



The AI Toolkit for Financial Advisors

Two Practical Guides. One Complete Toolkit.
Direct AI. Build Skills. Work Smarter.




BETTER
CLIENT OUTCOMES


SAVE
TIME


WORK
SMARTER


STAY
AHEAD

From IFA Academy · Panos Leledakis
Written in plain English for advisors around the world

Directing AI

Most advisors ask AI questions. This book teaches you to direct it — to organise the work, challenge the thinking, verify the result, and turn it into action.

From IFA Academy · Panos Leledakis

Written in plain English for advisors around the world

What is inside

From your first prompt to directing AI like a team.

1	The one idea in this book	3
2	The ladder	4
3	Why vague requests fail	6
4	RISEN: the five-part prompt	7
5	The Director's Six	9
6	The commands	11
7	Looping	16
8	The reverse prompt	19
9	The regulated professional's layer	20
10	The compliance gate	21
11	Print this page	22
12	What to do on Monday	23

START HERE

The one idea in this book

You are not talking to a search engine. You are briefing a very fast, very literal colleague.

Most advisors use AI the same way they use Google. They type a question. They read the answer. If the answer is thin, they blame the tool.

The tool is rarely the problem. The brief is.

When you give work to a new assistant, you do not just say "write something about retirement." You tell them who it is for, what it has to achieve, how to approach it, what to watch out for, and what you want back. You would never expect good work from a one-line instruction.

AI is the same. The difference is that AI will never tell you the brief was unclear. It will simply produce something average, confidently, and hand it to you.

IN ONE LINE

A good AI user asks questions. An excellent AI user **directs the work** — organising it, challenging it, verifying it, and turning it into action.

That shift is what this book teaches. It is not about memorising a hundred clever prompts. It is about learning perhaps twenty short instructions that change how the AI approaches any task at all.

Everything here works on any major tool — ChatGPT, Claude, Gemini, Copilot and the rest. The words are the skill. The tool will keep changing.

PART 1

The ladder

Five levels of working with AI. Almost everyone is stuck on the first one.

There is a natural progression in how professionals use AI. Most people never climb past the bottom rung, and then conclude the technology is overrated.

Level	What it is	How long it lasts
1. Ask	A prompt. One instruction, one time.	One answer
2. Direct	You tell the AI how to organise the work, how to think, and how to check itself.	One conversation
3. Loop	Draft, critique, revise, capture. You work the answer instead of accepting it.	One session
4. Save	A skill. A saved procedure that fires by itself when the task matches.	Forever
5. Contain	A project. A workspace holding your knowledge, your instructions and your context.	A whole part of your business

Levels 1 to 3 are this book. Level 4 is Book 2. Level 5 is Book 3.

WHY THIS MATTERS MORE THAN PROMPTS

An advisor with a great prompt saves ten minutes once. An advisor who has set up level 4 and level 5 saves ten minutes on every similar task, forever, without thinking about it. The prompt is the smallest part of the return.

Where most advisors actually are

Be honest about which rung you are on today. There is no wrong answer — there is only the next step.

If you are on level 1

You open a new chat, type a question, copy the answer. Your next step is Part 3 of this book: give the AI a role, context and a format. That alone will double the quality of what you get back.

If you are on level 2

You already write long, detailed prompts. Your next step is Part 6: stop accepting the first output. Learn to make the AI criticise its own work before you read it.

If you are on level 3

You go back and forth until it is right, then lose it. Your next step is page 20: the reverse prompt. Capture the win so you never repeat the work.

If you are on level 4 or 5

You have skills and projects running. Your next step is the second half of Part 5: orchestration, red-teaming and quality-control loops.

THE TWO LEVELS IN THIS BOOK

Every method here is labelled. **BASELINE** means you can do it today with no experience and it works. **GO FURTHER** means it rewards more effort and gives more back. Both are real. Neither is better. Start where you are, and climb when you are ready.

PART 2

Why vague requests fail

The AI is not reading your mind. It is filling in your blanks with the average of everything it has seen.

An AI language model predicts what comes next. When your instruction is vague, it has no reason to choose your version of the answer over the most common version. So you get the most common version. Generic, safe, slightly American, faintly hollow.

Look at the difference a few extra words make.

Weak	Strong
"Make this better."	"Make this shorter and warmer. Cut every sentence over fifteen words. Replace the general claims with one specific example."
"Write a post."	"Find the weak points in my argument and suggest how to strengthen them, then write the post."
"Write a title."	"Write five headlines: two provocative, two practical, one emotional."
"Is this good?"	"Score this out of ten on clarity, warmth and persuasion. Then give me three specific fixes."
"Fix my email."	"Rewrite this email so the client agrees to a meeting. Address their concern about cost. Keep it under 120 words."

THE PATTERN

Every strong version does one of three things: it names the **outcome**, it sets a **constraint**, or it tells the AI **how to think**. That is the whole game.

PART 3 · BASELINE

RISEN: the five-part prompt

If you learn nothing else in this book, learn these five letters.

RISEN is a simple checklist for building a complete prompt. It exists so you do not forget a part. Five letters, five questions.

R — Role

Who should the AI be? "You are a licensed insurance advisor with ten years of experience in family protection." This sets the vocabulary, the assumptions and the tone before anything else happens.

I — Information

What does it need to know? The situation, the audience, the background, the constraints. This is the part most people skip and it is the part that matters most.

S — Steps

How should it approach the work? "First summarise the situation. Then give two options. Then recommend one." You are writing the recipe, not just naming the dish.

E — Expectation

What is the goal? Not the task — the outcome. "The goal is that they feel understood and reply to book a meeting."

N — Narrowing

What are the limits? Length, tone, what to avoid, what to emphasise. "About three short paragraphs. Warm and professional. Do not go into policy details."

RISEN in action

Here is a real advisor task. A couple attended your financial wellness workshop and showed interest in life insurance. You want a follow-up email.

The weak version is: *"Write a follow-up sales email about life insurance."* You already know what that produces.

Now the same request, built with RISEN.

Follow-up email — RISEN

[COPY & CUSTOMISE](#)

You are a licensed insurance advisor with ten years of experience in family financial planning.

Write a follow-up email to a couple I met at a financial wellness workshop who showed interest in life insurance. They mentioned they worry about securing their children's future.

Include, in this order: a warm greeting that mentions the workshop, a short line showing I remember their specific concern, a brief explanation of how the right cover gives peace of mind, and an invitation to a short call.

The goal is that they feel valued and reply to book a meeting.

Keep it to about three short paragraphs. Warm and professional. Do not go into policy details or numbers. Never promise a guaranteed outcome.

NOTICE WHAT HAPPENED

You did not write a longer question. You wrote a **brief**. Role, situation, structure, goal, limits. The AI now has nowhere to be generic.

PART 4 · GO FURTHER

The Director's Six

RISEN builds a good prompt. This builds a directed piece of work.

RISEN is a formula for asking well. But for real professional work — a decision, an analysis, a plan — you need two things RISEN does not have: control over **how the AI thinks**, and control over **how it checks itself**.

That is what the Director's Six adds.

#	Element	What you are telling the AI
1	Objective	What you are actually trying to accomplish — not the task, the result.
2	Context	What it needs to know about your situation to be useful.
3	Orchestration	How to organise the work. In parallel, in steps, research first, compare after.
4	Thinking	How to examine the problem. Challenge me. First principles. Find the risk.
5	Quality control	How to verify. Separate fact from assumption. Show uncertainty. Critique yourself.
6	Output	Exactly what you want back, in what shape, at what length.

WHY ELEMENTS 4 AND 5 ARE THE WHOLE POINT

Every prompt framework covers objective, context and output. Almost none covers **thinking** and **quality control**. That gap is the difference between a chatbot and a colleague who tells you when you are wrong.

The Director's Six in action

Compare these two ways of asking the same question.

THE WAY MOST PEOPLE ASK

"Should I launch this new service?"

That gets you a balanced, sensible, useless paragraph. Now the directed version.

A directed decision

[COPY & CUSTOMISE](#)

I need to decide whether to launch [NEW SERVICE] in my practice.
My objective is [e.g. 20 qualified appointments a month within six months].

Context: [size of practice, market, licence, resources, timeline].

Treat market demand, competitor position, financial opportunity and implementation risk as four separate workstreams. Research current information where it matters. Then bring the findings together.

Challenge my assumptions and give me the strongest argument against launching this.

Separate facts from assumptions. Tell me what you are uncertain about and what I should verify myself.

Then give me your recommendation, and finish with the five actions I should take next, in order.

You did not ask a better question. You directed the entire thinking process — and you will get back something you can actually act on.

PART 5

The commands

Short instructions that change how AI approaches the work. Learn these, not a hundred prompts.

The rest of this part is a working vocabulary. Each command is one line. Each changes the shape of what comes back. They are grouped by what they control.

Group 1 · Orchestration — how the work is organised

BASELINE "Break this into steps before executing."

For anything with stages or dependencies. The AI plans first, then works. *"I want to build an AI marketing system for my practice. Break this into steps before executing."*

GO FURTHER "Run these as parallel workstreams."

For several independent tasks. It stops the AI mixing everything into one mush. *"Research the company, analyse its competitors, identify opportunities and review risks. Treat these as four parallel workstreams and bring the findings together at the end."*

GO FURTHER "Work backward from the desired outcome."

Starts from the result, not the activity. *"My objective is 20 qualified appointments per month. Work backward from that outcome and design the process."*

GO FURTHER "Show me your plan before you execute. Wait for my approval."

The most underrated command in this book. On any long task, it saves you twenty minutes of beautifully produced work in the wrong direction.

Group 2 · Thinking — how the problem is examined

BASELINE "Challenge my assumptions."

Also: "Tell me where my thinking may be wrong." "Do not agree with me automatically." "Give me the strongest argument against this idea." This single command turns AI from an agreement machine into a thinking partner.

GO FURTHER "Analyse this from first principles."

Tells the AI to ignore how the industry normally does it and look at the underlying facts. *"Forget how our industry traditionally does this. Analyse from first principles and propose a better approach."*

GO FURTHER "Red-team this idea."

Deliberately hunt for the weaknesses before reality finds them. Follow with: "What would cause this plan to fail?" "What risks am I underestimating?" "What would a sceptical competitor say?"

GO FURTHER "Act as three different experts and debate this in front of me."

One persona gives you one view. A panel gives you the argument. Name the three — for example a compliance officer, a top producer and a sceptical client — and let them disagree.

BASELINE "Go one level deeper."

Also: "Go beyond the obvious." "What am I missing?" "What would an expert notice that most people would overlook?" Use whenever an answer feels shallow.

Group 3 · Quality control — how the answer is verified

BASELINE "Separate facts, assumptions, interpretations and recommendations."

One of the most valuable commands for any professional decision. It shows you instantly what the AI actually knows versus what it is inferring.

BASELINE "Tell me what you are uncertain about."

Also: "Identify anything that needs verification." "Do not present assumptions as facts." AI can sound completely confident while being completely wrong. These instructions make the soft spots visible.

GO FURTHER "If your confidence is not high, say so and tell me exactly what to check myself."

Turns vague uncertainty into a to-do list you can act on.

BASELINE "Research this first and use current reliable sources before answering."

Essential for anything that changes: regulation, products, markets, tax, technology, companies, statistics. Add: "Verify anything that may have changed recently."

GO FURTHER "Answer only from the documents I gave you. If it is not there, say: not specified."

The single most effective line for preventing invented answers when you are working from your own files.

GO FURTHER "Critique your answer before finalising it, then improve it."

Also: "Review this for errors, weak assumptions, contradictions and missing information." This is a quality-control loop in one sentence. More on it in Part 6.

Group 4 · Shape — what you get back

BASELINE "Rank these from most important to least important."

Or: "Give me the 80/20 version first." AI loves long lists that tell you nothing about what matters. Force the priority.

BASELINE "Compare these using the same criteria."

Excellent for products, providers, policies, strategies and tools. Name the criteria yourself. *"Compare these three policies on price, what is covered, what is excluded, waiting periods and who each one suits."*

GO FURTHER "Give me three approaches: conservative, balanced and aggressive. Show me the trade-offs."

For any strategy, plan or decision. Three real options beat one recommendation.

BASELINE "Ask me only questions that would materially change your answer."

Stops the AI interviewing you for ten minutes before it does anything. The partner instruction: "If information is missing, make reasonable assumptions and clearly label them."

BASELINE "Do the work. Do not describe how you would do it."

Advisors lose hours to AI that explains its approach instead of producing the thing. Say this and it produces the thing.

BASELINE "Turn this into an action plan."

Also: "Give me the next five actions." "Identify the owner, the next action and the deadline." "What should I do first?" A good AI conversation should end in execution, not information.

Group 5 · Working style — the habits nobody teaches

BASELINE "Interview me. Ask me ten questions, one at a time."

Instead of you briefing the AI, the AI briefs itself. Superb for building a plan, a profile, a bio or a strategy when you do not yet know what you want to say. Add "one at a time" or you get all ten at once and answer none of them properly.

BASELINE "Here is a sample of my writing. Analyse my style and mirror it from now on."

Paste three things you actually wrote. Everything afterwards sounds like you instead of like a press release.

GO FURTHER "Define what an excellent answer looks like first. Then produce it. Then score it against your own criteria."

Setting the standard before the work is more powerful than criticising afterwards. This one command lifts quality more than any other in this book.

BASELINE "Give me this as a file I can send."

Ask for the deliverable, not the conversation. A document, a table, a one-page summary — something that leaves the chat.

BASELINE Context hygiene — start a fresh chat.

Long conversations get worse, not better. When a thread starts repeating itself or forgetting your rules, say: *"Summarise everything important from this conversation so I can continue in a new chat."* Paste that into a clean chat and carry on. This is the most practical habit in the book and almost nobody knows it.

PART 6

Looping

The first answer is a draft. Professionals work the draft.

Here is the most common mistake advisors make with AI: they read the first output, decide it is "quite good," and use it. The first output is almost never the best one available. It is simply the first.

Looping means working the answer. There are five loops worth knowing. All of them are short.

Loop 1 · The critique loop BASELINE

The simplest and most useful. Three moves.

- 1 Get the draft.
- 2 Say: *"Score this out of ten on clarity, warmth and persuasion. Explain each score."*
- 3 Say: *"Now rewrite it, fixing the two lowest scores."*

You will be surprised how honestly it marks its own work — and how much better the second version is.

Loop 2 · The variations loop BASELINE

"Give me five versions: two safe, two bold, one unexpected." Then: *"Take the structure of version two and the opening of version four, and combine them."* You are no longer editing. You are choosing.

Loop 3 · The layer loop [GO FURTHER](#)

"Improve this in three layers — structure, then clarity, then emotional impact. Show me each version separately."

Fixing everything at once produces mush. Fixing one thing at a time produces craft. This works especially well for a talk, a client letter or a piece of marketing.

Loop 4 · The devil's advocate loop [GO FURTHER](#)

This one is worth saving permanently. Paste it at the start of any important thinking session.

Make it argue with you

[COPY ALL](#)

From now on, do not simply confirm my opinions or assume my conclusions are correct. Your role is a thinking partner and an opponent of weak arguments, not a polite assistant who agrees.

Every time I give you an idea:

1. Analyse my assumptions. What am I taking for granted?
2. Give the counterargument. What would an intelligent, informed, well-intentioned sceptic say?
3. Check my logic. Are there gaps I am not seeing?
4. Offer alternative perspectives.
5. Prioritise truth, not agreement. If I am wrong, tell me clearly and kindly, with reasons.

Be demanding but constructive. If I drift into biased thinking or easy conclusions, point it out.

Loop 5 · The chain **GO FURTHER**

The output of one prompt becomes the input of the next. This is how one good answer compounds into a finished piece of work.

A REAL ADVISOR CHAIN

Prompt 1 builds your ideal client profile. Prompt 2 takes that profile and writes the LinkedIn content plan. Prompt 3 takes the plan and writes the first four posts. Prompt 4 takes the posts and checks them for compliance risk. Four prompts, one afternoon, a month of content.

And the loop that matters most: when to stop

Advisors who discover looping often over-loop. They polish for forty minutes to improve a client email by four percent.

THE STOP RULE

Two loops on most work. Three on something important. If the third loop is not clearly better than the second, you are finished — take the second and send it. Good and sent beats perfect and Thursday.

The exception is anything a client will see, anything with a number in it, and anything that touches compliance. Those get read by you, slowly, every time. Not looped — read.

PART 7

The reverse prompt

The bridge from this book to the next one. Do not skip this page.

You know the feeling. You asked. It answered. You corrected. You explained more. You fixed the tone. After four or five rounds, the result is exactly what you wanted.

Then you close the chat and it is gone. Next week you do the whole dance again.

SAY THIS INSTEAD

"Now that you understand exactly what I want, if I had to write this as one perfect prompt from the very start — what would that prompt be? Write it out."

The AI will hand you a clean, complete instruction that contains everything you fixed along the way. You did not have to guess the instructions. You **worked** until the result was right, then let the AI write the instructions for you.

That reverse prompt is now a permanent asset. You can save it in a note. Better, you can turn it into a skill so it fires automatically — which is exactly what Book 2 shows you how to do.

THE HABIT THAT COMPOUNDS

Every time you and AI finally get something right, spend thirty seconds capturing it. Over a year, that habit is worth more than every prompt library you will ever download — including ours.

PART 8

The regulated professional's layer

One extra instruction that belongs on anything client-facing.

Everything in this book applies to any professional. This page applies to you specifically, because you hold a licence.

For any task touching prospecting, marketing, client data, product comparison or automation, add this instruction:

The compliance layer

[COPY ALL](#)

Flag anything in your answer that may create compliance, privacy, regulatory, suitability, ethical or reputational concern.

Do not treat anything you produce as compliance, legal, tax, investment or professional approval. I am the licensed professional and I make the final decision.

Never state or imply a guaranteed return, a specific outcome, or past performance. If something needs a disclosure, tell me.

AND THE HARDER RULE

Do not place confidential client information, personal identifiers, health information or account information into any AI system unless your firm has specifically approved that system for that use. AI capability is not AI permission.

AI supports professional judgment. It does not replace it, and it does not carry your licence.

BEFORE YOU PRESS SEND

The compliance gate

This page appears in every book in this series. Read it once. Apply it always.

- **AI capability is not AI permission.** Because a tool can do something does not mean you are allowed to do it with real client data.
- **Use approved tools only.** Firm approval matters as much as technical security.
- **Give the minimum.** Before you paste, ask whether the AI actually needs that detail. Usually it does not.
- **De-identify first.** Changing the name is not enough. Replace exact ages, incomes, addresses and policy numbers with ranges.
- **Never paste** national ID or social security numbers, account numbers, passwords, credentials, identity documents or raw medical records.
- **A human reviews everything client-facing.** Always. You are the advisor. AI drafts; you decide.
- **Keep your records.** Marketing and client communications are usually regulated records, whoever drafted them.
- **Follow the strictest rule that applies** to your country, regulator, licence, firm and product provider.

BOOK 5

The full treatment — the traffic-light test for what you can paste, how to de-identify properly, and the questions to ask before you use any cloud AI tool — is in Book 5 of this series.

THE COMMAND CARD

Print this page

Every command in this book, on one sheet. Keep it next to your screen.

Organise the work

Break this into steps before executing. · Run these as parallel workstreams. · Work backward from the desired outcome. · Show me your plan before you execute, and wait for my approval. · Research first, then analyse.

Challenge the thinking

Challenge my assumptions. · Give me the strongest argument against this. · Do not agree with me automatically. · Analyse this from first principles. · Red-team this idea — what would cause it to fail? · Act as three different experts and debate this in front of me. · Go one level deeper. · What am I missing?

Verify the answer

Separate facts, assumptions, interpretations and recommendations. · Tell me what you are uncertain about. · Do not present assumptions as facts. · If your confidence is not high, say so and tell me what to check. · Research this first and use current reliable sources. · Answer only from the documents I gave you; if it is not there, say "not specified." · Critique your answer before finalising it, then improve it.

Shape the output

Rank these from most important to least. · Give me the 80/20 version first. · Compare these using the same criteria. · Give me three approaches: conservative, balanced, aggressive — with trade-offs. · Ask me only questions that would materially change your answer. · Do the work; do not describe how you would do it. · Give me this as a file I can send.

Loop it

Score this out of ten on [criteria], then fix the two lowest. · Give me five versions: two safe, two bold, one unexpected. · Improve this in three layers — structure, clarity, impact. · Interview me: ask ten questions, one at a time. · Define what an excellent answer looks like first, then produce it. · Summarise this conversation so I can continue in a fresh chat.

Capture it

Now that you understand what I want, write the one perfect prompt from the start.

Protect yourself

Flag anything that may create compliance, privacy, regulatory, suitability or reputational concern. Do not treat this as professional approval.

YOUR FIRST WEEK

What to do on Monday

Do not try all of this. Try five things.

- 1 **Print the Command Card.** Put it next to your screen, not in a folder.
- 2 **Pick one real task this week** — a client email, a policy comparison, a plan. Write it as a RISEN brief instead of a question.
- 3 **Use the critique loop once.** Ask for a score out of ten, then ask for the fix. Notice the gap between draft one and draft two.
- 4 **Use "challenge my assumptions" on a real business decision.** Not a test question. Something you are actually deciding.
- 5 **Capture one reverse prompt.** The next time something comes out exactly right, ask for the perfect prompt and save it. That is your first skill, waiting for Book 2.

THE HONEST PROMISE

You will not remember thirty commands. Nobody does. You will remember four or five, and those four or five will change your week. That is enough. Come back for the rest when you need them.

Simple word list

The few new words in this book, in plain English.

Prompt — What you type. One instruction, one time.

Directing — Telling the AI how to organise the work, how to think, and how to check itself — not just what to produce.

RISEN — A five-part checklist for a complete prompt: Role, Information, Steps, Expectation, Narrowing.

The Director's Six — Objective, Context, Orchestration, Thinking, Quality control, Output.

Looping — Working the answer instead of accepting it. Draft, critique, revise.

Red-teaming — Deliberately attacking your own idea to find its weaknesses first.

Grounding — Telling the AI to answer only from the documents you gave it.

Reverse prompt — Asking the AI to write the perfect prompt after you have finally got the result right.

Context hygiene — Starting a fresh chat when a long conversation starts getting worse.

Skill — A saved procedure the AI uses automatically. See Book 2.

Project — A workspace holding your files, instructions and context. See Book 3.

This guide is from IFA Academy for learning and educational purposes only. It is not legal, compliance, tax, investment or financial advice. AI tools change frequently — menu names, settings and capabilities described anywhere in this series should be checked against your own tool today. Everything AI produces is a first draft. You remain the licensed professional and you remain responsible for anything that reaches a client. Requirements differ by country, regulator, licence, firm and product provider. Always follow your own firm's compliance procedures and the law that applies to you.

© 2026 IFA Academy · #AdaptAndFly · Stop asking. Start directing.

Skills: Teach It Once

Save your method once. Never re-explain it again. What skills are, how to make them with no code, and eight ready-made skills to copy and use today.

From IFA Academy · Panos Leledakis

Written in plain English for advisors around the world

What is inside

Everything you need to build your first skill this week.

1	What is a skill?	3
2	Prompt, skill or project?	5
3	Why a skill beats a saved prompt	6
4	How a skill is built	7
5	Two easy ways to make one	9
6	The reverse prompt method	10
7	When NOT to build a skill	11
8	Best practices	12
9	The compliance line	13
10	Skill ideas for advisors	14
11	Eight ready-to-paste skills	15
12	Make your own in five steps	20
13	Simple word list	21

PART 1

What is a skill?

The simplest idea: a saved procedure that Claude uses on its own.

You know how you explain the same thing again and again? "Write it in my voice. Short sentences. No hard words. Add a friendly close." Every time. A skill ends that.

A skill is like an onboarding guide for a new team member. You write down how you want a job done — once. From then on, the assistant follows it. You do not have to repeat yourself.

The best part: you do not have to call the skill by name. When your task matches the skill, it kicks in automatically. You just ask for the work, and your method is already applied.

This is the difference between a tool and a team member. A tool waits for full instructions every time. A team member already knows how you work. A skill turns the tool into the team member.

IN ONE LINE

A skill is your method, made permanent. You teach it once. It is used every time — on its own.

Nothing here is code. A skill is mostly plain words, written the way you would explain the job to a new assistant on their first morning.

Where skills work, and what they change

Build a skill once, and it works wherever you work with Claude: the normal chat, the desktop app, and the developer tools. You do not rebuild it for each place.

Think about what you actually repeat in a week. The way you write a client reply. The way you lay out a policy comparison. The way you turn messy meeting notes into a clean summary. Each one is a procedure that lives in your head. Right now you re-explain it every single time.

THE REAL WIN

You stop spending the first five to ten minutes of every chat explaining what you want. Over a week, that is hours. Over a year, it is weeks.

Your first move

BASELINE Write down the one task you explain most often. Just the name of it. "Client email replies." "Policy comparisons." "Meeting notes." That single line is the skill you will build in this guide.

GO FURTHER Open your sent folder. Find three real examples of that task where you were happy with the result. Save them in one document. You will paste them into your skill later as examples. Examples are followed more faithfully than rules.

Do not build ten skills this week. Build one, use it for two weeks, then build the next.

PART 2

Prompt, skill or project?

Three tools. Three different jobs. Here is how to choose.

Advisors often confuse these three. They are not competitors. They do different work, and using the wrong one is why people give up.

Use	What it is	Choose it when
Prompt	A one-time ask, typed in the chat.	The task is new, rare, or you only need it once. Fastest to start.
Skill	A saved procedure that fires automatically when the task matches.	You do the same job again and again, and you always want it done your way.
Project	A bounded workspace with knowledge files, instructions and memory of past chats.	The work needs your documents as the source of truth, and several chats over time.

They stack. Inside a project, your skills still fire, and you still write prompts. A project holds the knowledge. A skill holds the method. A prompt holds today's request.

ABOUT OTHER PLATFORMS

Skills as described here are a Claude capability. Other AI platforms are moving toward the same idea with their own names and their own menus. Do not assume a feature exists because you read about it. Check your platform's current settings before you plan your workflow around it.

TRY IT

BASELINE Take your list of repeated tasks. Mark each one P, S or K — prompt, skill, project.

GO FURTHER For anything marked K, list the three documents that project would need to hold.

PART 3

Why a skill beats a saved prompt

Many advisors keep prompts in a notes app. A skill is far better. Here is why.

Most people who use AI have "a prompt they keep." They save it in Notes. They copy it, paste it, change a few words, and press enter. It works — but every step is manual, and manual steps get skipped when you are busy.

A saved prompt (old way)	A skill (new way)
You must find it every time	It is found for you
You copy and paste it	Nothing to paste
You must remember to use it	It kicks in on its own
Lives in a notes app	Lives inside the assistant, ready
Improves only when you edit the note	Improves once, for every future task

There is a second, quieter win. A saved prompt is used only by the person who saved it. A skill can be shared with your team, so everyone produces the same standard of work. Your best process stops being personal and becomes the house standard.

THE TEST THAT SETTLES IT

If you have ever thought "I know I have a good prompt for this somewhere" — that is the moment a skill would have saved you. The prompt you cannot find has no value.

PART 4

How a skill is built

A skill has three simple parts. You already understand all of them.

Do not worry — a skill is not code. It is mostly plain words. It has three parts, and only one of them is difficult to get right.

PART 1 — THE NAME

A short name, in small letters with dashes between words. Example: replying-to-clients. That is all.

PART 2 — THE DESCRIPTION (THE MOST IMPORTANT PART)

This is how the assistant decides when to use the skill. Say what it does and when to use it. Add the real words you would actually say. Example: "Use this when I need to reply to a client email, write back, or draft a message."

PART 3 — THE INSTRUCTIONS

The steps you want followed. Plus one example of a good result. This is where you put your voice, your rules, your format and your limits.

That is the whole structure. Name, description, instructions. Everything in this guide is a variation of those three parts.

What a skill looks like

Here is a complete, working skill. It is short on purpose. Read it once and you will see there is nothing technical in it.

a simple skill

PLAIN TEXT

```
--- the top part: name and description ---  
name: replying-to-clients  
description: Use this when I need to reply to a client or  
prospect email, write back, or draft a message.  
--- then the instructions ---  
  
# How to reply  
1. Read the email I paste.  
2. Reply in a warm but direct voice. Short sentences.  
3. End with a clear next step.  
  
Example of a good close: "Happy to help – shall we talk Tuesday?"
```

THE DESCRIPTION IS THE MAGIC

If the description is clear and contains the words you really use, the skill will fire at the right moment. If it is vague, the skill sits there and never wakes up. Most failed skills are failed descriptions. Spend your time here.

BASELINE Read the description above out loud. Would it match the way you ask for an email reply?

GO FURTHER Add five more trigger words you personally use — "respond", "answer this", "come back to him", and so on.

PART 5

Two easy ways to make one

No code. Really. Pick the way that feels right for you.

First, check that the feature is switched on. Look in Settings, under the section for capabilities or features, and enable code execution and file creation. Menu names change, so look for the nearest match on your account.

BASELINE Way 1 — let the assistant build it for you

There is a built-in helper called the skill-creator. It builds skills for you. You only talk.

- 1 Say: "Use the skill-creator to help me build a skill for replying to client emails."
- 2 It asks you questions. Your role. Your audience. Your tone. Examples of good and bad. Answer honestly and in your own words.
- 3 In a few minutes you get a ready skill. Review it. Ask for changes. Save it. Done.

WAY 2 — PASTE IT IN YOURSELF

- 1 Go to the Customize or Capabilities area and open Skills. Choose to create a new skill.
- 2 Paste your skill text — for example, one of the eight ready-made skills later in this guide.
- 3 Turn it on. Open a new chat. Ask for something the skill is made for, and watch it fire.

WHICH ONE FOR YOU?

New to this? Use Way 1. It does the hard thinking for you. Use Way 2 when you already have the text.

PART 6

The reverse prompt method

This one is gold. It turns a good chat into a perfect skill.

Here is the trick that bridges the gap between looping in a chat and owning a skill. Use it the moment you and the assistant finally "get it right" after going back and forth.

You know the feeling. You ask. You get an answer. You explain more. You fix the tone. You fix the format. After a few rounds, the result is exactly what you wanted. Most people copy that result and close the chat. The instructions are lost. Next week they start from zero again.

SAY THIS

"Now that you understand exactly what I want, if I had to write this as one perfect prompt from the very start — what would that prompt be? Write it out."

You will get back a clean, complete prompt that contains everything you corrected along the way — including the things you never managed to put into words. That is your instruction set, written for you.

THEN TURN IT INTO A SKILL

Say: "Great. Now turn that into a skill using the skill-creator." From then on, you get that result every time, with no back-and-forth.

BASELINE Use the reverse prompt on your next good chat. **GO FURTHER** Before you save it, ask: "What would make this fail? Add rules to prevent that."

PART 7

When NOT to build a skill

Skills are powerful. They are also easy to overuse. Four times to say no.

A badly chosen skill costs you more than it saves. It fires when you did not want it, or it locks in a method you have already outgrown. Before you build, check these four cases.

If the task is...	Do this instead
A one-off. You will do it once, maybe twice a year.	Write a normal prompt. Building takes longer than doing.
Different every time. No stable steps, no fixed format.	Keep it in the chat. A skill needs a repeatable procedure.
Dependent on current facts. Rates, tax rules, product terms, prices.	Use search or a project with up-to-date documents. A skill freezes what you wrote on the day you wrote it.
Still unclear. You do not yet know what "good" looks like.	Do it manually five times first. Then use the reverse prompt.

THE FROZEN-FACT TRAP

Never write figures, limits, premiums or tax thresholds inside a skill. They will be wrong within months, and they will keep being used with total confidence. Skills hold your method. Documents hold your facts.

PART 8

Best practices

Six simple rules. Follow these and your skills will fire when you need them.

- 1. One skill, one job.** Do not build one giant skill. Small, focused skills work better. One for emails. One for policy comparisons. One for notes.
- 2. Write a strong description.** This is what makes the skill kick in. Add the real words you would say — "reply", "email", "follow up". Say what it does and when to use it.
- 3. Give an example.** Show one good result inside the skill. A good example is copied more faithfully than a long list of rules.
- 4. Explain why, do not just command.** Instead of "ALWAYS use short sentences", say "use short sentences, because my clients are busy and read on the phone". Reasons are followed better than rigid rules.
- 5. Start simple. Test. Improve.** Begin with a few lines. Try it on real work. If it does not fire, make the description clearer. Skills are living documents.
- 6. Only use skills you trust.** Use your own, or ones from a source you trust. Never paste in a skill from a stranger without reading every line — it is instructions you are handing to your assistant.

A GOOD RHYTHM

Review your skills once a quarter. Delete the ones you never use. Improve the two you use daily.

PART 9

The compliance line

The skill shapes the writing. You still carry the responsibility.

This page matters more than any other in the guide. A skill makes output faster and more consistent. It does not make the output true, suitable, or approved.

YOU ARE STILL THE ADVISOR

Read every output before it reaches a client, a prospect, a regulator or a public feed. Speed is not a defence. If it goes out under your name, it is yours.

Never build a skill that...

- Promises guaranteed returns, or any specific investment result.
- Uses past performance as a promise about the future.
- Gives personal recommendations to an audience you have not assessed.
- Removes the disclosures your firm requires.
- Writes testimonials, or presents client results without the full context.
- Stores client names, numbers or health details inside the skill text.

That last one catches people. A skill is a document. Do not put a real client inside it. Use placeholders — [client name], [policy], [amount] — and paste the real details only in the chat, following your firm's data rules.

BEFORE YOU SHARE A SKILL

If you share a skill with your team, you are sharing a process that will be used without a second check. Have it reviewed the way you would have a template reviewed.

PART 10

Skill ideas for advisors

A menu worth building. The ones marked READY are written out in this guide.

Skill	What it does	Status
Client email replies	Reply in your voice, warm and clear.	READY
Policy comparison	Turn two or three policies into one clear table.	READY
Meeting notes to summary	Turn messy notes into a clean client summary.	READY
Follow-up messages	Write kind, short follow-ups that get replies.	READY
Prospect research brief	A short, useful brief before a first meeting.	READY
LinkedIn and social posts	A post in your voice, safe and on-brand.	READY
Content compliance check	Flag risky claims before you post or send.	READY
Annual review prep	Build the agenda and questions for a review.	READY
Referral request	Ask for a referral without feeling pushy.	IDEA
Plain-English quote explainer	Explain a quote or proposal in simple words.	IDEA
Claim support letter	Help a client through a claim, step by step.	IDEA

BASELINE Pick one READY skill and install it today. **GO FURTHER** Turn one IDEA row into a real skill using the reverse prompt method.

PART 11

Eight ready-to-paste skills

Copy each one. Paste it into Create skill, or hand it to the skill-creator.

Each skill below is plain text. The first lines are the name and the description. The rest are the instructions. Copy the whole block, then change the details so it sounds like you.

Skill 1 — Client email replies

replying-to-clients

[COPY ALL](#)

name: replying-to-clients

description: Use this whenever I need to reply to a client or prospect email, write back to someone, respond to a message, or draft a follow-up. Trigger words: reply, respond, write back, email, follow up.

How to write my replies

1. Read the email I paste.
2. Write a warm but direct reply. Short sentences. Easy words.
3. No hype. No hard promises. Never guarantee returns.
4. End with one clear next step (a call, a date, a document).
5. Keep it short – under 120 words unless I ask for more.

Example close: "Happy to help. Are you free Tuesday at 3pm for a quick call?"

Remember: I am the advisor. I read every reply before I send it.

Skill 2 — Policy comparison table

comparing-policies

[COPY ALL](#)

name: comparing-policies

description: Use this when I need to compare insurance policies, quotes, or plans, or when I paste 2–3 policy documents and want a clear comparison. Trigger words: compare, policy, quote, plan, difference.

How to compare

1. Read each policy or quote I give you.
2. Make one simple table. One column per policy.
3. Rows to include: price, what is covered, what is NOT covered, waiting periods, limits, and the biggest difference.
4. Below the table, write 3 short lines: who each plan fits best.
5. Use plain words a client can understand. No jargon.

Do not recommend one as "best" unless I ask. Just compare clearly.

Skill 3 — Meeting notes to client summary

summarizing-client-meetings

[COPY ALL](#)

name: summarizing-client-meetings

description: Use this when I paste messy notes or a transcript from a client meeting and want a clean summary. Trigger words: notes, meeting, summary, recap, call.

How to summarize

1. Read my notes.
2. Write a one-page summary with these parts:
 - What the client wants (their goals)
 - Their situation (family, money, risks)
 - What we agreed
 - Next steps (with who does what)
3. Keep it short and clear. Use simple words.
4. At the end, list any documents I still need from the client.

Skill 4 — Follow-up messages

writing-follow-ups

[COPY ALL](#)

name: writing-follow-ups

description: Use this when I need to follow up with a client or prospect who has not replied, or nudge someone politely. Trigger words: follow up, nudge, remind, check in, chase.

How to follow up

1. Ask me (or read) who I am following up with and why.
2. Write a short, friendly message. 2–4 sentences.
3. Be kind, not pushy. No guilt. No pressure.
4. Remind them of the value, then give one easy next step.
5. Offer two time options for a call when it fits.

Example: "Hi [name], just checking in. Still happy to walk you through the options whenever you're ready. Tuesday or Thursday?"

Skill 5 — Prospect research brief

researching-prospects

[COPY ALL](#)

name: researching-prospects

description: Use this before a first meeting when I give a person's or company's name and want a short brief. Trigger words: research, prospect, brief, prepare, first meeting.

How to prepare my brief

1. Use what I give you (and web search if it is on).
2. Write a short one-page brief:
 - Who they are, in 2 lines
 - What may matter to them (needs, risks)
 - 3 smart questions I could ask
 - 1 thing to avoid
3. Keep it factual. Do not guess private details.

If you are not sure about a fact, say so. Do not invent.

Skill 6 — LinkedIn and social posts

writing-linkedin-posts

[COPY ALL](#)

name: writing-linkedin-posts

description: Use this when I want a LinkedIn or social post about insurance, financial planning, or AI for advisors. Trigger words: post, LinkedIn, social, content, write a post.

How to write my posts

1. Start with a strong first line that makes people stop.
2. Use short lines. Simple words. One idea per post.
3. Motivational, with light, safe humor. No hype.
4. End with a question or a small call to action.
5. Never promise guaranteed returns or specific results.

Keep it safe for compliance. I review every post before I post it.

Skill 7 — Content compliance check

checking-content-compliance

[COPY ALL](#)

name: checking-content-compliance

description: Use this before I post, send or publish anything – a social post, an email, a newsletter, an ad or a client letter. Trigger words: check this, compliance, is this safe, review before I post, risky.

How to check my content

1. Read the text I paste. Do not rewrite it yet.
2. Flag every risky item, quoting the exact words:
 - Guaranteed returns, or any promised result
 - Past performance used as a promise about the future
 - "Best", "safest", "no risk", "always", "never"
 - Testimonials or client results without context
 - Missing disclosure (issuer, risks, fees, licence)
 - Personal advice given to a general audience
3. For each flag, say in one line why it is a risk.
4. Then give me one safer rewrite of the whole text.
5. End with: "Advisor to confirm against firm and local rules."

You are not my compliance officer. You help me spot problems early.

Skill 8 — Annual review preparation

preparing-annual-reviews

[COPY ALL](#)

name: preparing-annual-reviews

description: Use this when I am preparing for a client annual review or a yearly check-in and need an agenda and questions. Trigger words: annual review, yearly review, check-in, review meeting, prepare the agenda.

How to prepare my review

1. Use the notes or policy list I paste. Ask me for anything missing.
2. Build a one-page agenda in this order:
 - What changed this year (family, job, income, health, goals)
 - Cover and plans in force today, in plain words
 - Gaps or points worth checking
 - The client's own questions and concerns
 - Agreed next steps, with who does what and by when
3. Add 8 open questions I can ask. Simple, not technical.
4. Add a short list of documents to bring or to request.
5. One page. No jargon. No product pitch inside the agenda.

Do not assume facts about the client. If something is missing, ask me.

MAKE THEM YOURS

These are starting points, not finished products. Change the words to match your voice, your market and your rules. Then test each one on a real task before you rely on it.

PART 12

Make your own in five steps

When you need a skill for something not in this guide, do this.

- 1 **Pick your most-repeated task.** The one you re-explain the most. That is your best first skill — not the most interesting one, the most repeated one.
- 2 **Do it once, well, in a chat.** Go back and forth until the result is exactly right. Do not settle for "close enough" here.
- 3 **Use the reverse prompt.** Ask: "Write the perfect prompt for this, from the start."
- 4 **Turn it into a skill.** Say: "Now turn that into a skill with the skill-creator." Read every line before you save it.
- 5 **Test and improve.** Use it on real work for two weeks. If it does not fire, the description is too vague. If the output drifts, add one example.

THE HABIT THAT MAKES IT STICK

Build one skill this week. Use it for two weeks. Then build the next. In three months you will have a small set of skills doing your admin work — and you will not remember how you worked without them.

BASELINE Build one skill from this guide today. **GO FURTHER** Write a short standard for your team: naming, review, and who approves a shared skill.

GLOSSARY

Simple word list

The few new words in this guide, in plain English.

Skill — A saved set of instructions for a repeated job. It is used automatically when your task matches it.

Description — The part of a skill that says what it does and when to use it. It decides whether the skill fires. The most important part.

skill-creator — A built-in helper that interviews you and builds the skill for you. No code needed.

SKILL.md — The plain text file behind a skill. You rarely need to touch it by hand.

Reverse prompt — Asking for the perfect prompt after you have finally got the result right in a chat. The fastest way to write good instructions.

Trigger words — The real words you would say for that task, placed in the description so the skill kicks in at the right moment.

This guide is from IFA Academy for educational purposes only. It is general guidance, not legal, compliance, tax or financial advice. Product features, menu names and settings change often — always check your platform's current settings and help pages before you rely on a feature. The ready-made skills are starting points: test them, adapt them to your own voice, and check them against your firm's compliance rules and your local regulations before using them with real clients. You remain responsible for every output that reaches a client.