

Using Your Chairside AI Skills

A quick-reference card for each skill — what it does, what to say, what to give it, and what you'll get back.

How to use this guide

Find the skill you want. Use the "Say this" line to kick it off — or just describe what you need in plain English. Give Claude the information listed under "Give it this," and you'll get the output described under "You get."

Privacy reminder before every session: strip out patient names, birthdates, and account numbers before pasting anything. See the safety one-pager for the full picture.

Money-Leak Finder

What it's for: Scans your bank or card statement for silent recurring costs — duplicate subscriptions, zombie SaaS tools, overpriced print contracts, bloated supply bills — and hands you a prioritized savings list with ready-to-send renegotiation scripts.

Say this:

"Run a money-leak analysis on my practice. I'll upload my statement."

Give it this:

*Your bank or card statement exported as a CSV (60–90 days is ideal to catch quarterly charges). **Before uploading, delete the full account and card-number column** — this tool never needs it. Totals by vendor are all that's*

required; no patient names or account numbers are involved. (A bank statement is financial data, not patient health info — HIPAA does not apply. The concern is card safety, not PHI.)

You get: A ranked HTML report showing your top leak categories, the estimated dollar impact of each, and a plain-English action per item — Cancel, Keep, or Renegotiate — with the actual script attached to each renegotiation item. Most practices find \$5,000–\$30,000 in fixable leaks on the first pass. Every finding is labeled with a confidence level; unclear charges are flagged "verify" rather than asserted.

Case Closer

What it's for: Turns a de-identified treatment plan into a designed, patient-ready case presentation — plain-language clinical explanation, a tiered financial menu with phasing and payment options, and a four-touch follow-up sequence. The doctor reviews it before any patient sees it.

Say this:

"I have a treatment plan I'd like to turn into a patient presentation. I'll paste the de-identified details."

Give it this:

A de-identified description of the case (strip every identifier before pasting — no patient name, date of birth, address, phone, member ID, or claim number):

- Procedures recommended and their fees, each marked urgent or recommended
- A plain-language "why" for each area
- Age band (e.g., "late 50s"), insurance type, remaining annual max, and rough coverage %
- The patient's goal in their own words — this is the north star
- Financing options your practice offers

Turn off model training (Settings → Privacy) before pasting any clinical information.

You get: A branded HTML handout the patient can decide from, with three tiers (Comprehensive / Phased / Essential-now), computed patient portions, and a deferred-interest disclosure on any promo financing. Plus a four-touch follow-up sequence (Day 0/3/7/14) and a short internal coaching note for your team — who leads the conversation, which tier to open with, and the one cost-of-waiting line.

Fee Auditor

What it's for: Analyzes your fee schedule against directional regional benchmarks, quantifies the profit leverage of a fee increase (at your actual overhead), flags the codes most likely underpriced, and produces a re-key CSV ready for import into your practice management software.

Say this:

"Audit my fee schedule against what other practices in my area are charging."

Give it this:

Your fee schedule exported from Dentrix, Eaglesoft, Curve, or similar — CDT code, description, and current fee, one code per row (CSV, XLSX, or a plain paste). No patient data is involved; this is your own pricing. Also helpful: your ZIP code, your overhead percentage (default 65% if unknown), your annual collections (for a real-dollar projection), the year your fees were last updated, and whether you are in-network with PPOs.

You get: A clean HTML report with a profit-leverage snapshot ("a 10% fee increase is a 29% profit increase at 65% overhead"), a flagged-codes table with directional benchmark ranges, a recommended increase schedule with two tiers (baseline and targeted), the math shown in plain English, talking points for your team and front desk, and the re-key CSV. Benchmarks are stamped "directional — confirm vs. Patterson/FAIR Health before adjusting insurance contracts."

EOB Decoder

What it's for: Translates a dense, confusing Explanation of Benefits into plain English — what the insurance paid on each line, what they denied and why, the accurate patient balance (computed correctly: deductible first, then coverage %, then annual-max cap), and a reusable coverage-book entry so the next patient on this plan skips the back-and-forth.

Say this:

"Decode this EOB — tell me what happened and what I should do."

Give it this:

*The text of a **de-identified** EOB — paste or upload a redacted image or PDF. Strip every identifier before sharing: patient name, member ID, subscriber ID, claim number, date of birth, address. Keep procedure codes (D2750, etc.), tooth numbers, billed/allowed/paid amounts, coverage percentages, annual max, deductible status, and remark or reason codes — none of those are PHI under Safe Harbor.*

Turn off model training (Settings → Privacy) before pasting any EOB text. The skill runs an automated PHI detector before reading anything back — if an identifier slips through, it stops and asks you to re-share.

You get: A line-by-line plain-English decode for your front desk (what each code means, what was paid, what was denied and why, and the concrete next step — resubmit, appeal, collect from patient, or write off). A coverage-book entry the team can file and reuse for every future patient on this plan. And, optionally, a patient-facing estimate sheet in plain language — two lines: pay in full today and a monthly estimate, no insurance jargon.

Cyber Self-Audit

What it's for: Walks you through an 18-question guided security checkup across 8 risk domains — backups, cloud restore, cyber insurance, MFA, IT partnership quality, OS hygiene, phishing training, and vendor BAAs — then delivers a scored risk-tier badge, a per-domain traffic-light dashboard, a prioritized roadmap, and two copy-paste call-sheets for your IT provider and insurer.

Say this:

"Run the cyber self-audit for my practice."

Give it this:

Answers to 18 yes/no/idk questions that Claude will walk you through one domain at a time. No patient data is required or accepted — every question is about your systems and policies. "I don't know" is always a valid answer and is scored as a gap, never skipped. Plan for 10–15 minutes. Optionally, upload your cyber insurance declarations page (policy number blacked out) and the skill will flag ransomware sublimits or exclusions.

You get: A printable HTML scorecard with your overall risk tier (Critical / High / Moderate / Low-watch), a color-coded domain dashboard, red-flag callouts, and a three-horizon roadmap (This Week / This Month / This Quarter) — every critical item paired with a cheap, specific first step and an owner (You / IT / Insurer). Plus two verbatim call-sheets: exactly what to ask your IT provider and exactly what to ask your insurer.

Morning Huddle

What it's for: Turns tomorrow's schedule into a one-page printable huddle brief in under a minute — scheduled production vs. daily goal, high-value case routing, open-chair fill plan, and a team talking point. No more reactive days.

Say this:

"Set up my morning huddle for today. I'll paste the schedule."

Give it this:

*Today's schedule in a simple pipe-delimited format — one appointment per line: Time | Provider/Op | Initials | Appointment Type | Scheduled \$ | Tx-Plan Status. **Use initials only — no full names, no DOB, no phone or insurance ID.** Also give Claude your daily production goal (or monthly goal plus working days) and your high-value case threshold (default \$3,000). Turn off model training before pasting.*

You get: A one-page printable huddle sheet: scheduled production vs. goal and the dollar gap; production by provider; a high-value money list with routing flags (high-value cases in hygiene slots get an explicit "route to TC" flag); unscheduled treatment plans for patients coming in today; an open-chair fill plan; and a rotating team talking point. Claude echoes the appointment count and total back to you for a 5-second confirm before generating — if the total is wrong, fix the paste before the day starts.

Hiring Kit

What it's for: Turns your role description and practice values into a complete, EEO-safe hiring packet — job post, phone-screen questions, structured interview scorecard with anchored ratings, a paid working interview with a rubric, reference-check questions, a 30/60/90-day onboarding checklist, and a one-page decision summary.

Say this:

"Help me hire a dental assistant — write a job post and give me interview questions."

Give it this:

Tell Claude: the role (hygienist / treatment coordinator / front desk / dental assistant / associate dentist, or describe a hybrid), your 3–5 practice values, your compensation range or use "[owner to set]" for any figure you are not ready to commit to, your must-haves (licenses, experience), nice-to-haves,

and any relevant context (solo vs. group, why the last person left). No patient data is ever needed — applicant information is employment data, not PHI. If you paste a résumé, a candidate-privacy reminder fires before any scoring.

You get: A single Artifact with all seven sections: an EEO-clean job post with a real EEO statement, structured phone-screen questions, a 1–5 anchored scorecard with behavioral questions per competency, a paid working interview task with a rubric, five targeted reference-check questions with red-flag signals, a 30/60/90-day onboarding checklist with role-specific milestones, and a one-page decision summary you fill in after interviews. Compensation slots show your confirmed figures or "[owner to set]" — the skill never invents a number.

Vision Coach

What it's for: Helps you work through a difficult people situation on your team — an underperformer, a culture issue, a missed expectation — by first diagnosing what is actually happening under the surface (through a Socratic reframe), then handing you a word-for-word conversation script, a 15-minute meeting agenda, anticipated pushback with ready replies, and an offer to role-play the opening before you walk in.

Say this:

"I'm struggling with a situation on my team. Help me think it through and prepare for the conversation."

Give it this:

Describe the situation in behaviors, not labels ("she missed the morning huddle twice this week," not "she doesn't care"). Claude will ask you six questions one at a time: what's happening, what success looks like, what you've already tried, who's involved (role and first name only — no last names or identifying details), how confident you are that you are not part of the problem, and what happens if nothing changes in 90 days. No patient data is ever involved.

Note: if the situation involves potential termination, discrimination, harassment, wage/hour disputes, or a team member in crisis, the skill stops and refers you to an employment attorney or HR professional. It does not script employment consequences.

You get: A printable one-pager you carry into the room: the plain-language reframe of what is really happening, the manage-vs-lead distinction applied to your situation, a word-for-word opening script with specific behavioral asks, a tight 15-minute meeting agenda, 3–4 anticipated pushback lines with word-for-word replies, and a concrete follow-up checkpoint. If the reframe surfaces a vision vacuum, it also builds a practice vision statement. After delivering the one-pager, Claude offers to play the team member so you can rehearse the opening line.

All eight skills are live. Install any of them via Settings → Customize → Skills → Upload a skill (one ZIP at a time), then toggle it on. Full steps in the Quick-Start guide.