

The Chugai Marketing AI Toolkit

Foundation Edition

13 MAY 2026 · FOR HELEN MILLAR · AI COMMERCIAL EXCELLENCE PROGRAMME

The Chugai Marketing AI Toolkit

Foundation Edition

A practical playbook for getting elite-level work out of Claude Enterprise. Grounded in the 4Cs framework from Chugai PACE training. Built on real worked examples from oncology launch practice. Designed both as a personal playbook and as the foundation for wider Chugai Marketing AI training.

Version 1.0 — 13 May 2026

Executive summary (read this in five minutes)

There are three things that separate elite AI users from average ones, and they are all teachable.

1. **They front-load context.** The first message contains the audience, the format, the role, the constraints and the brief. Six messages of correction become one well-aimed shot.
2. **They challenge the workings.** Every number is interrogated. Every claim is asked to cite its source. Confidence is treated as evidence of nothing.
3. **They use a second pair of eyes.** Critique happens in a clean chat, never in the same one that produced the draft. This is the single biggest hidden lever in advanced prompting.

Everything else in this document is in service of those three behaviours, structured against Chugai's own **4Cs framework: Clear, Context, Constraints, Check.**

If you take only three things from this document, take these:

1. **Save your reusable context block** (Section 3) into a Claude Enterprise Project's custom instructions. Stop typing it into every chat.

2. **Use the 30-second pre-send check** (Section 5) before you hit send on anything that matters. It maps directly to the 4Cs.
3. **Critique drafts in a new chat, not the same one** (Section 6, the Two-Model Critique). Fresh context kills anchoring.

The rest is depth and breadth. Build the three habits first.

Section 1. Why this exists

This document has two audiences.

Audience 1: a single skilled user (Senior Brand Lead, Oncology) who wants the next level. You are already an effective prompter. You front-load constraints, you assign personas, you challenge the maths, you set persistent rules, you escalate complexity progressively. The Opus profile you generated on 13 May was accurate on all counts. This document gives you the next set of moves.

Audience 2: the rest of Chugai Marketing. As the Chugai AI capability initiative scales, this document is the spine of the training. Sections 11 and 12 in particular are written to be lifted directly into team enablement material.

The two audiences point to the same document because elite practice is what we want to teach. The worked examples throughout are drawn from real oncology launch prompting (the Akynzeo OS launch, the Japanese sales team brief, the Ireland business case, the ExCo deck), suitably abstracted. The lessons travel.

Section 2. The 4Cs framework as the spine

Chugai's PACE training already gave us the right structure. Every technique in this document maps to one or more of the 4Cs.

C	Question to ask before sending	What it covers
Clear	What exactly do I want, in what shape?	Output format, length, structure, deliverable type
Context	What does Claude need to know about the situation?	Role, audience, business background, prior work
Constraints	What must be true (or not true) about the output?	Compliance, scope, must-haves, must-not-haves, standing rules
Check	How will I verify the output is right?	Sources, calculations, sceptical second look, leak test

A useful exercise: take any prompt and score each C from 1 to 5. The weakest C is where the next correction round will come from.

Every advanced technique in Sections 6 onwards is labelled with which Cs it strengthens. Use the 4Cs as the assessment lens whenever you are coaching another team member.

Section 3. The reusable context block (the single biggest unlock)

Most users drip-feed context across five or six messages. The fix is one block of text saved permanently into a Claude Enterprise **Project**'s custom instructions, or pasted at the top of any cold chat.

Here is a template, calibrated for an Oncology BUL:

You are a senior commercial strategist with deep oncology launch experience.

I am the Senior BUL for Oncology at Chugai UK & Ireland, and the launch lead for Akynzeo OS. My audience defaults to ExCo, Japan HQ, or Helsinki unless I say otherwise.

House style: plain British English, no marketing fluff, show formulas in Excel, use the Chugai business case template when producing business cases, exclude NEC unless asked, no qualifying caveats like "consult your legal team".

Compliance lens: the ABPI Code of Practice (current edition) and UK Competition Law are non-negotiable. Flag risks proactively.

Working style: before drafting, ask up to 5 questions only if my answers would materially change the output. Otherwise proceed. Show your workings on any calculation. Cite sources or mark claims as inferred.

That is roughly six lines. It saves six messages of correction per chat. Across a year of practice it saves hundreds of iteration rounds.

4Cs coverage: Context (role, audience, situation), Constraints (house style, compliance), Check (show workings, cite or mark inferred).

Quick orientation: Projects and Custom Instructions

If you have only ever used Claude Enterprise as a chat interface, two terms used throughout this guide are worth understanding before you start.

A **Project** is a workspace inside Claude Enterprise. Think of it as a folder you set up once and reuse for one area of your work — for example, your *Akynzeo OS Launch* Project, your *Brand Planning* Project, your *ExCo & Japan HQ Communications* Project. Each Project holds three things in one place: custom instructions, reference documents you upload, and all your chats on that topic. Every chat you start inside a Project automatically inherits its context.

Custom Instructions are a saved preamble attached to a Project. Anything you write in that field loads automatically into every new chat you start inside the Project. You write your role, your audience defaults, your house style, and your compliance rules once. From that point on, every chat begins with all of it already loaded.

Where to find them: open Claude Enterprise. In the left sidebar, look for **Projects**. Click **Create Project**, give it a name (for example, "Akynzeo OS Launch"). Inside the Project's settings, find the **Custom Instructions** field. Paste your context block there. Save. Done.

Every future chat you start from inside that Project begins with your context block already applied. You will never need to retype it.

A team adaptation pattern: each team member writes their own version of the context block, kept inside their own Project. The skeleton is shared; the specifics (role, audience defaults, house style) are personal.

Section 4. Use Projects, not loose chats

This is the single biggest productivity move available on Claude Enterprise, and most users never make it.

Spin up a Project per ongoing workstream. For oncology launch practice that is probably:

- Akynzeo OS Launch
- Oncology Brand Planning (current year)
- Competitive Intelligence (Oncology)
- ExCo & Japan HQ Communications
- AI Commercial Excellence Programme

Each Project holds three things:

1. **The reusable context block** in the custom instructions field
2. **Reference documents** uploaded once and reused everywhere (current brand plan, forecast model, ABPI Code, competitor intel pack, business case template)
3. **All chats related to that workstream**, with continuity of memory across them

Stop starting new chats from a cold context. By the time you type the first word, your reference material and your rules should already be loaded.

Pharma-specific reference docs worth uploading once into the right Project:

- The current ABPI Code of Practice
- Chugai's house style guide (if one exists)
- The Chugai business case template
- The current Akynzeo OS forecast and assumption log
- The latest competitive intelligence pack

- Prior ExCo papers (for tone calibration)

Section 5. The 30-second pre-send check

Before you hit send on any prompt that matters, run this seven-point check. It maps directly to the 4Cs.

#	Check	Which C
1	Have I named the audience?	Context
2	Have I specified the output format and length?	Clear
3	Have I assigned a persona?	Context
4	Have I included the must-haves and must-not-haves?	Constraints
5	Have I flagged which compliance lens applies?	Constraints
6	Have I told Claude how to handle uncertainty (cite, flag, ask)?	Check
7	Have I told Claude what to do if it disagrees with my framing?	Check

If all seven are answered in the prompt, you have already won most of the iteration rounds before they start.

The aim is not to be slow. With practice this check runs in roughly thirty seconds. The result is that prompts which used to take five rounds now take one.

Section 6. The advanced moves (eight techniques)

Each technique below follows the same teaching structure: name, which Cs it strengthens, when to use it, the template, a worked example from real practice, and the common mistake to avoid. Pick two or three. Build the habit. Add more once those are automatic.

6.1 The Two-Model Critique

Strengthens: Check.

When to use: any draft that will go to ExCo, Japan HQ, Helsinn, or any external audience.

Template:

*[Open a **new chat**. Clean context. Paste the draft.]*

"You are a sceptical [ExCo member / Helsinn commercial lead / Japanese sales VP]. You have five minutes before the meeting. Find three weaknesses in this paper, and the one question that would derail the recommendation in the room."

Worked example: after Claude produces a one-pager for ExCo on the Ireland part-time contractor business case, open a new chat, paste the one-pager, and prompt it as the sceptical ExCo CFO. The original chat is anchored on the framing that the contractor is justified. The new chat is not. The new chat will find the holes faster than you can.

Common mistake: asking the same chat to critique its own work. It will be polite, anchored, and gentle. A fresh chat is brutal.

6.2 The Output Template

Strengthens: Clear, Constraints.

When to use: any repeating format you will produce more than twice.

Template: paste an explicit fenced template at the end of your prompt.

Output exactly as below. No preamble. No closing remarks.

TITLE (max 8 words):

BLUF (bottom line up front, max 30 words):

KEY POINTS (3, one sentence each):

RISKS (max 3, each flagged as commercial / ABPI / competition law):

DECISION NEEDED:

Worked example: for ExCo summaries, the BLUF + three points + three risks + decision-needed pattern is reusable across forecasts, launch updates, business cases, and competitive intelligence briefings. Save it as a template once. Reuse forever.

Common mistake: describing the format in prose instead of showing the format as a literal template. Claude follows literal templates almost perfectly. It interprets prose descriptions loosely.

6.3 The Quote-the-Paragraph Check

Strengthens: Check.

When to use: any output that cites internal documents, peer-reviewed evidence, market data, or competitive intelligence.

Template:

"For each claim above, quote the exact sentence from the source document. If you cannot quote it, mark the claim as inferred."

Worked example: when Claude summarises the current Akynzeo OS forecast assumptions and says something like "the model assumes 40% conversion in Wave 1", ask it to quote the sentence in the forecast doc that says so. If it cannot, it has invented the figure. This catches a category of hallucination that no other technique catches as fast.

Common mistake: asking "where did this come from?" generically. Claude will produce a plausible-sounding citation. Force it to **quote**, verbatim.

6.4 Negative Prompting

Strengthens: Constraints.

When to use: anywhere Claude has a default behaviour you do not want.

Template: a short list of "do not" instructions appended to the brief.

"Constraints:

- *No marketing buzzwords.*
- *No "consult your legal team" caveats.*
- *No invented statistics.*
- *No qualifying phrases like 'it is worth noting' or 'it is important to recognise'.*
- *No bullet points longer than one line.*
- *No closing sentence that summarises the document."*

Worked example: when you ask Claude for an ExCo paper, the first draft typically comes back padded with hedging phrases — "it is worth noting that...", "while it is important to recognise...", "this is not without its challenges", "in some respects". Three or four of these per paragraph. They add bulk without adding meaning, and you end up editing each one out by hand. A negative-prompt block tells Claude not to produce them in the first place. You save five minutes per paper.

Common mistake: assuming Claude will infer that you do not want marketing-speak. It will not, unless you say so explicitly.

6.5 The Worst-Headline Test

Strengthens: Check (pharma-specific).

When to use: any draft that might face external eyes, including investor updates, partner communications, conference materials, and press-adjacent content.

Template:

"If this content leaked to PMLive or Pharma Letter tomorrow, what is the worst possible headline a journalist could write? Rewrite to eliminate that risk."

Worked example: a draft Akynzeo OS competitive positioning paragraph contains a phrase comparing the product favourably to Aloxi by name. The worst-headline test surfaces it as a potential disparagement / comparative-claims risk under the ABPI Code. The rewrite removes the comparison and lands on a safer claim grounded in label data.

Common mistake: only running this on press-facing material. The risk surface is wider than press releases. Internal slides leak. ExCo decks are forwarded. Run the test on anything that could plausibly end up in front of an outsider.

6.6 The Pre-Mortem

Strengthens: Check.

When to use: any significant decision document going to ExCo or Japan HQ.

Template:

"Imagine this recommendation was approved, then failed twelve months later. What is the single most likely reason it failed? What are we missing now that we will regret then?"

Worked example: the Ireland part-time contractor business case approved at 0.5 FTE for two years. Twelve months in, the contractor has left, no replacement is in place, Wave 2 conversion stalls. Pre-mortem in advance would have flagged single-person dependency risk and a knowledge transfer plan. The decision still goes ahead, but with the mitigation built in.

Common mistake: running the pre-mortem only on documents you already suspect are weak. Run it on the strong ones too. Strong cases that have failed are the most painful, because the post-mortem reveals everyone agreed and nobody pushed.

6.7 Few-Shot by Screenshot

Strengthens: Clear.

When to use: any format-sensitive output where you have a reference example. Slide layouts, table styles, specific deck templates, brand-deck visuals.

Template:

"Here is what good looks like for this slide format. Match this layout, this density of text, and this style of heading. [paste screenshot]"

Worked example: the recurring "boxes on top of each other" PPT problem in the II May prompt sequence. A single screenshot of an ExCo slide that landed well, with the instruction to match the layout, eliminates several iteration rounds.

Common mistake: describing the format in words when an image would carry it. One image is worth a thousand instructions on this category of task.

6.8 The Ask-Me-Questions Preamble

Strengthens: Context.

When to use: complex business cases where assumptions will materially change the output (forecasts, ramp rates, FTE planning, scenario modelling).

Template:

"Before you start, ask me up to five questions where my answer would materially change your output. Wait for my answers before drafting."

Worked example: a fresh Ireland forecast scenario request. Without this preamble Claude will quietly assume a default Wave 1 ramp curve, a default FTE productivity rate, and a default conversion latency. With this preamble, it surfaces those assumptions as questions, you correct three of them upfront, and the first draft is closer to right than the third draft would have been otherwise.

Common mistake: using this on every prompt. Reserve it for prompts where assumption errors will be expensive. For routine tasks the questions waste a round.

Section 7. Worked examples from real practice

Three strong prompts, lightly abstracted, with annotations on what makes them strong. Three weaker prompts with the fix applied.

7.1 Strong example: the conversion ramp Excel request (11 May)

The prompt (abstracted):

"For the attached conversion ramp, please show the conversion flow by account with assumptions/workings in Excel so I can share with the forecaster. Have a tab for units. On a separate tab please show [the per-month breakdown]."

What makes it strong:

- **Clear:** Excel as the output format, explicit tab structure, exactly what to show
- **Context:** the audience (the forecaster) and the purpose (sharing for review)
- **Constraints:** assumptions and workings exposed, multi-tab structure
- **Check:** implicit, because the forecaster will check the workings

The lesson for the team: a multi-layered request can still be crystal clear if every layer is named. This is one of the strongest prompt shapes in this entire document.

7.2 Strong example: the Akynzeo OS launch role context (11 May)

The prompt (abstracted):

"I'm the launch lead for Akynzeo OS. We don't have a process for launch at Chugai yet so we're building the prototype as we go along. I've set up nine workstreams with a lead for each..."

What makes it strong: this is pure context-setting. It tells Claude who is asking, what the business situation is, and what has already been done. Anything Claude produces next is calibrated against the right reality.

The lesson for the team: investing a paragraph in role context at the start of a complex chat pays for itself within two messages.

7.3 Strong example: the recalibration prompt (12 May)

The prompt (abstracted):

"This is too ambitious for a part-time contractor, looks more like the expectation for a current established full-time KAM. Please adjust expectations accordingly so they reflect 0.5 FTE and time to onboard..."

What makes it strong: domain expert challenge. Claude produced something plausible-looking, the BUL recognised it was unrealistic, and the correction is specific, expert, and bounded.

The lesson for the team: the willingness to push back on the AI is the single most underrated prompting skill. Domain experts do this naturally. New users tend to accept the first output. Teach new users to challenge.

7.4 Weak example: the contextless "1 page exec summary" (11 May)

The prompt:

"1 page exec summary"

Why it falls short: zero context. No audience, no tone, no purpose. Claude has to infer everything.

The fix:

"1-page exec summary for ExCo. Focus on financial impact and the decision required. Keep to five or six bullet points maximum. BLUF up front in 30 words. Plain English. Flag the one risk that would change the decision."

The lesson for the team: the first instinct when typing a short prompt should be to add five more words of context. The marginal cost is tiny. The marginal value is large.

7.5 Weak example: the two-issues-in-one-prompt (11 May)

The prompt:

"Can you re-do as 30% or 40%. The ppt formatting didn't work very well so boxes on top of each other."

Why it falls short: two unrelated issues (a data change and a formatting bug) bundled into one prompt with no priority. Claude will likely address one well and the other partially.

The fix: split into two messages.

*"1) Please re-do the scenario using 30% and 40% conversion rates as alternative cases.
2. The PPT boxes overlapped on the layout. Can you simplify to a single-column layout?
Reference the attached screenshot for what worked previously."*

The lesson for the team: one issue per prompt is a simple rule that prevents this whole category of partial outputs.

7.6 Weak example: the request for an impossible capability (13 May)

The prompt:

"Please share the link to the slide deck."

Why it falls short: Claude cannot generate shareable links to files it creates. Repeating the request does not change the capability.

The fix: change the request shape to something Claude can deliver.

"Please regenerate the slide deck content as a structured outline I can paste into PowerPoint. One section per slide. Include slide titles, bullet points, and speaker notes."

The lesson for the team: when Claude says it cannot do something, do not repeat the same request. Reshape the request to fit what it can do.

Section 8. Common failure patterns (extracted from real practice)

These nine patterns account for the majority of avoidable iteration rounds. They are presented here as a checklist for self-review and for coaching others.

#	Pattern	The fix
1	Bare "Yes" confirmations that lose context across turns	Echo what you are confirming. "Yes, proceed with the 0.5 FTE scenario."
2	Short prompts missing audience/purpose	Add five words of context. Who is it for? What is it for?
3	Format requests with no source reference	Always specify which data, which model, which prior message
4	Two unrelated issues in one prompt	Split into separate messages
5	Repeating a request for a capability Claude does not have	Reshape the request to fit what Claude can do
6	Standalone factual queries with no context	Add why you need it. The framing changes the answer's usefulness.
7	Instructions without rationale	Include the why. Claude can then flag downstream issues.
8	Vague qualitative requests ("make it better", "more impact")	Specify which paragraph, what aspect, against which benchmark
9	Error reports without diagnostic info	Say what you saw vs what you expected. Ask if there was a generation error.

A practical training exercise: take a team member's last ten prompts from a real chat. Score each against this list. Coaching beats lecturing.

Section 9. Pharma-specific compliance hygiene

This is the section the generic prompting guides do not include. Pharma marketing has compliance constraints that AI tools do not understand by default. The hygiene below is non-negotiable for any output that could face external eyes.

9.1 Data classification before paste

Treat the chat window like an email you might forward by mistake. Before pasting, check the source against Chugai's data handling policy. Do not paste:

- Identifiable HCP data
- Patient-level data
- Anything classified Chugai Confidential without confirming the policy allows AI processing
- Partner commercial terms (Helsinn or otherwise) without confirming sharing rights
- Pre-publication study data
- Anything under embargo

When in doubt, redact first, prompt second.

9.2 Triple compliance pass

For any output that might face external eyes:

1. **Draft chat:** generate the content.
2. **New chat (clean context):** "Review this draft line by line against the current ABPI Code of Practice. Flag every potential breach. For each, cite the relevant clause."
3. **New chat (clean context):** "Act as a UK CMA lawyer specialising in pharmaceutical competition. Flag every potential competition law risk in this draft, particularly around comparative claims, market definition, and pricing communications."

Three passes, three clean contexts, no anchoring across them. This catches things a single chat will miss because it has already committed to the framing.

9.3 Standing constraints for any pharma marketing prompt

Build these into your context block (Section 3) so they apply automatically:

- Disease-area claims must be on-label
- Comparative claims must be substantiated by referenced data
- Off-label discussion is flagged, not produced
- Patient-facing content respects health literacy norms
- HCP-facing content respects the Code distinction between promotional and non-promotional material

9.4 The leak test

For any internal document that could plausibly be forwarded:

"If this document was screenshot and posted to LinkedIn tomorrow, what is the worst implication for Chugai? Rewrite to remove anything that could not survive that exposure."

Section 10. Tooling discipline (Claude Enterprise specifics)

10.1 Match the model to the task

Claude Enterprise gives you three model tiers. Using the largest for everything is slow and wasteful.

Model	When to use
Opus	Forecast modelling, ExCo narratives, strategic memos, scenario design, anything where the answer materially matters
Sonnet	80% of daily work. Drafting, summarising, table generation, brief expansion
Haiku	Quick lookups, bulk text cleanup, simple reformatting, single-paragraph tasks

You are not penalised for using Haiku on simple tasks. You are penalised (in waiting time and in attention) for using Opus on everything.

10.2 Use artifacts for deliverables

When you want a document, table, or chart you will iterate on, ask Claude to produce it as an **artifact**. You can then edit it directly in the side panel instead of having Claude regenerate the whole thing each round. This alone removes most of the slow "no, change just that one bit" cycles.

10.3 File connectors

If your Claude Enterprise tenant has SharePoint or OneDrive connectors enabled, learn the syntax for referencing internal documents in prompts. This is the difference between Claude searching the live source and Claude inferring from your description.

10.4 Voice and tone control

Claude Enterprise supports custom style profiles. Train one against three or four of your best previous ExCo summaries and Claude will match the tone automatically on future generations.

Section 11. The promptbook starter pack

The European Prompt Library goal starts with a personal promptbook of the ten prompts you will reach for most. Below is a starter pack calibrated for Chugai Oncology Marketing. Adapt and extend.

Each entry uses this structure:

```
PROMPT NAME: [short label]
USE CASE: [when to grab this]
PROMPT TEMPLATE: [the actual text with {placeholders}]
EXAMPLE INPUT: [a real example]
EXAMPLE OUTPUT: [trimmed to show shape]
NOTES: [what to watch for]
```

11.1 ExCo one-pager summary

You are a senior commercial strategist. Produce a one-page ExCo summary of the attached {document}.

Audience: ExCo (commercial, medical, finance, legal in the room).

Output format:

TITLE (max 8 words):

BLUF (max 30 words):

KEY POINTS (3, one sentence each):

RISKS (max 3, each flagged as commercial / ABPI / competition law):

DECISION NEEDED:

Constraints: plain British English, no marketing buzzwords, no qualifying caveats. Flag any claim that you cannot quote directly from the source document.

11.2 Forecast assumption check

Below is a forecast scenario for {product} for {market} over {timeframe}. List every assumption baked into the numbers, classified as:

- (a) explicit (stated in the document)
- (b) implicit (inferable from the maths)
- (c) inherited (carried over from prior modelling without revalidation)

For each assumption, rate the sensitivity: would a 20% change move the headline number by more than 10%? If yes, flag for revalidation.

11.3 ABPI Code review

You are a compliance reviewer specialising in the current ABPI Code of Practice. Review the attached {content} clause by clause. For each line or paragraph, flag any potential breach. Cite the specific Code clause. Rate severity as high / medium / low. Do not soften your assessment.

11.4 UK competition law review

You are a UK CMA lawyer specialising in pharmaceutical competition. Review the attached {content} for competition law risk. Pay particular attention to comparative claims, market definition, pricing communications, and any language that could be characterised as exclusionary or anti-competitive. Cite specific risk areas. Suggest specific rewordings where applicable.

11.5 Competitor positioning paragraph

You are a senior brand strategist for {our product} in {therapy area}. Produce a positioning paragraph against {competitor product}.

Constraints:

- Stay on-label for {our product}
- Substantiate every comparative claim with referenced data
- No disparaging language
- ABPI Code compliant
- No more than 120 words

11.6 Internal change comms (UK team)

You are an internal communications partner. Produce a {short / medium / long} message to {audience} about {change}.

The change context: {paragraph}.

Tone: clear, direct, honest about what is changing and why. No corporate filler. Acknowledge the human impact. End with the single most important action the reader needs to take.

British English. No jargon. Plain enough for the newest team member to understand on first read.

11.7 Meeting prep brief

You are my chief of staff. I have a {meeting type} with {attendees} at {time}. The purpose is {purpose}. Attached: {prior materials}.

Produce a one-page prep brief:

- Three key points I should land
- The two most likely questions from {attendees}
- The one thing I should not say
- My desired outcome from this meeting (one sentence)

11.8 Sceptical ExCo critique (use in a new chat)

You are a sceptical ExCo member. You have five minutes before the meeting starts. The attached paper is in front of you.

Find three weaknesses in this paper, and the one question that would derail the recommendation in the room. Be direct. Do not soften.

11.9 Worst-headline test

The attached content {may / will} be {internal-only / externally visible}. Imagine it leaks to PMLive or Pharma Letter tomorrow.

What is the worst possible headline a journalist could write from this content? Rewrite the content to eliminate that risk while preserving the business message.

11.10 Brand plan narrative scaffold

You are a senior brand planner. Produce a narrative scaffold for the {year} brand plan for {product} in {market}.

Inputs: attached {market context, performance to date, competitive landscape, ambition for the year}.

Output format:

- Where we are (one paragraph)
- Where we are going (one paragraph)
- The three battlegrounds (one paragraph each)
- How we will win on each (one paragraph each)
- The one thing we are not doing this year (one paragraph)
- The single biggest risk (one paragraph)

Constraints: British English, plain language, no marketing buzzwords, no "unlock", "leverage", "synergy" or "journey".

Section 12. Training the team (the rollout playbook)

This section is written to be lifted directly into Chugai Marketing AI capability training material.

12.1 Three maturity tiers

Adoption is not uniform. Plan for three distinct user populations.

Tier	Description	What they need
Beginner	Has used Claude a handful of times. Tends to write one-line prompts. Accepts most outputs as given.	The 4Cs framework, the 30-second pre-send check, the reusable context block. Aim for them to write five-line prompts by end of week one.
Capable	Uses Claude regularly. Front-loads context. Sets standing rules. Occasionally challenges outputs.	The eight advanced techniques in Section 6. The triple compliance pass. The promptbook habit. Aim for them to internalise the Two-Model Critique as a default by end of month one.
Expert	Uses Claude as a daily collaborator. Treats it like a junior analyst. Challenges every number. Sets up Projects, uses artifacts, manages model selection deliberately.	The teach-others layer. They become the team coaches.

12.2 The rollout sequence

A six-week structure that has worked elsewhere:

Week	Focus	Outcome
1	Mindset and the 4Cs	Every team member can score a prompt against the 4Cs
2	The reusable context block	Every team member has a personal context block saved in a Project
3	Output templates and the 30-second check	Three standard output templates adopted across the team
4	Compliance hygiene	Every team member has used the triple compliance pass at least once on a real document
5	The Two-Model Critique	Every team member has critiqued one of their own drafts in a new chat
6	Promptbook and team review	Each team member contributes three prompts to the team promptbook

12.3 Common stumbling blocks

Block	What you will hear	What to do
"Writing the prompt is harder than just doing the work"	True for the first ten prompts. Stops being true after that.	Show them the time math: 6 messages × 5 chats per day × 200 days = 6,000 messages saved per year on a single saved context block.
"I don't trust the output"	Correct instinct.	Teach the Check column of the 4Cs as the answer, not as a barrier.
"It hallucinates"	Correct observation.	Teach the Quote-the-Paragraph check as the primary defence.
"I can't put confidential data in"	Sometimes correct.	Walk through Section 9.1 with them. Distinguish what is genuinely sensitive from what feels sensitive.
"It writes in marketing-speak"	Default behaviour.	Teach negative prompting (Section 6.4).

12.4 How to assess prompting skill

A simple rubric for line managers reviewing a team member's prompts:

Score	What it looks like
1 — Beginner	One-line prompts. Generic outputs. Accepts first draft.
2 — Aware	Three to four sentences of context. Sometimes specifies format. Occasionally challenges output.
3 — Capable	Front-loads audience, format, role, constraints. Uses Projects. Challenges calculations routinely.
4 — Skilled	Uses 4+ techniques from Section 6 regularly. Has a personal promptbook. Critiques drafts in new chats.
5 — Expert	Designs prompt templates others use. Coaches other team members. Contributes to the team promptbook.

Aim: the median team member reaches Level 3 within six weeks. Level 4 within six months.

12.5 The team promptbook

A shared promptbook is more valuable than ten personal ones in isolation. Suggested structure:

- One folder per team workstream (oncology brand planning, market access, medical, commercial ops)
 - Each prompt in the Section II format
 - Author named, date stamped
 - Quarterly review: prune the ones that have stopped working, promote the ones in heaviest use
-
-

Section 13. Master rules (the north star)

Pin these somewhere visible. They are the one-sentence distillations of every other section.

1. **Front-load:** audience, format, length, persona, constraints. Before you hit send.
 2. **One brief, not six messages:** paste the full picture in turn 1.
 3. **Challenge the maths, always:** "how did you calculate that?"
 4. **Quote or flag:** no claim without a source you can verify.
 5. **Two contexts beat one:** critique in a new chat, never the same one.
 6. **Save what works:** promptbook discipline compounds.
 7. **You are accountable:** never copy-paste raw output without re-moulding.
-
-

Section 14. When to push back on Claude

Hallucination tells to watch for:

- Suspiciously specific numbers with no source ("the European oncology supportive care market grew 14.3% in 2023")
- Citations that look right but you cannot find them in the source document
- Confident answers to questions where you would expect "it depends"
- "Industry standard" claims attributed to no specific industry body
- Polished writing about subjects where you suspect the model is reasoning from a thin base
- Round-number coincidences across multiple categories (every one of three options at "approximately 30%")

When you see those, respond with the Quote-the-Paragraph prompt (Section 6.3). If Claude cannot quote, mark the line, verify externally before using it. The phrase you want to internalise: *the polish of the output is not evidence of its accuracy*.

Appendix A. Quick-reference card

A pocket version of this document, suitable for printing as an A5 card and handing to team members in week 1.

THE 4Cs (Chugai PACE)

Clear What do I want, in what shape?
Context What does Claude need to know?
Constraints What must (and must not) be true?
Check How will I verify the output?

THE 30-SECOND PRE-SEND CHECK

1. Audience named?
2. Format and length specified?
3. Persona assigned?
4. Must-haves and must-not-haves listed?
5. Compliance lens flagged?
6. Uncertainty handling stated?
7. Disagreement handling stated?

THE MASTER RULES

1. Front-load.
2. One brief, not six messages.
3. Challenge the maths.
4. Quote or flag.
5. Two contexts beat one.
6. Save what works.
7. You are accountable.

THE THREE HABITS THAT MATTER MOST

1. Save your reusable context block in a Project.
 2. Run the 30-second pre-send check before send.
 3. Critique drafts in a new chat, not the same one.
-

Appendix B. The reusable context block (template)

Lift this, edit the personalisation, save it in a Claude Enterprise Project's custom instructions field.

You are a senior commercial strategist with deep oncology [or equivalent therapy area] launch experience.

I am [role] at Chugai [market]. I work on [products]. My audience defaults to [audiences] unless I say otherwise.

House style: [British English, plain language, no marketing buzzwords, show formulas in Excel, use the Chugai business case template, etc.]

Compliance lens: the ABPI Code of Practice (current edition) and UK Competition Law are non-negotiable. Flag risks proactively. [Add other relevant codes for international audiences.]

Working style: before drafting, ask up to 5 questions only if my answers would materially change the output. Otherwise proceed. Show your workings on any calculation. Cite sources or mark claims as inferred.

This is a living document. Update it as the practice evolves. The European Prompt Library and the wider Chugai Marketing AI capability start here.