AIPG

SESSION A WORKSHOP KIT

Chairside AI *Day-1 Implementation Guide.*

Six workflows. End-to-end. Ready Monday.
Everything you need after you walk out of Session A.

EVENT

Bulletproof Dental CE Summit

August 9, 2026

AUTHOR

Blake McClellan · All In Practice Growth
aipg.co

RULE THREADED THROUGH EVERYTHING

Segment-level, not patient-level. Human review before anything ships.

CONTENTS

Table of Contents

—

HOW TO USE THIS GUIDE

- You just walked out of Session A. You have one workflow on paper.
- This guide gives you six workflows — each end-to-end, ready to install Monday.
- **Chapter 1** is the foundation — set up your Claude Desktop Project first. Every other chapter assumes you did this.
- **Chapter 2** is the mental model — the 4-move framework, so you can build workflows that aren't in this guide.
- **Chapters 3–6** are workflows in increasing operational depth.
- Every chapter has: what you're building, why it matters, prerequisites, step-by-step setup, the exact prompt, expected output, troubleshooting, variations, success criteria.
- **PHI rule threaded through everything:** segment-level, not patient-level. Every setup step names its PHI check.
- **This guide is PHI-free zone.** For PHI-adjacent workflows (clinical documentation, chart-level data), see Zach Hanson's 12:30 session materials.

—

CHAPTER

CHAPTER 1 — Setting Up Your First Claude Desktop Project

What you're building

A Claude Desktop **Project** that acts as your Practice Memory — a persistent workspace with your practice context loaded so every conversation starts pre-informed.

Why it matters

Without a Project, every new conversation with Claude starts from zero. You re-explain who you are, what your practice does, what your voice sounds like, and what your rules are — every single time. Five minutes wasted per conversation, and every conversation drifts because you never say the rules the same way twice.

With a Project, you load all that context once. Every conversation inside the Project inherits it. Same voice, same rules, same PHI guardrails — every time. This is the single highest-leverage setup step in the entire kit.

Prerequisites

- **Software:** Claude Desktop (or claude.ai in a browser). Download from `claude.ai/download`.
- **Account:** Free tier works for Chapters 1, 2, and 6. Chapters 3, 4, 5 require Pro tier (~$20/mo) because they use Files + Python + Skills.
- **PHI check:** none for this chapter — you're not pasting any data yet.

Step-by-step setup

Step 1 — Install Claude Desktop

- **Action:** Go to `claude.ai/download`. Choose Mac or Windows. Run the installer.

- **Expected visual:** After install, a Claude icon appears in your Applications folder (Mac) or Start Menu (Windows). Open it. You'll see a chat interface with a sidebar on the left and a "New chat" button at the top.
- **Gotcha:** If you already had a claude.ai browser tab open, log out there before installing. Multiple simultaneous logins occasionally trip the sync.

Step 2 — Sign in

- **Action:** Click "Sign in" in Claude Desktop. Use email + password, or "Continue with Google."
- **Expected visual:** After sign-in, you land on a blank chat screen. Sidebar on the left shows "Chats" and "Projects" sections.
- **Gotcha:** Use your practice email, not a personal one. This account will hold your practice data — you want it separate from your personal Claude use.

Step 3 — Turn off model training

- **Action:** Click your profile initials in the top-right → **Settings** → **Privacy**. Toggle OFF "Help improve Claude."
- **Expected visual:** The toggle turns gray. A confirmation message appears: "Your conversations will not be used to train Claude."
- **Gotcha:** This is a per-account setting, not a per-Project setting. Do it once, done.
- **PHI check:** This is a low-stakes privacy hygiene step — but it's the first move because if you skip it and later paste something you shouldn't, you can't take it back.

Step 4 — Create your Project

- **Action:** In the left sidebar, click the **"+ Project"** button (or "New Project" — the label varies by version).
- **Expected visual:** A dialog opens asking for a Project name and (optionally) a description.
- **Gotcha:** Some browsers show "Create a Project" as a full-page flow; the Desktop app shows it as a dialog. Either works.

Step 5 — Name the Project

- **Action:** Name it `[Your Practice Name] Ops`. Examples:
  - "Summit Family Dental Ops"
  - "Hilltop Perio Ops"

    - “Kensington Endodontics Ops”

- **Expected visual:** The new Project appears in your sidebar. Click it to open.
- **Gotcha:** Don’t put version numbers (“v1”, “v2”) in the name. You’ll iterate on the instructions, not the Project.

Step 6 — Paste the Practice Memory template into Custom Instructions

- **Action:** In the Project view, click **”Set custom instructions”** (right panel, top). Paste the template below. Edit the bracketed placeholders.

The exact prompt (Practice Memory template):

```
You are the AI assistant for [Practice Name], a [general/specialty] dental
practice
in [City, State]. Our patient base is [describe: e.g., family + geriatric,
suburban, PPO-heavy]. Our brand voice is [warm, plain-spoken, no dental
jargon
patients don't use, never salesy].

WORKFLOWS YOU HELP WITH:
- Patient communication drafts (welcome, follow-up, reactivation, review
response)
- Marketing copy review and rewrite
- Hiring rubrics and interview scorecards
- Vendor and tool evaluation
- Fee schedule and financial analysis
- SOP drafting

NON-NEGOTIABLE RULES:
- Never include or ask for patient names, dates of birth, or clinical de-
tails.
- Work at the segment level ("patients overdue 6+ months"), not patient
level.
- Sound like a dentist talking to a patient they know — never corporate.
- Always draft; never send. I review every output before it leaves the
practice.
- Never use these banned words in patient materials: best, guaranteed,
  painless, permanent.

OUTPUT PREFERENCES:
- Default length: as short as possible while complete.
- Default tone: warm, plain-spoken, direct.
- When asked for options, give me 3 versions — not one.
- When I paste a spreadsheet or CSV, run it in the code sandbox
  and show me the code. Never guess numbers.
```

- **Expected visual:** The custom instructions box now shows your template. Click **Save.** A green confirmation appears.
- **Gotcha:** The template limit in some tiers is ~8,000 characters. This template is ~1,200 characters — plenty of room to add more rules over time.
- **PHI check:** the non-negotiable rules block is your PHI guardrail — you're baking the "no patient names, segment-level only" rule into the Project itself so it applies to every conversation.

Step 7 — (Optional) Attach files to the Project

- **Action:** In the Project view, look for **"Add content"** or **"Project files"** (right panel, below custom instructions). Upload:
  - Your brand guide PDF (colors, fonts, voice examples)
  - Your fee schedule CSV (once you've stripped patient-specific columns)
  - Your letterhead SVG or PNG
  - Any style-guide document your team uses
- **Expected visual:** Files appear as tiles in the Project. Each has a filename and a size.
- **Gotcha:** File limit is 20 in most tiers. Don't upload every patient handout you've ever written — pick 3–5 canonical examples.
- **PHI check:** review every file before upload. Redact patient names, account numbers, DOBs. If you're unsure, don't upload it.

WHAT GOOD OUTPUT LOOKS LIKE

To test your Project: click "New chat" **inside the Project** (not the sidebar's global "New chat"). Type:

> *"Introduce yourself and remind me of the rules you follow."*

Claude should respond with something like:

> *"I'm the AI assistant for Summit Family Dental in Tulsa, OK. Family + geriatric patient base, PPO-heavy. Warm, plain-spoken voice.*
>
> *I help with: patient communication drafts, marketing copy, hiring rubrics, vendor evaluation, fee schedule analysis, SOP drafting.*
>
> *My non-negotiable rules: no PHI, segment-level only, dentist voice not corporate, always draft never send, and I never use best/guaranteed/painless/permanent in patient materials.*
>
> *Where do you want to start?"*

If you got that (or something close), the Project is live.

If Claude ignored the template and said "I'm Claude, an AI assistant by Anthropic" — the instructions didn't save. Go back to Step 6 and click Save again.

TROUBLESHOOTING

1. **Claude introduces itself generically, not as my assistant.** The custom instructions didn't save. Go back to Settings → Custom Instructions → re-paste → Save.
2. **My files won't upload.** File may be too big (>30MB) or in a rejected format. Try converting PDFs to text and CSVs to standard comma-separated (not tab-separated).
3. **Custom instructions are truncated in Claude's replies.** Your template is over the character limit. Trim the "workflows you help with" list or move examples to project files.
4. **The Project button is missing from the sidebar.** You're on a legacy tier. Upgrade to Free-current or Pro, or use claude.ai in a browser instead.
5. **Claude asks for a patient's name.** Your rules didn't stick. Re-paste the template and explicitly say in your first message: "Follow the non-negotiable rules I set in the Project instructions."
6. **Every new chat feels different.** You're starting chats outside the Project (using the global "New chat" button). Always click into the Project first, then click "New chat" from within it.
7. **The sidebar shows two Projects with similar names.** Delete the older one. Rename the newer one to be crystal clear.
8. **Custom instructions dropdown doesn't show a "Save" button.** Scroll down inside the dialog — Save button is at the bottom, easy to miss.

Variations

- **Variation 1 — Multi-location practice:** Create one Project per location. Same template, different practice name and city. Use identical rules so voice stays consistent.
- **Variation 2 — Solo doctor personal use:** Create a second Project called "[Your Name] Personal" for non-practice tasks (personal writing, health tracking, side interests). Keep them separate so the practice Project stays focused.
- **Variation 3 — Practice + specialty offshoot:** If you're a GP with an implant surgery focus, create a second Project called "[Practice] Implant Cases" with a modified voice and additional financing rules.

Success criteria (End of Week 1)

- [ ] Claude Desktop installed and signed in with your practice email
- [ ] Model training toggled OFF
- [ ] One Project created, named `[Practice] Ops`
- [ ] Practice Memory template pasted into custom instructions and saved
- [ ] At least one file uploaded to the Project (brand guide, fee schedule, or letterhead)
- [ ] Test conversation completed — Claude correctly identifies itself as your practice's assistant
- [ ] You've told at least one team member "this is our AI assistant" and shown them how to open a new chat inside the Project

—

CHAPTER

CHAPTER 2 — The 4-Move Workflow Framework

What you're building

A repeatable mental model for turning any repetitive practice task into a workflow. This chapter is theory, not click-through. Come back to it whenever you're staring at a task and thinking "should this be a workflow?"

Why it matters

A prompt is a tool. A workflow is a system. Random prompts produce random results. Workflows produce reliable results. Every dental practice runs on clinical protocols — same steps, every time, auditable. Your AI should run on the same principle. The 4-move framework is how you get there.

The 4 moves

Every workflow — no matter the pattern (Communication, Quality Review, Decision Framework), no matter the tool (Claude, ChatGPT, Grok) — has these four moves:

Move 1 — TRIGGER

What starts the workflow? Something has to happen for the workflow to begin. Examples:

- A new lead comes in from a Meta ad → Trigger for the New Patient Welcome workflow
- A patient completes a treatment plan discussion → Trigger for the Treatment Plan Follow-Up workflow
- A review posts on Google → Trigger for the Review Response workflow
- Monday at 8 AM → Trigger for the Weekly Huddle workflow

The Trigger is small and specific. "Patients need welcomes" is not a trigger. "A new lead just submitted the intake form" is a trigger.

Move 2 — AI STEP

What does AI do? The heavy lifting. Examples:

- **Draft** — write an email, a text, a script, a letter
- **Analyze** — process a spreadsheet, find patterns, flag outliers
- **Compare** — put two vendors side-by-side, surface the trade-offs
- **Summarize** — reduce a long document to key points

Notice: the AI Step never says “send.” It says “draft” or “analyze” or “compare.” Sending is a human decision — see Move 3.

Move 3 — HUMAN REVIEW

You approve or tweak. This is the non-negotiable gate. Every workflow has one. Examples:

- Review a drafted email before it goes out — one-minute check
- Approve a flagged vendor list before you contact anyone
- Sign off on a fee-increase schedule before your PMS admin re-keys it
- Confirm a huddle sheet before the team reads it

The Human Review is not a rubber stamp. It’s your judgment applied to AI output. This is what keeps the AI a lever, not a liability. Skip this and eventually something goes out that shouldn’t.

Move 4 — OUTCOME

What ships? The concrete deliverable. Examples:

- Email arrives in patient’s inbox
- Text lands on patient’s phone
- Vendor contract gets signed
- Huddle sheet is on the printer

The Outcome is the point. Every workflow exists because you wanted this thing to happen. If you can’t name the outcome in one sentence, the workflow isn’t real.

Worked example — Review Response Workflow

Let's build one all the way through.

- **Move 1 — Trigger:** Every Monday at 8 AM, I check Google, Yelp, and Healthgrades for reviews posted in the last 7 days.
- **Move 2 — AI Step:** I paste each review into Claude Desktop (inside the Project). I ask Claude to draft a response in my practice voice — grateful for positive, professional for negative, no defensive language, no clinical details.
- **Move 3 — Human Review:** I read each draft. I edit if needed. I never send Claude's first draft without at least one word change — even if it's just to swap a comma. That mental gate keeps me from ever coasting.
- **Move 4 — Outcome:** By 8:20 AM Monday, all reviews from the past week have a public response.

The workflow above takes 20 minutes on Monday. Without it, review responses sit for 3 weeks and negative ones snowball. That's what "workflow beats prompt" looks like in a real practice.

The three patterns (recap from Session A)

Every workflow you'll ever build fits one of three patterns:

1. **Communication** — every message that repeats 2+ times per week (welcome, follow-up, review response, reactivation, insurance verification).
2. **Quality Review** — anything you personally review (ad copy, hiring materials, vendor comparisons, SOPs, financial projections).
3. **Decision Framework** — comparing options, surfacing red flags (should I hire this associate, switch PMS software, buy this equipment).

If your workflow doesn't fit one of these three, you're overengineering. Simplify.

Common mistakes

1. **Skipping Move 3 (Human Review).** The most tempting shortcut. Don't. If you're not reviewing, you're not accountable.
2. **Combining Move 2 with Move 4.** "AI writes AND sends" is not a workflow — it's a runaway train.

3. **Trigger too vague.** “When I feel like it” is not a trigger. Pick a specific event, a specific time, or a specific state change.
4. **Outcome too vague.** “Something happens” is not an outcome. Name the ship: email sent, doc printed, decision made.
5. **Building all workflows at once.** Build one. Prove it. Then build the second. Never plan a “workflow system” before you’ve run one workflow three times.

Success criteria (End of Week 1)

- [ ] I have written down my first workflow using the 4-move framework
- [ ] I have identified which of the 3 patterns it fits (Communication, Quality Review, Decision Framework)
- [ ] I can name my Trigger, AI Step, Human Review, and Outcome in one sentence each
- [ ] I’ve run the workflow at least once in Claude Desktop
- [ ] I’ve refined the AI Step prompt at least once (based on what the first output missed)

—

CHAPTER

CHAPTER 3 — Files + Python: The Fee Schedule Audit

Deep-dive skill: Chairside Fee Auditor. This chapter walks the exact live demo from Session A end-to-end.

What you're building

A fee schedule audit that computes real profit-leverage math from your actual fee schedule and produces a re-key CSV your PMS admin can import Monday morning.

Why it matters

Most dental practices haven't raised fees in 2–4 years. Why? Because it feels risky, feels salesy, and requires math nobody wants to do. Meanwhile, at 65% overhead, every dollar of fee revenue adds $0.35 of profit — the leverage is enormous and it's costing you weekly.

The Fee Auditor Skill turns that math from scary into obvious. Five minutes of work, decision-ready output, front-desk scripts included. This chapter walks it end-to-end.

Prerequisites

- Claude Desktop installed, signed in, Project set up (Chapter 1).
- **Pro tier required** — you need Files, Python, and Skills, all Pro-only.
- Chairside Fee Auditor Skill installed (see Chapter 5 for Skill install steps; you can install Fee Auditor the same way).
- Your fee schedule exported as CSV from your PMS (Dentrix, Eaglesoft, Curve, Open Dental, etc.).
- **PHI check:** a fee schedule is your pricing data, not patient data. Not PHI. Safe to upload.

Step-by-step setup

Step 1 — Install the Chairside Fee Auditor Skill

- **Action:** Download `chairside-fee-auditor.zip` from `aipg.co/chairside`. Claude Desktop → **Settings** → **Capabilities** → turn ON "Code execution and file creation." Then **Settings** → **Customize** → **Skills** → **Upload skill** → drop the ZIP in → toggle ON.
- **Expected visual:** Fee Auditor appears in your Skills list with a green ON toggle.
- **Gotcha:** If you didn't turn on Code Execution first, the Skill will install but silently fail when it tries to run Python.

Step 2 — Export your fee schedule from your PMS

- **Action:** Every major PMS supports fee schedule export. Quick paths:
  - **Dentrix:** Reports → Practice Analysis → Fee Schedule Export
  - **Eaglesoft:** Utilities → Fee Schedules → Export
  - **Curve:** Reports → Fee Schedules → Download CSV
  - **Open Dental:** Setup → Fees → Export
- **Expected visual:** A CSV file with three columns at minimum: CDT code, description, current fee.
- **Gotcha:** Some PMSs export XLSX by default. Save-as CSV before uploading. Claude reads XLSX too but CSV parses more reliably.
- **PHI check:** open the CSV before uploading. Make sure there are no patient names, no provider-specific identifiers, no confidential vendor markup. If any of those are present, delete those columns.

Step 3 — Open a new chat inside your Practice Ops Project

- **Action:** Click into your Project (from Chapter 1). Click "New chat."
- **Expected visual:** Fresh chat window opens. Your Project name shows at the top.
- **Gotcha:** Don't open the chat from the global sidebar — always from inside the Project so the Practice Memory template applies.

Step 4 — Paste the exact prompt and attach the CSV

The exact prompt (edit the bracketed values):

```
Audit my fee schedule.

Practice: [Your Practice Name], ZIP [Your ZIP].
Overhead: [65]%.
Annual collections: $[1,800,000].
Fees last updated: [2022].
In-network: [yes/no].

I've attached the fee schedule as a CSV.
Flag underpriced codes and give me a tiered increase recommendation.
```

- **Action:** Paste the prompt. Edit the five bracketed values (practice name, ZIP, overhead, collections, last updated year, in-network status). Click the paperclip icon to attach your fee schedule CSV.
- **Expected visual:** File chip appears below the text input showing the CSV filename. Send the message.
- **Gotcha:** If you don't know your overhead %, use 65% (national average). If you don't know your last-updated year, skip that line — the analysis still runs, you just won't get the stale-fee gap calculation.

Step 5 — Confirm the parse count

- **Expected visual:** Claude responds first with a privacy note ("this is your pricing, not PHI"), then runs Python (you'll see a code block), then echoes back: **"I parsed 25 codes — confirm this matches your schedule?"**
- **Action:** If the parse count matches your CSV, reply "yes." If not, reply "no — the CSV has X codes" and Claude will retry.
- **Gotcha:** The parse count is the Skill's built-in error-catcher. Don't skip this. If it says 25 and your CSV has 25, say yes. If it says 25 and your CSV has 400, something in the CSV is malformed — usually a delimiter issue (semicolons instead of commas).

Step 6 — Read the rendered Artifact

- **Expected visual:** An Artifact panel opens on the right side. Sections include:
  - **Headline** — "A 10% fee increase = a 29% profit increase for your practice."
  - **Profit-leverage snapshot** — your overhead, 5% increase result, 10% increase result, dollar lift, stale-fee gap since your last update.
  - **Flagged codes table** — CDT | Description | Current Fee | Directional Range | Gap% | Confidence | Basis.

- **Recommended increase schedule** — CDT | Description | Current Fee | Recommended Fee | % Increase | Tier.
- **The math, shown** — leverage formula in plain English for your numbers.
- **Front-desk talking points** — team briefing, patient response scripts.
- **Re-key CSV** — downloadable, ready for PMS import.
- **Caveats footer** — the directional-only stamp and disclaimers.

- **Action:** Read the headline. Read the flagged codes. Note the front-desk scripts.

Step 7 — Export the Artifact + download the re-key CSV

- **Action:** In the Artifact panel, click the export dropdown. Choose PDF for the full report. Click the CSV download link for the re-key file.
- **Expected visual:** Two files download: the PDF report and the re-key CSV.
- **Gotcha:** The re-key CSV is not the whole fee schedule — it's only the codes with recommended changes. Your PMS admin imports it as an update, not a wholesale replacement.

WHAT GOOD OUTPUT LOOKS LIKE

Here's a segment-level example of what the Artifact should contain (numbers are illustrative — yours will differ based on your inputs):

> ***Headline: A 10% fee increase = a 29% profit increase for Summit Family Dental.***
>
> ***Profit-leverage snapshot:***
>
> - *Your overhead: 65%*
> - *+5% fees → +14% profit / +$90,000 top-line*
> - *+10% fees → +29% profit / +$180,000 top-line*
> - *Stale-fee gap since 2022: 12.3% behind CPI*
>
> ***Flagged codes (7 of 25 flagged as likely underpriced — directional):***
>
> *| CDT | Description | Current Fee | Directional Range | Gap% | Confidence | Basis | |—-|—-|—-|—-|—-|—-|—-| | D2740 | Crown - porcelain/ceramic | $1,150 | $1,400–1,700 | 22% low | High | Cost-tier median, 2026 | | D2790 | Crown - full cast high noble | $1,050 | $1,300–1,550 | 24% low | High | Cost-tier median, 2026 | | D4341 | Perio scaling 4+ teeth per quad | $210 | $260–320 | 24% low | High | Cost-tier median, 2026 | | … | … | … | … | … | … | … |*
>
> ***Front-desk script — "Did prices go up?"***
>
> *"Yes, we updated our fee schedule this year. We hadn't adjusted in [X] years, so we brought fees in line with what it costs us to deliver the level of care you've come to expect. Your insurance benefits and coverage percentages haven't changed. Is there anything about your specific treatment I can clarify?"*

TROUBLESHOOTING

1. **Claude just wrote a paragraph with numbers, no code block.** Fee Auditor didn't fire — the Skill isn't installed or Code Execution is off. Go to Settings → Skills → confirm Fee Auditor is toggled ON. Settings → Capabilities → confirm Code Execution is ON.
2. **The parse count is way off (e.g., 25 rows read as 1 row).** CSV delimiter issue. Open the CSV in Excel, save as "CSV (Comma delimited)" — not "CSV UTF-8" or "CSV Macintosh." Re-upload.
3. **Claude asks for a patient name.** Your Project rules didn't stick. Say: "Follow the segment-level rule — never ask for patient names." Also: this shouldn't happen with Fee Auditor since it never touches patient data. If it does, the Skill malfunctioned — reinstall.
4. **The recommended fees look absurdly high or low.** Fee Auditor uses directional ranges by ZIP tier — a rural ZIP will suggest lower fees than urban. Confirm your ZIP is right in the prompt. If it still feels off, cross-check with the Patterson Dental Fee Survey.
5. **No codes flagged.** Either your fees are already at market, or your CSV is missing description text (the flagging engine keys off descriptions). Ensure your CSV has three columns: CDT code, description, fee.
6. **The Artifact panel doesn't render — I just see a wall of text.** Claude occasionally responds inline instead of building an Artifact. Reply: "Render this as an HTML Artifact." That usually forces the panel.
7. **The re-key CSV won't import to my PMS.** Different PMSs expect different column headers. Open the re-key CSV, match the column headers to what your PMS expects (usually `CDTCode`, `Description`, `FeeAmount`), save, retry.
8. **Fees flagged that shouldn't be (e.g., D2740 flagged in a low-cost-tier ZIP).** The directional ranges are national medians by cost tier. Your local market may run higher or lower. Fee Auditor stamps every recommendation "directional" for exactly this reason — cross-reference before you change contracts.

Variations

- **Variation 1 — Annual audit cadence:** Run Fee Auditor every January. Same prompt, updated collections. This becomes your annual review protocol.
- **Variation 2 — PPO vs. UCR comparison:** Run Fee Auditor twice — once with `in-network: yes` and once with `in-network: no`. Compare the recommended increases. This shows you what your contracts are costing.
- **Variation 3 — Multi-provider practice:** Run once per provider if your fee schedule varies by provider. Otherwise run once per location.

Success criteria (End of Week 1)

- [ ] Fee Auditor Skill installed and toggled ON
- [ ] Fee schedule CSV exported from PMS
- [ ] Audit run at least once with your real numbers
- [ ] Artifact reviewed and exported to PDF
- [ ] Re-key CSV downloaded
- [ ] Front-desk scripts read out loud to your team lead
- [ ] Decision made: implement in Q3 or Q4

References cited from the Skill

- Skill definition: `products/chairside-ai/chairside-fee-auditor/SKILL.md`
- Python engine: `products/chairside-ai/chairside-fee-auditor/scripts/fee_audit.py`
- Demo CSV: `products/chairside-ai/chairside-fee-auditor/assets/demo-fee-schedule.csv`
- Reference documents (bundled with Skill):
  - `references/profit-leverage-math.md`
  - `references/benchmark-ranges.md`
  - `references/fee-raise-psychology.md`
  - `references/annual-increase-playbook.md`
  - `references/talking-points-templates.md`

—

CHAPTER

CHAPTER 4 — MCP Setup for a Dental Practice

What you're building

Model Context Protocol (MCP) connections between Claude Desktop and the five most-used tools in a modern dental practice: Google Drive, Gmail, Google Calendar, Slack, Notion. Once connected, Claude can read from and write to these tools directly — no copy-paste.

Why it matters

Your practice has a dozen tools that don't talk to each other. Your team is the nervous system copying data between them. MCP makes AI the nervous system instead. Concrete wins:

- Ask Claude "read my brand guide from Drive and rewrite this welcome letter to match" — no upload, no re-explaining.
- Ask Claude "draft a Slack message to my team for the huddle" — Claude sends the draft to Slack for you to approve.
- Ask Claude "what's on my calendar tomorrow?" — Claude checks and tells you.

Prerequisites

- Claude Desktop installed, signed in, Project set up (Chapter 1).
- **Pro tier required** for MCP connectors.
- Google Workspace account for Drive, Gmail, Calendar (personal Gmail works too, but Workspace has better admin controls).
- Slack workspace admin permissions (or ask your admin).
- Notion workspace with at least one page.

- **PHI check:** MCP gives Claude access to real systems. Before connecting, audit what's in each system:
  - **Google Drive:** any patient-identifying files? If yes, move them to a folder Claude can't access, or use Google Drive folder-level sharing.
  - **Gmail:** patient emails likely contain PHI. Do NOT connect Gmail if you use it for patient communication — use a separate marketing-only Gmail.
  - **Calendar:** patient appointments in Calendar? If names are in event titles, disconnect Calendar or rename events with initials only.
  - **Slack:** internal channels only. Do not connect Slack if patient-identifying messages ever go through it.
  - **Notion:** SOPs and marketing docs only. Do not connect Notion if you store patient records there.

Step-by-step setup — do these one at a time, not all at once

Connection 1 — Google Drive

- **Action:** Claude Desktop → **Settings** (gear icon, top-right) → **Connectors** → find **Google Drive** → click **Connect**.
- **Expected visual:** Browser opens to Google OAuth. Sign in with your practice Google account.
- **OAuth scopes to approve:** "See your Drive files" and "See metadata about your Drive files." Approve.
- **Return to Claude:** connector shows "Connected" with a green checkmark.
- **Test:** in a chat inside your Project, type: `List the folders in my Drive.` Claude should reply with 5–10 folder names.
- **Gotcha:** If Claude says "I don't have access to Drive," refresh Claude Desktop. If still broken, revoke the connector, sign out of Google in your browser, reconnect.
- **PHI check for this connector:** review your Drive folders before connecting. If patient-identifying files exist, move them to a folder shared only with your PMS integration, not your general Drive.

Connection 2 — Gmail

- **Action:** Same path: Settings → Connectors → Gmail → Connect.

- **OAuth scopes:** “Read, compose, send email” and “Modify labels.”
- **Test:** `List my last 5 unread emails.` Claude should return 5 subject lines.
- **Gotcha:** Do NOT connect a Gmail account that receives patient emails. Use your marketing Gmail (`marketing@yourpractice.com`) not your `office@yourpractice.com`.
- **PHI check:** Gmail is the highest-PHI-risk connector. Skip it unless you have a dedicated marketing inbox. Even then, tell Claude in your Project: “Only draft emails to the marketing inbox. Never read or draft emails that mention patient names.”

Connection 3 — Google Calendar

- **Action:** Settings → Connectors → Google Calendar → Connect.
- **OAuth scopes:** “See events on all your calendars” and “Change events.”
- **Test:** `What’s on my calendar tomorrow?` Claude should list your appointments.
- **Gotcha:** If your event titles include patient names, do NOT connect Calendar. Or connect a separate marketing calendar you use only for team meetings, vendor calls, and planning.
- **PHI check:** appointments with patient names in the title are PHI. Do the audit before connecting.

Connection 4 — Slack

- **Action:** Settings → Connectors → Slack → Connect.
- **OAuth scopes:** “Read messages,” “Send messages,” “Read channels.”
- **Test:** `List the channels I’m in.` Claude should return your channel list.
- **Gotcha:** Slack app permissions are workspace-wide by default. Restrict Claude’s access to specific channels using Slack’s App Manager.
- **PHI check:** if patient-identifying messages ever go through Slack, do not connect. If they don’t (SOPs, team ops, vendor coordination only), connect and set a rule in your Project: “Only send Slack messages to channels I explicitly name.”

Connection 5 — Notion

- **Action:** Settings → Connectors → Notion → Connect.
- **OAuth scopes:** “Read content,” “Insert content.”
- **Test:** `List my Notion pages.` Claude should return top-level pages.
- **Gotcha:** Notion access is granted per-workspace. Make sure you’re connecting the right workspace if you have multiple.

- **PHI check:** if Notion holds patient records (many DSOs use Notion for team-facing patient trackers), do not connect. If it's SOPs, marketing docs, hiring materials only, connect.

WHAT GOOD OUTPUT LOOKS LIKE

Once all 5 are connected, try this prompt inside your Project chat:

> *"Draft me an email to my team announcing next Monday's huddle. Reference the Q3 practice goals I have in Notion. Send the draft to Slack #ops as a message for me to approve."*

Claude should:

1. Read your Q3 goals from Notion.
2. Draft the email.
3. Post the draft to your Slack #ops channel as a message.
4. Confirm all three actions with a summary.

If you got all four steps, MCP is live. Welcome to your new nervous system.

TROUBLESHOOTING

1. **OAuth flow loops back to sign-in without connecting.** Sign out of ALL Google accounts in your browser, then retry. Chrome multi-account is the #1 cause.
2. **Connector shows "Connected" but Claude says it can't access.** Refresh Claude Desktop (Cmd+R on Mac). Still broken? Revoke the connector, wait 60 seconds, reconnect.
3. **Slack messages aren't sending.** Check Slack's App Manager — the Claude app may be limited to specific channels. Add the channels you want Claude to post to.
4. **Notion pages missing.** Notion's OAuth asks you to explicitly grant access to each page/workspace. Re-run the OAuth flow and expand access.
5. **Calendar events won't create.** Check the calendar you're granting access to — you may have granted read-only access to your primary calendar but Claude tried to write to a different one.
6. **"Rate limit exceeded" on Gmail.** You're asking Claude to read too many emails at once. Ask for the last 10, not the last 100.
7. **Claude reads a PHI-adjacent Drive folder it shouldn't.** Restrict Drive folder sharing at the Drive level, not the Claude level. Move the folder into a "Claude no-access" parent and reshare only what Claude should see.
8. **Multiple connectors all break at once.** Your Claude session token expired. Sign out of Claude Desktop, sign back in, connectors auto-refresh.

Variations

- **Variation 1 — Solo practice:** Skip Slack (no team channel to post in) and Notion (no external workspace). Drive + Gmail + Calendar is enough.
- **Variation 2 — Multi-location DSO:** Connect one Slack workspace but restrict Claude to a specific channel per location.
- **Variation 3 — Practice management via SaaS:** If you use Aloe, Curve, or another cloud PMS with an API, ask your PMS vendor if they have an MCP server. Several are shipping in 2026.

Success criteria (End of Week 1)

- [ ] Google Drive connected and tested
- [ ] (Optional) Gmail connected — only if you have a marketing-only inbox
- [ ] (Optional) Calendar connected — only if event titles are PHI-free
- [ ] Slack connected and restricted to non-patient channels
- [ ] Notion connected (if you use it for SOPs)
- [ ] Test prompt run successfully: “Read [file from Drive], draft a message, post to Slack for approval”

—

CHAPTER

CHAPTER 5 — The Chairside AI Kit Skill Install

Walkthrough example: Chairside Case Closer. Same install path applies to all 8 Skills in the Kit.

What you're building

An installed Claude Skill that gives Claude a specialty. Case Closer specifically turns a de-identified treatment plan into a patient-ready case presentation with tiered financing, phasing, and a follow-up sequence.

Why it matters

Skills are the difference between "Claude can help me draft this" and "Claude has a specialty in this." A prompt is you explaining what you want every time. A Skill is Claude already knowing.

Case Closer specifically solves a common revenue leak: the $5,000 case that dies in the hygiene chair because nobody connected the treatment to what the patient came in wanting. Case Closer produces a doctor-reviewed, on-brand handout the patient can actually decide from.

Prerequisites

- Claude Desktop installed, signed in, Project set up (Chapter 1).
- **Pro tier required** — Skills need Code Execution + File Creation, both Pro-only.
- Chairside Case Closer ZIP downloaded from `aipg.co/chairside`.
- Your practice brand config (see Step 3).
- **PHI check:** Case Closer only works on de-identified treatment plans. You must strip all patient identifiers before pasting. The Skill has a privacy gate that will halt if it detects identifiers — but you should not rely on it. Do the redaction first.

Step-by-step setup

Step 1 — Enable Code Execution

- **Action:** Claude Desktop → **Settings** → **Capabilities** → toggle ON "Code execution and file creation."
- **Expected visual:** Toggle turns green. Message says "You can now upload Skills and run code."
- **Gotcha:** This is a one-time setting per account. If you skip it and install a Skill, the Skill silently fails when it tries to run Python.

Step 2 — Install the Skill

- **Action: Settings** → **Customize** → **Skills** → **Upload skill**. Drop `chairside-case-closer.zip` into the upload area, or click to browse.
- **Expected visual:** Case Closer appears in the Skills list. Toggle it ON.
- **Gotcha:** If you get "invalid Skill package," you downloaded a corrupted ZIP or extracted it first. Re-download from `aipg.co/chairside` and upload the ZIP file directly — do not unzip first.

Step 3 — Set up your practice brand config

Case Closer needs to know your practice colors, fonts, and logo for the patient-ready handout.

- **Action:** Open a new chat in your Project. Ask: `Show me the brand config template from Case Closer.`
- **Expected visual:** Claude returns a JSON template with fields: practice name, primary color, secondary color, Google fonts, phone, booking URL, financing partners, practice voice.
- **Fill it in:**

```
{
  “practice_name”: “Summit Family Dental”,
  “primary_color”: “#2C5F7C”,
  “secondary_color”: “#F4A261”,
  “font_heading”: “Playfair Display”,
  “font_body”: “Inter”,
  “font_fallback”: “Georgia”,
  “phone”: “918-555-0100”,
  “booking_url”: “summitfamilydental.com/book”,
  “financing_partners”: [
    { “name”: “CareCredit”, “term”: “6 mo no-interest promo”, “long_term”:
“24 mo APR” },
    { “name”: “Sunbit”, “term”: “3 mo no-interest promo”, “long_term”: “12
mo APR” }
  ],
  “voice”: “Warm, plain-spoken, direct. Never salesy. Anchor to the
patient’s stated goal — never upsell past it.”
}
```

- **Action:** Save this JSON to a file called `brand-config.json`. Upload it to your Project as a project file.
- **Expected visual:** File appears in your Project files sidebar.
- **Gotcha:** Google Font names have to be exact. “Inter” not “inter” or “Intra.” Case Closer uses these to render the handout — a typo means the fallback font kicks in.
- **PHI check:** brand config is your practice’s marketing data. Zero PHI. Safe to upload and share.

Step 4 — Test the install with a demo plan

- **Action:** Open a new chat in your Project. Type:

```
Turn a de-identified treatment plan into a case presentation using Case
Closer.
```

- **Expected visual:** Claude should respond with the Case Closer privacy gate:

> *"Before we start — Case Closer works on a de-identified plan. HIPAA's Safe Harbor applies when a plan is stripped of identifiers.*
>
> *Before you paste, confirm:*
>
> 1. *You've stripped patient name, DOB, address, phone/email, member/subscriber ID, claim number.*
> 2. *You've kept: tooth numbers, procedures + fees, age band, insurance type + remaining max + rough coverage %, and the patient's stated goal.*
> 3. *Model training is off in Settings → Privacy.*
>
> *Ready when you are. Paste the plan (or use the demo from `assets/demo-treatment-plan.txt`).*"

- **Gotcha:** If Claude just started writing a generic case presentation without showing the privacy gate, the Skill didn't install. Go back to Step 2 and confirm the toggle is ON.

Step 5 — Run the demo plan

- **Action:** Paste this demo plan (this is the bundled demo from Case Closer, not real patient data):

```
Age band: late 50s.
Insurance: PPO, ~$1,200 remaining annual max, ~50% major/80% basic cover-
age.
Patient's stated goal: "I want to be able to chew on both sides again
without pain."

Recommended treatment:
- Urgent:
  - Tooth #14: extraction + immediate implant + custom abutment + crown.
Total: $5,800.
  - Tooth #19: buildup + crown. Total: $1,850.
- Recommended (not urgent):
  - Tooth #3: replace old amalgam with composite. Total: $320.
  - Tooth #30: bruxism-related wear, night guard recommended. Total: $650.
```

- **Expected visual:** Claude runs Python (`financial_menu.py`), then renders an Artifact with:
  - The patient's goal restated as the anchor.
  - Three tiers: Comprehensive / Phased / Essential-now.

- Each tier: subtotal, insurance offset, patient portion, monthly estimate range (labeled "estimate, subject to approval").
- A phased breakdown showing what happens in phase 1 vs. phase 2.
- A 4-touch follow-up sequence (Day 0, 3, 7, 14).
- A short internal coaching note for the team.
- Closing line: "Your doctor reviews this before any patient sees it; the monthly figures are estimates, not a bill or a guarantee of approval."

Step 6 — Export the handout

- **Action:** In the Artifact panel, click Export → PDF.
- **Expected visual:** A branded PDF handout downloads with your practice logo, colors, and fonts.
- **Gotcha:** If the logo doesn't appear, you didn't upload it as an inline SVG or base64 PNG. Re-upload the logo as an SVG for best results.

WHAT GOOD OUTPUT LOOKS LIKE

The exported Case Closer handout should:

- Open with the patient's goal in plain language ("You told us you want to be able to chew on both sides without pain — here's how we get you there").
- Present three tiers side-by-side with total, insurance estimate, patient portion, monthly estimate range.
- Never lead with the monthly payment — always show total first, then break down.
- Include a phased option for patients who can't do everything at once.
- Contain a follow-up sequence (Day 0, 3, 7, 14) that anchors every message to the patient's stated goal.
- NOT contain any of the banned words: best, guaranteed, painless, permanent.
- End with the doctor-review disclaimer.

TROUBLESHOOTING

1. **The privacy gate doesn't fire.** Skill isn't installed. Go to Settings → Skills → check Case Closer is toggled ON.
2. **Claude skipped the code sandbox and just wrote numbers.** Code Execution is off. Go to Settings → Capabilities → toggle it ON.
3. **The handout looks like plain HTML, no colors or fonts.** Your brand config wasn't loaded. Re-upload `brand-config.json` to the Project.
4. **The logo doesn't appear in the PDF.** Logo needs to be an inline SVG or base64-embedded PNG. Convert your logo, re-upload.
5. **Case Closer used a patient name.** You didn't fully de-identify. Case Closer's privacy gate should have caught this — reinstall the Skill and re-paste with the identifier stripped.
6. **Monthly estimates seem too aggressive.** Case Closer uses illustrative ranges based on the financing partners you named. If your CareCredit rate differs from the default, update `brand-config.json` with your actual term and re-run.
7. **The 4-touch sequence sounds generic.** Add more detail to the voice field in `brand-config.json`. "Warm, plain-spoken, direct" is a starting point — add specifics like "Doctor sometimes references her golden retriever in messages" to inject character.
8. **Claude used one of the banned words.** Report a bug at `aipg.co/chairside/support` — the Skill has a built-in ban list. If it slipped through, the Skill needs updating.

Variations

- **Variation 1 — Specialty practice:** Case Closer works for perio, prosth, ortho, endo. Change the voice field to match your specialty style.
- **Variation 2 — Multi-doctor practice:** Create a brand config per doctor. Load the specific one at the start of each conversation.
- **Variation 3 — Financing-heavy practice:** List up to 3 financing partners in the config. Case Closer will show relevant options per tier.

Success criteria (End of Week 1)

- [ ] Code Execution turned ON
- [ ] Case Closer Skill installed and toggled ON
- [ ] Brand config JSON filled in and uploaded to Project
- [ ] Demo plan run successfully — privacy gate fired, Artifact rendered
- [ ] Handout exported to PDF with your branding
- [ ] Doctor and 1 team member reviewed the sample handout
- [ ] Decision made: which patient's plan (de-identified) to run first in Week 2

References cited from the Skill

- Skill definition: `products/chairside-ai/chairside-case-closer/SKILL.md`
- Python engine: `products/chairside-ai/chairside-case-closer/scripts/financial_menu.py`
- Sample handout: `products/chairside-ai/chairside-case-closer/assets/sample-output-handout.html`
- Demo treatment plan: `products/chairside-ai/chairside-case-closer/assets/demo-treatment-plan.txt`
- Reference documents (bundled with Skill):
  - `references/case-acceptance-psychology.md`
  - `references/tiered-menu-framework.md`
  - `references/financing-language.md`
  - `references/followup-sequence-templates.md`

—

CHAPTER

CHAPTER 6 — Your 3 First Workflows

Three complete workflows, ready to run Monday. Every one is Communication pattern. Every one is PHI-free zone. Every one has a copy-paste prompt.

—

Workflow 6.1 — New Patient Welcome Sequence

What you're building

A 5-touchpoint welcome sequence (text + email) for new patients between "first contact" and "first appointment" — designed to reduce no-shows and set expectations.

Why it matters

The gap between "patient submits form" and "patient sits in your chair" is where 30–50% of new-patient revenue leaks. A patient who booked an appointment two weeks out and heard nothing from you until the reminder text is a patient who half-forgot they booked. A patient who got a warm welcome sequence in that window is a patient who shows up prepared.

Prerequisites

- Claude Desktop + Project set up (Chapter 1).
- Nothing else — this workflow runs on free tier.

Step-by-step setup

1. Open a new chat inside your Project.
2. Paste the exact prompt below.
3. Read Claude's output.
4. Refine once (see below).
5. Save the refined output somewhere your team can access (Notion, Google Doc, a printed SOP).

The exact prompt

```
Build me a new patient welcome sequence.

Trigger: patient submits the new-patient intake form on our website.
Time span: from form submission to first appointment (average 10 days
out).
Format: 5 touchpoints — 3 texts + 2 emails.

Voice rules:
- Warm, plain-spoken, direct. Not corporate.
- Sound like a dentist talking to someone who trusted us with their care.
- Never salesy. Never fake-cheerful.

Content rules:
- Touch 1 (within 1 hour of form submission): confirm the form was re-
ceived, name the appointment date/time slot, tell them what to bring.
- Touch 2 (24 hours later): personal welcome from the doctor. Set expecta-
tions for the first visit.
- Touch 3 (5 days before appointment): what to expect the day of. Insur-
ance verification confirmation.
- Touch 4 (24 hours before): friendly reminder + parking/entry instruc-
tions.
- Touch 5 (2 hours before): warm confirmation. Optional link to fill out
any remaining forms.

Constraints:
- No PHI in the templates — use [First Name] placeholders.
- Segment-level: this is for all new patients, not one specific person.
- Never use: best, guaranteed, painless, permanent.
```

WHAT GOOD OUTPUT LOOKS LIKE

Claude should return 5 message drafts with clear labels. For example:

> ***Touch 1 — Text (within 1 hour):*** *"Hi [First Name] — it's Sarah at Summit Family Dental. We got your intake form and your appointment is confirmed for [date/time]. Please bring your insurance card and a photo ID. If you're a driver's-license state, we recommend getting here 15 minutes early for parking. Any questions in the meantime, text this number."*
>
> ***Touch 2 — Email (24 hrs later):*** *"Hi [First Name] — I wanted to send a quick welcome. I'm Dr. Patel and I'll be seeing you at your first appointment. A few things to know… [continues]"*

If the tone feels like a corporate marketing template ("Dear Valued Patient, Thank you for choosing us…"), that's a bad output. Refine.

The refinement prompt

Paste this if the first output feels too corporate:

```
Make these more conversational. Less "campaign," more "my team and I no-
ticed you booked." Tighten every message — remove any sentence that a real
dentist wouldn't say out loud in the operatory.
```

PHI check

- Everything in this workflow is segment-level. No patient-specific data ever enters Claude.
- Your team fills in the actual `[First Name]` values inside your CRM (GHL, HubSpot, whatever) — Claude never sees the individual patient.

TROUBLESHOOTING

1. **Every message reads identical.** Ask: "Give me 3 different angles for each touchpoint — I'll pick." That forces variety.
2. **The doctor voice feels wrong.** Add to your Project instructions: "Doctor sometimes references her dog Rufus" or "Doctor grew up in the same town as most patients." Character injects voice.
3. **Messages are too long.** Add: "Every text is under 160 characters. Every email is under 120 words."
4. **Claude wrote about a specific procedure.** You didn't say "segment-level." Repeat: "This is for all new patients. Don't mention procedure-specific detail."
5. **Sequence timing feels off for our practice.** Change the intervals in the prompt (1 hr / 24 hr / 5 days / 24 hr / 2 hrs) to match your booking flow.
6. **Team keeps asking for edits.** Great — that's the workflow working. Refine the prompt, not the output. Your goal is a prompt that produces team-ready first drafts.
7. **Insurance confirmation timing conflicts with our verification process.** Change the touchpoint sequence to fit your ops. This is your workflow.
8. **Messages sound generic across practices.** Add practice-specific detail to your Project: your patient base description, common procedures, seasonal notes.

Variations

- **Variation 1 — Cosmetic-only welcome:** Modify Touch 3 to include "the day of your consultation, we'll take photos and imaging so we can walk through your options" — but never use the banned words.
- **Variation 2 — Emergency patient welcome:** Compress to 3 touchpoints (Touch 1, Touch 4, Touch 5) — the timeline is shorter.
- **Variation 3 — Multi-location welcome:** Add "[Location]" placeholder and generate parking/entry instructions per location.

Success criteria (End of Week 1)

- [ ] Prompt run and 5-touchpoint sequence generated
- [ ] At least one refinement pass completed
- [ ] Team reviewed the sequence and approved for pilot

- [ ] Sequence loaded into your CRM (GHL / HubSpot / etc.)
- [ ] First real new patient received the sequence

—

Workflow 6.2 — Review Response Workflow

What you're building

A weekly workflow to respond to every new Google/Yelp/Healthgrades review — grateful for positive, professional for negative, never defensive.

Why it matters

The average dental practice ignores 60% of reviews and responds to the other 40% in a way that hurts more than helps (defensive, salesy, or corporate). A Monday-morning workflow that gets every review responded to within 7 days changes your public-facing brand more than any marketing campaign.

Prerequisites

- Claude Desktop + Project set up (Chapter 1).
- Access to your review platforms (Google Business Profile, Yelp, Healthgrades logins).
- Nothing else — this workflow runs on free tier.

Step-by-step setup

1. Every Monday at 8:00 AM, spend 5 minutes pulling new reviews from the past 7 days.
2. Open a new chat inside your Project.
3. Paste the exact prompt below.
4. For each review, paste the review text and Claude drafts a response.
5. Review + tweak + post to the review platform.

The exact prompt (the "workflow starter")

```
I'm about to paste dental practice reviews from the past week. For each
one:

1. Draft a response in my practice voice.
2. Keep it under 60 words.
3. Follow these rules:
   - Positive reviews: warm, specific, mention one detail they mentioned.
   - Negative reviews: professional, never defensive, never argue facts
publicly.
     Acknowledge, apologize where appropriate, invite them to call the of-
fice
     to make it right. Do not mention specific procedures or clinical de-
tails.
   - Never use: best, guaranteed, painless, permanent.
   - Never disclose PHI publicly. If a review mentions a specific proced-
ure or
     complaint, respond in general terms — never confirm or deny a treat-
ment.
   - Sign every response with "Warmly, [Doctor First Name] and the team at
[Practice Name]"

Ready — I'll paste the first review.
```

The per-review paste

After the setup prompt, paste each review individually with a header:

```
Review from Google, 4 stars, posted 2 days ago:
"Great experience with Dr. Patel and her team. Very thorough cleaning.
Only ding is the wait time — was 20 minutes past my appointment.
Otherwise excellent."
```

Claude drafts. You review. You edit. You post to the platform.

WHAT GOOD OUTPUT LOOKS LIKE

For the review above, Claude should produce something like:

> *"Thanks [First Name] — we're grateful for the trust and thrilled you had a good visit with Dr. Patel. The wait time on that morning is on us — we had an emergency case run long and it cascaded. We're actively working on our scheduling buffer so this doesn't happen again. See you at your next visit. Warmly, Dr. Patel and the team at Summit Family Dental."*

For a 1-star with a specific complaint about a procedure:

> *"Thanks for the feedback [First Name] — we take every concern seriously. We'd like to talk with you directly to understand what happened and make things right. Please call our office at [phone] and ask for [name]. Warmly, Dr. Patel and the team at Summit Family Dental."*

Notice what's NOT in either response: no clinical details, no argument, no "you must be mistaken," no upselling ("come see us again for our best crown!").

PHI check

- Never confirm or deny a specific treatment publicly. Even if the reviewer named it.
- Never use the patient's last name — reviewers use first name only.
- Never respond with clinical detail that suggests you accessed a chart to write the response.
- If a review mentions a procedure, respond in general terms only.

TROUBLESHOOTING

1. **Claude wrote a defensive response.** The prompt failed. Re-paste and add: "Never argue facts publicly, even if the reviewer is factually wrong."
2. **Response is too long.** Add "Under 40 words for negatives, under 60 for positives."
3. **Claude mentioned a specific procedure.** Add to your Project rules: "In review responses, never confirm or deny a specific procedure was done."
4. **Response feels generic.** Ask Claude to mention one specific detail from the review (name, star rating, one word from their comment) to signal you actually read it.
5. **Team wants to A/B test tone.** Ask Claude for 3 versions per review. Pick the winner. Track which style gets thank-you replies.
6. **Owner wants to sign as the doctor personally, not "the team."** Change the signature line in the setup prompt.
7. **Reviewers rarely give names.** Use "friend" or "our neighbor" or nothing — "Thanks so much for the kind words…"
8. **Negative reviews get responses days late.** Move the weekly cadence to twice weekly (Mon + Thu) if you get more than 5 reviews/week.

Variations

- **Variation 1 — Per-doctor responses:** Sign each response as the specific doctor. Requires tagging reviews by which doctor treated them (your team does this offline).
- **Variation 2 — Multi-language:** Add "Some reviews may be in Spanish. Respond in the same language the reviewer used."
- **Variation 3 — Escalation flag:** Add to the prompt: "If a review mentions injury, malpractice, or legal action, flag it 'escalate' and do not draft a response. I'll handle it offline."

Success criteria (End of Week 1)

- [ ] Setup prompt saved as a template (Notion, Google Doc, or the top of your Monday chat)
- [ ] First Monday morning session completed — all reviews from the past week responded to
- [ ] At least one negative review handled without defensiveness

- [ ] Team member briefed on the workflow so it survives if you're on vacation

—

Workflow 6.3 — Treatment Plan Follow-Up

What you're building

A same-day follow-up email that summarizes the treatment plan discussion, options presented, and next steps — sent within 1 hour of the patient leaving the operatory.

Why it matters

Case acceptance drops 15% for every 24 hours of delay after the chairside discussion. A patient who left with a printed treatment plan and no follow-up is a patient who's already forgetting the doctor's rationale by dinner. A patient who got a warm, specific follow-up email while they were still in the parking lot is a patient who's already talking to their spouse about it.

Prerequisites

- Claude Desktop + Project set up (Chapter 1).
- **Recommended:** Chairside Case Closer Skill installed (Chapter 5) — Case Closer produces a richer handout AND the follow-up email. If you don't have Pro tier, this workflow's basic version below works fine.
- **PHI check:** if you use Case Closer, you paste a de-identified plan. If you use the basic version below, you paste ZERO patient data — you paste only the segment description (age band, insurance type, stated goal).

Step-by-step setup — basic version (free tier)

1. Right after the chairside treatment plan discussion, before the patient reaches the front desk, open a new chat in your Project.
2. Paste the exact prompt below.
3. Fill in the segment description.
4. Read Claude's draft.
5. Edit in your voice.
6. Send from your practice email (not your personal one) within 1 hour of the patient leaving.

The exact prompt (basic version)

```
Draft a treatment plan follow-up email for a patient I just saw.

Segment details (no patient identifiers):
- Age band: [late 50s / early 30s / etc.]
- Insurance type: [PPO / self-pay / etc.]
- Stated goal (their words): "[e.g., I want to chew on both sides without
pain]"
- Recommended treatment (general category, not specifics):
  [e.g., implant + crown for #14, buildup + crown for #19]
- Financing options discussed: [CareCredit / Sunbit / in-house]

Voice:
- Warm, plain-spoken, direct.
- Not salesy. Not fake-cheerful.
- Reference their stated goal in the opening line.
- Doctor voice, not corporate.

Constraints:
- No PHI in the email — use [First Name] placeholder.
- Segment-level rationale, not patient-specific chart data.
- Under 200 words.
- Include: (a) reference to their stated goal, (b) brief rationale for the
recommended treatment,
  (c) financing options mentioned, (d) invitation to call/reply with ques-
tions,
  (e) doctor signature.
- Never use: best, guaranteed, painless, permanent.
```

The exact prompt (Pro tier — with Case Closer)

If you have Case Closer installed:

```
Turn this de-identified treatment plan into a Case Closer follow-up se-
quence.

Age band: [late 50s]
Insurance: [PPO, $1,200 remaining, ~50%/80%]
Stated goal: "[I want to chew on both sides without pain]"

Recommended:
- Urgent: [tooth #14 extraction + implant + crown - $5,800]
- Urgent: [tooth #19 buildup + crown - $1,850]
- Recommended: [tooth #3 composite - $320]
- Recommended: [night guard - $650]

Financing: [CareCredit, Sunbit]

Give me the full Case Closer output: handout + 4-touch follow-up sequence
(Day 0, 3, 7, 14) + internal coaching note.
```

Case Closer will produce a handout + sequence — you use Touch 1 (Day 0) as the same-day follow-up email.

What good output looks like (basic version)

> ***Subject:*** *Following up on today, [First Name]*
>
> *Hi [First Name],*
>
> *Thanks for coming in today. You told us you want to chew on both sides again without pain — I want to make sure you leave with a clear picture of how we get you there.*
>
> *We recommend addressing #14 first with an implant and crown — that's what solves the chewing issue. #19 needs a buildup and crown to keep it stable long-term. Both are the priority.*
>
> *The other items I mentioned (composite on #3, night guard) are worth doing but can wait if you'd rather stage it.*
>
> *On financing — we offer CareCredit and Sunbit. Both have no-interest promotional plans if paid off in the promo window. My team can walk you through the numbers when you're ready.*
>
> *Reply to this or call the office if any questions come up. No pressure — just wanted to make sure you have what you need to decide.*
>
> *Warmly, Dr. Patel*

If the draft mentions a specific chart detail ("the deep cavity on the mesial") or uses banned words ("this is the best option"), edit it.

PHI check

- The prompt intentionally uses segment-level detail only. Age band, insurance type, general procedure category. Not name, DOB, chart-level clinical detail.
- You paste no chart data into Claude. Your team writes the final email in your CRM with the actual patient's first name manually filled in.
- The mental model: Claude drafts the shell. Your team assembles the letter.

TROUBLESHOOTING

1. **Email is too long.** Add: "Under 150 words. Cut anything a doctor wouldn't say out loud."
2. **Sounds pushy.** Add: "Never use urgency language. Never reference 'limited-time' anything."
3. **Financing feels salesy.** Add: "Mention financing in one sentence, no dollar figures, no promises."
4. **Doctor wants a different signature.** Change the sign-off line in the prompt.
5. **Different doctors want different voices.** Create one Project per doctor with a slightly different voice field.
6. **Team is skipping the same-day send.** The workflow value drops 15% per day of delay. Consider setting an alarm on the front desk POS for 45 minutes after each chair-out.
7. **Case Closer output feels too formal for a follow-up email.** Use only Touch 1 (Day 0) from the sequence — that's the warmest touch.
8. **Patient replies with questions.** Great — you have a live conversation. Respond personally, not with another AI-drafted email.

Variations

- **Variation 1 — Cosmetic case follow-up:** Add a photo attachment expectation to Touch 1 — "I'll send imaging and mockups by Thursday" — then deliver on Thursday.
- **Variation 2 — Split-financing patients:** If the patient asked about splitting between insurance and financing, add a numbers-broken-out section (using Case Closer to compute).
- **Variation 3 — Non-committal patients:** For patients who said "I'll think about it," modify Touch 1 to end with "No pressure to decide anything by email — call or reply if any specifics come up as you think it through."

Success criteria (End of Week 1)

- [ ] Prompt saved as a template accessible to the front desk
- [ ] First real patient received a same-day follow-up within 1 hour of leaving
- [ ] Team briefed: the workflow is your responsibility, not the front desk's judgment
- [ ] By Friday: 3 patients received the follow-up. Track which ones replied.

—

APPENDIX

APPENDIX A — The PHI-Free Zone Cheat Sheet

The rule

Consumer Claude, ChatGPT, and Grok do not sign BAAs. You never paste PHI. You work at the segment level.

Segment-level is safe (no BAA needed)

- "Patients overdue 6+ months"
- "New patients this quarter"
- "PPO patients with a remaining max of $500+"
- "Patients who booked but haven't shown"
- Marketing copy, brand voice, tone examples
- Hiring rubrics, interview questions
- Vendor comparisons, financial projections
- SOP drafts, team communication
- Fee schedules, overhead analysis
- Review response drafts (no PHI in reviews)

Patient-level is NOT safe (unless you use a BAA-backed tool)

- Any full patient name + any medical or dental detail
- DOB + medical detail
- Chart numbers, claim numbers, member IDs
- Insurance letters that name the patient
- Treatment plans that identify the patient
- Photos or X-rays of identifiable patients

- Voice recordings of appointments

When you MUST touch PHI

Use the redaction-before-prompt pattern:

1. Strip all identifiers before you paste.
2. Work on the de-identified data.
3. Get output.
4. Human review before anything goes out.
5. Re-attach patient context manually in your PMS or CRM — Claude never sees it.

For anything more PHI-adjacent than the above, use a BAA-backed tool (many practice management vendors are shipping BAA-covered AI features in 2026). See Zach Hanson's 12:30 Bulletproof session for the current landscape.

—

APPENDIX

APPENDIX B — Where to Get Help

- **Chairside AI Kit downloads:** `aipg.co/chairside`
- **AIPG resources hub:** `aipg.co/resources`
- **Blake's email for Bulletproof attendees:** `blake@aipg.co`
- **Book a call:** `aipg.co/call` (30-minute intro)
- **Community:** the Bulletproof Slack workspace — search for `#ai-implementation`

—

APPENDIX

APPENDIX C — Chairside AI Kit Index

All 8 Skills. All PHI-free. All free at `aipg.co/chairside`.

1. **Chairside Case Closer** — treatment plan → patient-ready case presentation with tiered financing, phasing, and follow-up sequence. Pro tier only.
2. **Chairside Money-Leak Finder** — bank/card statement CSV → ranked list of wasted spend with renegotiation scripts pre-filled. Pro tier only.
3. **Chairside Fee Auditor** — fee schedule CSV → profit-leverage math, flagged codes, tiered increase recommendation, re-key CSV, front-desk scripts. Pro tier only. (Featured in Chapter 3.)
4. **Chairside Morning Huddle Brief** — de-identified schedule → 1-page huddle sheet with high-value case routing, open-chair fill plan, and team talking point. Pro tier only.
5. **Chairside Hiring Kit** — role description → interview rubric + scorecard + calibration questions. Free tier works.
6. **Chairside EOB Decoder** — de-identified EOB → plain-English explanation for team training. Pro tier recommended.
7. **Chairside Cyber Self-Audit** — 15-question audit → prioritized remediation list. Free tier works.
8. **Chairside Vision Coach** — practice goal → 90-day plan with weekly check-ins. Free tier works.

Each Skill lives at `products/chairside-ai/[skill-name]/` in the AIPG-OS repo, with a `SKILL.md`, `scripts/`, `references/`, and `assets/` folder. Every Skill is auditable — read the SKILL.md before you install if you want to see exactly what it does.

—

Built for the Bulletproof Dental CE Summit — August 9, 2026 — by AIPG. Blake McClellan.