input: one ERP export. Time: about four minutes.

The rep briefs came back with it — “they buy shaft seals every 4 weeks, last order was 9 weeks ago, their new buyer started in May” — ready to forward.

THE FOLLOW-UP

Then I asked for the cheat sheet my rep actually needed before the call.

THE PROMPT

My outside sales rep is following up with Trelleborg Fort Wayne this week. Give me a customer cheat sheet so we have talking points: (1) Estimate the inventory they keep on the shelf, by product line, from their order quantities and cadence. (2) Show every product line that has dipped from its usual cadence, sorted by most at risk. (3) For each, the largest order drop on record and the last date the cadence was ever this low. (4) Three talking points for the rep.


THE OUTPUT — CUSTOMER CHEAT SHEET (ILLUSTRATIVE PART NUMBERS)

Product line	Usual cadence	Last order	Est. on shelf	Largest drop on record	Last time this low	Risk
Hydraulic rod seal · P/N 25001250-375B	4 wks	9 wks ago	~0 (-5 wks)	7 wks (Oct-2023)	Never	Highest
John Deere · RE517892 rod wiper	6 wks	11 wks ago	~0 (-5 wks)	12 wks (Apr-2024)	Apr-2024 (plant shutdown)	High
Caterpillar · 4J-0527 seal	8 wks	10 wks ago	~2 wks	14 wks (Jul-2022)	Jul-2022	Watch
Cameron · 2160021-02 seat seal	12 wks	9 wks ago	~3 wks	18 wks (Dec-2023)	—	On cadence
O-rings, FKM 2-2xx series	3 wks	2 wks ago	~1 wk	5 wks (Dec-2024)	—	Healthy

Talking point #1, straight from the sheet:

“By your own history you ran out of the 2-inch rod seals five weeks ago — and the cadence has never been this low. Either the line is down, or someone else is filling it.”

Sizing a market: the three words and the two methods.



TAM — Total Addressable Market

Every dollar spent on the product, by everyone, in the market you defined. The ocean.

SAM — Serviceable Addressable Market

The slice you could actually serve with your capabilities, certifications and geography. The water you can reach.

SOM — Serviceable Obtainable Market

The slice of the SAM you can realistically win in ~3 years given competition and sales capacity. The fish you'll actually catch.

TOP-DOWN

Start big, multiply down. All hydraulic components × share that is seals × share that is custom-molded × share in North America. Fast; uses published numbers; every multiplier is a guess.

BOTTOM-UP

Start small, add up. Count the buying companies × what each spends on seals per year. Slower; grounded in real buyers; you can miss some.

When both land near the same number, trust it. When they don't, the gap tells you which assumption is wrong.

THE QUESTION

How big is the market — and how much of it can we address with the tooling we already own?

THE PROMPT


DEEP-RESEARCH MODE — 15-20 minute run.

You are a McKinsey market analyst. I need a TAM analysis for [custom-molded hydraulic seals for North American mobile & industrial hydraulics OEMs].

Provide: top-down and bottom-up estimates; a TAM / SAM / SOM breakdown with dollars; 5-year CAGR; every assumption; a comparison against 3 published analyst figures; the 10 largest prospective accounts. Cite sources.

Context: [what we make, capabilities, certifications, geography].

THE OUTPUT — TAM / SAM / SOM (\$M, ILLUSTRATIVE)



Useful. Also exactly what a consultant would have charged \$40K for.

Cited, reconciled two ways, with the assumptions listed so you can argue with them. And it's still generic — any shop in the segment gets this same report. The next slide is where it stops being generic.

THE RUBBER-SHOP LAYER

Then I gave it our mold vault — and asked how much of Parker's seal business our tooling could already serve.

THE FOLLOW-UP PROMPT

"Here is the list of molds and product lines we currently own, and the molds we are building this year: [LIST]. Give me the specific addressable market for the molds we ALREADY own — revenue we could win with zero new tooling — versus what the new tooling unlocks. Which mold in our vault is sitting on the largest untapped market?"

Then: estimate what percentage of North American hydraulic seals Parker Hannifin sells, and what percentage of that business we currently get from Parker given the molds we own."

PARKER HANNIFIN — ADDRESSABLE WITH TOOLING WE OWN VS. WHAT WE INVOICE TODAY (ILLUSTRATIVE)

Mold / product family	Molds	Parker spend our tooling can serve	Our Parker rev. today	Untapped
Rod seals, 1–3" families	14	\$6.1M	\$0.55M	\$5.55M
Piston seals, FKM high-temp	6	\$3.4M	\$0.15M	\$3.25M
Wipers, polyurethane	9	\$2.8M	\$0.30M	\$2.50M
Back-up rings, PTFE	11	\$1.7M	\$0.10M	\$1.60M
Total, existing tooling	40	\$14.0M	\$1.10M	\$12.9M
<i>2026 new tooling (3 molds)</i>	3	<i>+\$2.3M unlocked</i>	—	—

Parker Hannifin, est. share of NA hydraulic seals

~22% ≈ \$90M

Of which our owned molds can serve \$14.0M

What we get from Parker today

\$1.1M ≈ 8%

of what our existing tooling could already serve.
The other 92% needs a sales call, not a mold.

THE QUESTION

Where is margin actually made and lost across the book?

THE PROMPT

ROLE: You are an FP&A analyst.

CONTEXT: Attached is P&L detail by product line and by customer for June 2025 and June 2026: revenue, cost, units.

TASK: (1) Build a margin league table by product line. (2) Trend margin and flag deterioration. (3) Flag every line below 25% GM. (4) Quantify concentration. (5) Name the three biggest margin leaks.

FORMAT: League table, five findings in plain English, three actions.

THE OUTPUT — MARGIN LEAGUE TABLE BY PRODUCT LINE, JUN-2026 (ILLUSTRATIVE)

Product line	Revenue (ann.)	GM %	GM \$	vs. Jun-2025
Oil & gas valve seals	\$1.3M	34%	\$442K	▲ rising
Custom molded (in-house)	\$1.8M	30%	\$540K	▼ falling
O-rings, FKM / NBR	\$2.1M	28%	\$588K	— flat
Hydraulic seals (OEM programs)	\$3.4M	20%	\$680K	▼ falling hard
Heavy-duty seal kits	\$2.6M	18%	\$468K	▼ falling hard
Lathe-cut / PTFE	\$0.9M	14%	\$126K	▼ falling
Total	\$12.1M	23.6%	\$2,856K	▼ -13.3 pts YoY

Finding #1 in plain English: the two biggest lines are the two worst margins, and both are falling hard.

That sentence took the controller two days to write last year. It took the model forty seconds — and it's only the setup for the question that matters: why?

THE MARGIN WALK — PRICE VS. COST VS. MIX

Margin fell 13.3 points. It was price: we gave away 8.4 points on identical parts.

GROSS MARGIN WALK, JUN-2025 → JUN-2026 (MARGIN POINTS; AXIS FROM 20%)



Follow-up: "Now break the cost bar into every driver, largest first, and make it tie to the reported P&L." → next slide.

Three effects, three different plays.

Price: hold or push. Cost: attack the buy, freight and scrap. Mix: steer the sales team. Before the walk, the room blamed freight and an inventory adjustment. Both were real. Neither was the story.

The rubber-shop layer, for the molders in the room: re-rank by margin per press-hour.

FKM: 46% margin, but a long cure on a 1,000-ton press → ~\$184 per press-hour. NBR: 31% margin, fast cycle, high cavitation → ~\$212 per press-hour. Margin percent lies in a molding shop. The press-hour is the scarce resource.



THE BREAKOUT — EVERY DRIVER, LARGEST FIRST, TIES TO THE P&L

Behind the three bars: eight drivers, one owner each, and a tie check that reads flat.

GROSS MARGIN WALK — REPORTED JUN-2025 TO JUN-2026 (ONE MONTH; FIGURES REBALANCED TO TIE — CONFIRM AGAINST THE P&L)

Driver (largest effect first)	Margin pts	Bar	Note
Reported Jun-2025 GM%	36.9%		per P&L
Price on identical parts	-8.4	Price	THE story: same part numbers, 8.4% less price, cost flat
Account adjustments (inventory)	-7.0	Cost	Jun-25 booked under-shipped — inflated prior year
Material cost (net)	+2.2	Cost	Modestly favorable
Inbound freight (vendor + brokerage)	-1.7	Cost	Brokerage back
Net customer shipping (markup)	+1.3	Cost	Modest
Production supplies	+0.9	Cost	Favorable
Direct labor	+0.7	Cost	Labor \$ +3.9% vs revenue +9.4% — next slide
Product mix	-1.3	Mix	More kit volume at lower margin
Sum of drivers	-13.3		
Reported Jun-2026 GM%	23.6%		per P&L
Tie check (2025 + drivers – 2026)	flat		walk reconciles to the financials

Why the tie check matters

A walk that doesn't reconcile is a story. This one starts at the reported number, ends at the reported number, and every point in between has an owner. Your CFO will trust it; your board will too.

The “Bar” column

Every driver rolls up to one of the three bars on the previous slide: Price -8.4, Cost net -3.6 (six lines, some favorable), Mix -1.3. Two views of the same month, and they agree.

ONE MORE FOLLOW-UP

Then I uploaded the payroll register: why did direct labor move?

THE PROMPT


Attached: the payroll registers for June 2025 and June 2026, direct labor only (names removed, employee IDs only).

Break the year-over-year change in direct labor dollars into: wage increases on existing staff, overtime, new hires, departures, and hours / shift changes. Tell me which pieces are structural and which are one-time.

THE OUTPUT — DIRECT LABOR BRIDGE, JUN-2025 → JUN-2026 (ILLUSTRATIVE)

Driver	\$ change	Type
Wage increases (merit / COLA)	+\$3.1K	Structural
Overtime	+\$2.4K	Fix it
New hires (2)	+\$4.8K	Structural
Departure (1)	-\$4.7K	One-time
Hours / shift changes	-\$0.3K	Noise
Total direct labor change	+\$5.3K (+3.9%)	
Revenue change, same month	+\$90K (+9.4%)	
Labor % of revenue: 14.0% → 13.3%	= +0.7 pts	ties to walk

WHERE THE +\$5.3K CAME FROM



The finding: labor was “favorable” in the margin walk only because revenue grew 9.4% while payroll grew 3.9%. Underneath, overtime is up \$2.4K — and 80% of it is on one shift.

Structural: raises and net headcount, +\$3.2K, expected. The fixable piece: overtime on second shift, where a departure wasn't backfilled for six weeks. Small dollars this month; a pattern if it repeats.

THE QUESTION

A \$1.2M customer wants NET120. Another offers NET10 for a 5% discount. What do those really cost?

THE PROMPT

ROLE: You are our treasury analyst.

CONTEXT: Customer does \$1.2M/yr with us at 30% gross margin. Our cost of capital is 9%.

SCENARIO A: NET30 → NET120. Carrying cost as an effective rate, as margin points, and the price increase that offsets it.

SCENARIO B: NET10 for a 5% early-pay discount. Annualized cost of the discount vs. our cost of capital, and the maximum discount that still makes sense — from NET30, NET60 and NET90 standard terms.

FORMAT: One-page memo, math shown, walk-away lines.

THE OUTPUT — COST-OF-TERMS MEMO (ILLUSTRATIVE NUMBERS)

Scenario A — NET30 → NET120	
Annual carrying cost ($\$1.2M \times 9\% \times 90/365$)	\$26,630
As % of revenue / as gross-margin points	2.2% / -2.2 pts
Price increase that fully offsets	+2.3%

Scenario B — 5% discount to pay NET10	from NET30	from NET60	from NET90
Days paid early	20	50	80
Annualized cost $(5/95) \times (365/\text{days})$	96%	38%	24%
Cash cost per year on \$1.2M	\$60,000	\$60,000	\$60,000
Max discount worth offering at 9% capital	0.5%	1.2%	2.0%

The longer your standard terms, the more an early-pay discount is worth.

From NET30, 5% is borrowing at 96%. From NET90, a 2% discount for NET10 is roughly a wash against a 9% line. Same discount, three different answers — which is why it's a table, not a rule of thumb.



THE DECISION IT DROVE

We stopped negotiating terms in the meeting and started walking in with the walk-away line.

96%

the effective annual rate on a 5%-for-20-days early-pay discount

+2.3%

the price adjustment that makes NET120 cost-neutral on a 30% margin account

0.5%

the largest early-pay discount that beats our line of credit

What changed

Two follow-up prompts became policy: a one-page terms matrix — what we accept, at what price adjustment, by customer size — and a counter-offer email template. My CFO runs this one himself now.

Why this matters for distributors especially

You live between OEM terms on the buy side and customer terms on the sell side. Ten seconds of math on both sides is the difference between a working-capital squeeze and a margin lever.

SAME QUESTION, NO CREDIT LINE

No debt? Your cost of capital isn't zero — it's what that cash would earn back in your business.

For a private company with no line of credit, the right rate is your return on invested capital — what a dollar earns when you put it into inventory, tooling or people. For most owner-operated shops that's 30–35%. The big companies in the room can ask finance for their WACC; everyone else, use 30–35% and move on.

SCENARIO A — NET30 → NET120, \$1.2M ACCOUNT, 30% GM

	9% (line of credit)	30%	35%
Annual carrying cost	\$26,630	\$88,767	\$103,562
As gross-margin points	-2.2	-7.4	-8.6
Price increase that offsets	+2.3%	+8.0%	+9.4%

SCENARIO B — MAX EARLY-PAY DISCOUNT WORTH OFFERING FOR NET10

Standard terms	at 9%	at 30%	at 35%
From NET30 (20 days early)	0.5%	1.6%	1.9%
From NET60 (50 days early)	1.2%	4.1%	4.8%
From NET90 (80 days early)	2.0%	6.6%	7.7%

NET120 goes from a 2-point problem to an 8-point problem.

For a no-debt company, letting a \$1.2M customer stretch to NET120 costs as much margin as a bad pricing year. The offsetting price increase isn't 2.3% — it's 8%. Most owners would never quote it, which means most owners are giving it away.

And the early-pay answer flips.

At 9%, a 5% discount for NET10 is a bad deal from anywhere. At 30–35%, the same 5% from NET90 is a bargain — the annualized cost is 24%, below what the cash earns in your business. Same discount, opposite answer, and the only input that changed is what your money is worth to you.

THE QUESTION

A key account now requires part marking on 2.6M seals a year. Three quotes. Which machine?

THE PROMPT

[3 vendor quotes attached]

A key account now requires us to mark the parts in their kits – about 2.6M rings a year. Here are three quotes for a part-marking machine. Help me decide which one is best.

Attached is my direct labor costs.

That was the whole prompt. Everything in the table — the cost per 1,000, the labor split, the capacity math — the model decided I needed on its own. The honed version of this prompt is in the library; this is the one I actually typed.


THE OUTPUT — COST PER 1,000 RINGS (REAL QUOTES, VENDORS ANONYMIZED)

	A — Inkjet	B — CO ₂ laser	C — UV laser
Quote total	\$23,582	\$36,029	\$50,845
Installed (quote + fixturing)	\$41,582	\$54,029	\$68,845
Effective rings / hour	1,800	2,500	2,500
Labor per 1,000	\$12.22	\$8.80	\$8.80
Capital per 1,000	\$2.28	\$2.97	\$3.78
Consumables per 1,000	\$3.00	\$0.15	\$0.15
Service per 1,000	\$0.82	\$0.83	\$0.77
TOTAL per 1,000 rings	\$18.33	\$12.75	\$13.50
Annual cost at 2.6M rings	\$47,650	\$33,153	\$35,103

THE DECISION IT DROVE

The machine price was rounding. The whole decision was rings per hour.

COST PER 1,000 RINGS, CO₂ LASER OPTION — BY COMPONENT



Labor is 3× the machine. Fixturing honesty moved installed cost from \$44K to \$55–71K — and at 2.6M rings that is ~0.1¢ per ring. The cell is rounding. The cost is handling.

1 — It told me what to go test.

The vendor demo stopped being a sales pitch and became a stopwatch test: time the real load-mark-unload cycle. That number decides everything; the price sheet decides nothing.

2 — It labeled every number FIRM, ESTIMATE, or MINE TO CONFIRM.

The quotes were firm. The rings-per-hour was the vendor's claim. The labor rate and volume were mine. Knowing which numbers were guesses told me exactly where the model could be wrong — and where to spend my diligence.

3 — The reason to buy a second machine isn't volume. It's downtime.

One laser does 2.6M rings in about half a shift — volume wouldn't justify a second until ~5M. But once the contract requires the mark, a broken laser means we can't ship. The second unit is insurance, not expansion, and that changes when we buy it.

THE QUESTION

What are my real strengths and weaknesses as a leader — according to the evidence?

THE PROMPT

THE HABIT: one ongoing chat per direct report. Every recorded 1:1 (with consent) goes in. Every debrief happens there. For months.

THE PROMPT, word for word, in every chat: "We have been working together for a while in these 1:1s. Based on everything you know about my 1:1 chats, what are my biggest strengths and weaknesses as a leader and CEO?"

EIGHT RELATIONSHIPS. ONE QUESTION. ONE CONVERGING ANSWER.

1:1 chat 1

1:1 chat 2

1:1 chat 3

1:1 chat 4

1:1 chat 5

1:1 chat 6

1:1 chat 7

1:1 chat 8



WHAT EIGHT INDEPENDENT REVIEWERS AGREED ON

Strengths, every single time:

Structural thinking — I separate the role from the person and the system from the symptom. I update fast on evidence, including against myself. I develop people by handing them the pen.

Weaknesses, every single time:

The player reflex — I fill the silence, I do the work myself “for time's sake.” Dates I set for myself slip. I'm soft in the room and hot after it. I delay the personally uncomfortable decision and call it judgment.

“Your strengths and weaknesses are the same trait seen from two sides.”



IT ALSO COACHES THE MOMENT BEFORE I WALK INTO IT

An employee brought me an accusation from a heated argument. My instinct was to referee it.

THE PROMPT

[1:1 transcript pasted]

What do I do? How do I address this? I don't want to deny race exists, but I can't have people calling others racists.

That's the actual prompt. No role, no format, no framework. It doesn't have to be pretty when it matters.

The advice

Don't adjudicate the he-said / she-said — you weren't there, and you'll lose either way. Set the standard for how concerns get raised in your company, for everyone, and hold it. You're not deciding who was right. You're deciding how this works here.

THE STANDARD I SET

"If you believe someone is treating you differently because of who you are, you bring that to me privately and I will take it dead seriously.

What you don't do is raise it in the middle of a heated argument with a co-worker.

That's not raising a concern — that's weaponizing one."

Not a ruling on who was right. A rule for how it works here.

WHAT EIGHT CHATS SAID ABOUT ME

It does not flatter. That is the point.

“Fourteen minutes on laser wavelengths in a leadership meeting. ‘Have him reach out to me and I’ll show him how to connect the dots’ — said in front of the controller whose job that was. Your attention is the single point of failure across three companies, and you keep spending it on the interesting problem instead of the important one.”

“Soft in the room, hot after it. You said ‘I love how we root-caused it’ in the meeting, then called the same exchange ‘babying’ two days later and wanted shame. The standard you suppress in the room returns later with interest — and your people only ever see the suppressed version.”

“You were ready to downgrade a man’s trajectory on honesty — and the facts cut the other way at every single step. You did to him precisely what you coach him not to do: led with a conclusion you couldn’t back.”

“Your patience is real leadership — and other people pay for it. You ran a fifty-five-day experiment on one manager while his team absorbed the mess. You spent five weeks deciding a call you had already made. The structural thinking that stops you blaming a person before you’ve decomposed the P&L is the same instinct that lets you defer a people decision into a process fix and call it patience. Humaneness, curdling into delay.”

“The most load-bearing seat in the company has been empty for a year. Four candidates reached the finish line; four times you didn’t close. Meanwhile you redrew the org chart around a chair nobody sits in and counted the re-org as progress. You delay the personally uncomfortable piece and call it judgment.”

THE ONE I KEEP

“You build people by handing them the pen, and it’s your best skill. You never touched the mouse through fifteen minutes of him fumbling. You refused to build the board and built the tool that let him build it. You gave him authorship of a solution you’d seeded weeks earlier and never claimed the seed.

The result is measurable: a man who said ‘you can’t automate a place like this’ in July was designing his own systems and running his own adoption meetings by August.”

Both lists are true. That’s what owning a company is.



WHAT TO DO MONDAY

One prompt. One real problem. One owner. Then scale.

Don't build a strategy. Run one prompt.

Pick the one with the most obvious dollar payoff — payment terms or reorder cadence are the fastest wins. Have one person run it on real data this week. Review the output together. Win once.

Then build a shared prompt library inside your company so the wins compound instead of living on one laptop. That library is the asset; the tool is a subscription.

If you remember one prompt, remember the one that writes your prompts.

"I want help with [fuzzy goal]. Before you answer, ask me the 7 questions you'd need answered to give a genuinely useful analysis. Then write the prompt I should have given you — and run it."

It removes the blank page. Every other prompt you saw today started as a version of this one.

THE GIVE — AND THE ASK

This deck is yours. Everyone. Free.

Every prompt you saw today, verbatim, plus the speaker notes. Get with Lori and she'll send it — or find me after.

THE FULL LIBRARY — 34 PROMPTS, 43 PAGES

Bring one more person from your company to the 2027 Summit on the Summit Pro Package, and the library comes with them.

1

The Executive AI Prompt Library

Every prompt you saw today plus 28 more — the Rubber-Shop Layers, the deep-research strategy suite, and the ones I didn't have time for.

2

Early-bird registration for everyone your company sends in 2027

For most companies this alone covers the package.

3

One complimentary ad in ISD inBrief

The monthly newsletter, in front of every member company.

I'm asking as a favor — I just joined the board and want to make a dent. One more seat gives someone on your team this toolkit, and it funds more sessions like this one. If this was worth your hour, I'll do Part 2 next year with whatever I've built by then.

To get either one: get with Lori.

QUESTIONS

Appendix

Josh Petrawski · President, Elastotech Southwest · ISD Board Member

jpetrawski@elastotechsouthwest.com · elastotechsouthwest.com · For the deck or the package: see Lori





BEFORE YOU UPLOAD ANYTHING

Three rules. Non-negotiable.

1

Business-tier account. Never a personal free account.

Business and enterprise tiers don't train on your data by default (ChatGPT Business is ~\$25/user/month). Free accounts may. The tool is fine — the account type is the decision.

2

Know what never goes in — even on a business tier.

PII, anything under NDA, signed contracts, employee records. Nothing you saw today needed a customer's real name. Anonymize before you paste. 35% of what employees paste into AI tools is sensitive (Cyberhaven, 2025).

3

Every output is a first draft from a fast, occasionally wrong analyst.

It will state a wrong number with total confidence. A human signs off on anything that leaves the building. For math, make it show its work — every prompt today did.