

C · T · R · L — A FRAMEWORK FOR DELEGATING TO AI AGENTS

# The CTRL Framework.

For leaders who are done asking AI questions and ready to hand it work. A plain-language method for delegating to agents without losing control.

15 SECTIONS · ~15 MIN READ

CONTENTS

# What's inside.

Four controls. One framework. A short reference for handing real work to an AI agent without losing the wheel.

– *Foundations*

---

01	<b>Introduction</b>	03
	DELEGATION ISN'T PROMPTING	
02	<b>The CTRL Framework</b>	04
	FOUR CONTROLS, AT A GLANCE	

---

– *The Four Controls*

---

03	<b>C Context</b>	05
	SET THE WORLD THE AGENT WORKS INSIDE	
04	<b>T Task</b>	06
	DEFINE WHAT "DONE" MEANS	
05	<b>R Rules</b>	07
	KEEP THE HUMAN IN AUTHORITY	
06	<b>L Loop</b>	08
	TELL THE AGENT HOW TO WORK, CHECK, AND STOP	

---

– *In Practice*

---

07	<b>Agent vs Chatbot</b>	09
	THE RELATIONSHIP HAS CHANGED	
08	<b>The Handoff Card</b>	10
	FOUR QUESTIONS FOR ANY TASK	
09	<b>Weak vs Strong Delegation</b>	11
	SIDE-BY-SIDE EXAMPLES	
10	<b>Human Review Points</b>	12
	WHERE JUDGEMENT BELONGS	

---

– *Closing*

---

11	<b>The Impact</b>	13
	HOW EACH CONTROL AFFECTS OUTCOMES	
12	<b>Next Steps</b>	14
02	WHERE YOU GO FROM HERE	• AIBRAIN.COACH
13	<b>Credits</b>	15
	CONTRIBUTORS & CONTACT	

## 01 INTRODUCTION

# Delegation isn't prompting.

Most leaders learned AI by prompting: ask better questions, get better answers. That worked when AI was mostly a conversation partner.

Agentic AI changes the relationship. An agent doesn't just answer. It plans, searches, compares, drafts, revises, calls tools, triggers workflows, and moves through a task across many steps.

You are no longer asking. You are *delegating* — and delegation needs a different discipline.

*CTRL* is that discipline. A structured, repeatable system for handing work to AI agents without losing the wheel. Four controls, written in plain language. No tooling expertise required.

## THE THESIS

“ *AI agents don't need better vibes. They need better handoffs.* ”

## WHAT CHANGES

From asking for answers to delegating real work.

## WHAT IT COSTS

About thirty seconds more clarity before the agent moves.

## COMPANION TO PROMPT

PROMPT helps you speak to LLMs. CTRL helps you direct agents. Use both — one for asking, one for assigning.

## READ THIS GUIDE IF

You're a leader, operator, founder, consultant, or assistant whose AI tools are starting to act on their own — and you want them to act on *your* terms.

## TIME ON PAGE

~15 min, cover to cover

02 THE FRAMEWORK

# Four controls for agentic work.

A prompt asks for a response. A CTRL handoff assigns work. Four controls cover what the agent needs to know, what it must complete, what it must obey, and how it should check, correct, and stop.

<p><b>C</b></p> <p>CONTEXT</p> <p><b>What does the agent need to know?</b></p> <p>Agents fail when they work without real-world background. Give them the situation, audience, goal, files, prior decisions, and constraints that shape the work.</p>	<p><b>T</b></p> <p>TASK</p> <p><b>What must the agent complete?</b></p> <p>Agents need a finish line. Define the deliverable, the outcome, and what "done" looks like — otherwise the agent produces activity instead of completion.</p>
<p><b>R</b></p> <p>RULES</p> <p><b>What must the agent obey?</b></p> <p>Rules cover permissions, limits, tone, privacy, tools, approvals, brand, legal boundaries, and escalation. This is where you keep authority.</p>	<p><b>L</b></p> <p>LOOP</p> <p><b>How should it proceed and stop?</b></p> <p>Tell the agent whether to plan first, ask before acting, verify sources, work in stages, or stop for approval. This is what prevents runaway execution.</p>

Each control earns its letter. Skip one and the agent fails in a predictable way. The next four pages take them one at a time.

• TURN THE PAGE →

# C

01 / 04 • CONTEXT


## Set the world the agent works in.

Tell the agent what's real before it starts. The more autonomy it has, the more dangerous vague context becomes.

CONTROL 01 • CONTEXT

A chatbot can survive a vague question. An agent can't. Context is the operating environment — business goal, audience, situation, files, prior decisions, preferences, market, budget, stakeholders, risks, tone. Without it the agent guesses. With it, the agent works inside reality.

THE MOVE	HOW THE AGENT REACTS	EXAMPLE HANDOFFS
<p><b>No context given</b> DEFAULT STATE</p>	<p>The agent fills gaps with assumptions — and may optimise for the wrong audience, goal, or business reality entirely.</p>	<p>● WEAK</p> <p>"Help me prepare for a client meeting."</p>
<p><b>Basic context provided</b> SITUATED</p>	<p>The agent understands the assignment well enough to produce relevant work.</p>	<p>● GOOD</p> <p>"I'm meeting a mid-sized Canadian manufacturer about AI adoption. They're cautious, non-technical, and worried about cost. Prepare a one-page briefing I can use to guide the conversation."</p>
<p><b>Operational context provided</b> EXECUTION-READY</p>	<p>The agent understands the goal, constraints, audience, assets, and decision environment before acting.</p>	<p>● BETTER</p> <p>"Same brief — likely use cases are sales admin, reporting, and customer service. Position me as practical, not hype-driven. Include three use cases, likely objections, and smart questions to ask."</p>



02 / 04 • TASK

# Define what "done" looks like.

A task is not a topic, an intention, or "look into this." It is an outcome the agent can complete and you can use.

CONTROL 02 • TASK

Agents do work. Work needs a finish line. The task tells the agent what to produce, for whom, in what form, and by when — and makes completion obvious. If the agent can't tell whether it's done, the task isn't clear enough. Good tasks turn an energetic intern into a useful operator.

THE MOVE	HOW THE AGENT REACTS	EXAMPLE HANDOFFS
<b>Vague task</b> DEFAULT STATE	The agent produces activity: notes, options, summaries — output without a clear deliverable.	● WEAK "Research AI tools."
<b>Defined deliverable</b> OUTCOME-LED	The agent knows what to produce and can stop when it's complete.	● GOOD "Compare three AI scheduling tools for a five-person consulting team."
<b>Decision-ready task</b> BUSINESS-READY	The agent produces something a human can review, approve, send, publish, or use.	● BETTER "Same comparison — return a table with pricing, core features, integrations, risks, and a recommendation for which one to test first."

# R

03 / 04 • RULES

## Set the boundaries before it moves.

Rules are how humans stay in authority. Without them, the agent decides its own limits — and may oversell, overshare, or act too soon.

CONTROL 03 • RULES

Rules define what the agent can do, can't do, must avoid, and when it must escalate. Tone, brand, budget, privacy, sources, legal limits, approval requirements, tools, geography, risk tolerance. This is the difference between *delegation* and *abdication*.

THE MOVE	HOW THE AGENT REACTS	EXAMPLE HANDOFFS
<p><b>No rules</b> DEFAULT STATE</p>	<p>The agent decides its own limits and frequently oversteps.</p>	<p>● WEAK</p> <p>"Draft replies to these client emails."</p>
<p><b>Basic rules</b> CONSTRAINED</p>	<p>The agent avoids the obvious mistakes and keeps to the surface of what it's allowed to do.</p>	<p>● GOOD</p> <p>"Draft replies. Do not promise pricing, timelines, discounts, or legal terms."</p>
<p><b>Authority rules</b> CONTROLLED EXECUTION</p>	<p>The agent knows what it can do alone, what needs approval, and what must never happen.</p>	<p>● BETTER</p> <p>"Draft replies in a calm, professional tone. Don't promise pricing, timelines, discounts, legal terms, or technical capabilities. Flag anything involving money, contracts, complaints, or risk for human review."</p>



04 / 04 • LOOP

# Tell it how to work, check, and stop.

The loop is the operating rhythm. Without it the agent rushes, drifts, overproduces, or keeps working past the useful point.

CONTROL 04 • LOOP

Agents can move across many steps — useful when the loop is clear, risky when it isn't. The loop tells the agent whether to plan first, ask before acting, work in stages, verify sources, revise after feedback, or stop at a defined output. A strong loop turns a task into a controlled workflow.

THE MOVE	HOW THE AGENT REACTS	EXAMPLE HANDOFFS
<p><b>No loop</b> DEFAULT STATE</p>	<p>The agent may rush, drift, overproduce, or keep working beyond the useful point.</p>	<p>● WEAK</p> <p>"Create a launch plan."</p>
<p><b>Review loop</b> CHECKPOINTED</p>	<p>The agent pauses at key moments for human approval before continuing.</p>	<p>● GOOD</p> <p>"Create a launch plan. First give me the structure. Wait for approval before writing the full plan."</p>
<p><b>Quality loop</b> SELF-CORRECTING</p>	<p>The agent checks its own work against the task and rules before returning the output.</p>	<p>● BETTER</p> <p>"...then review the draft against audience, budget, timeline, and risks. Stop after the final draft and list the three decisions I still need to make."</p>

03 IN PRACTICE · 01

# The relationship has changed.

A chatbot answers. An agent acts. That single shift changes the skill the human needs to bring. Casual works for chatbots. Agents need a deliberate handoff.

A CHATBOT CONVERSATION

## Asking.

The main risk is a weak answer. Stakes are low; revision is cheap.

"Explain this."

"Summarize this."

"Give me ideas."

"Rewrite this."

"Answer this question."

• PROMPTING

AN AGENT HANDOFF

## Acting.

The risk is weak execution: wrong steps, wrong assumptions, wrong tools, wrong audience — work that looks complete but fails in practice.

"Research this, compare the options, prepare a recommendation, draft the email, check for risks, and stop before sending."

• DELEGATION

*“ The more autonomy the system has, the more deliberate the handoff has to be.*

CTRL EXISTS FOR THE SECOND WORLD.

# Four questions anyone can use.

CTRL should fit on a sticky note. Brief the agent the way you'd brief a competent assistant — no orchestration jargon required.

- C** **What does it need to know?**  
Context — the situation, audience, files, and constraints.

---

- T** **What must it complete?**  
Task — the deliverable, in plain language.

---

- R** **What is it not allowed to do?**  
Rules — limits, brand, privacy, escalation.

---

- L** **When should it check back or stop?**  
Loop — the checkpoints, review points, and finish.

HANDOFF CARD
EXAMPLE

## A one-page executive briefing

**CONTEXT**  
Canadian executive audience. Interested in AI but sceptical of hype.

**TASK**  
A one-page briefing on how agentic AI will affect administrative work in 2026.

**RULES**  
Plain English. No technical jargon. No legal, HR, or financial claims. Don't exaggerate.

**LOOP**  
Outline first. Wait for approval. Then full draft, with three risks to review.

The user doesn't need to understand agents, orchestration, memory, retrieval, or tool calling. They only need to brief the machine the way they would brief a competent assistant.

03 IN PRACTICE · 03

# Most failures begin with lazy handoffs.

Weak delegation sounds efficient. It isn't. When you say "handle this," the agent invents the missing structure — context, priorities, boundaries, finish line. That's where the errors begin.

• WEAK

"Plan my content."

• STRONG

"Create a five-day LinkedIn content plan for Canadian executives who are curious about AI but not technical. Position me as a practical AI advisor. Avoid hype, fear, and jargon. Give me post topics first; wait for approval before writing the posts."

• WEAK

"Deal with these emails."

• STRONG

"Sort these emails into urgent, waiting, informational, and ignore. Draft replies only for urgent and waiting. Don't send anything. Flag anything involving money, contracts, complaints, or personal data."

• WEAK

"Research competitors."

• STRONG

"Find five Canadian AI consulting firms serving mid-market businesses. Compare positioning, services, pricing signals, proof points, and gaps. Return a table and a short recommendation on how we should differentiate."

*“ Strong delegation isn't longer instructions. It's complete instructions.*

WEAK DELEGATION CREATES  
CLEANUP.  
STRONG DELEGATION  
CREATES USABLE WORK.

03 IN PRACTICE · 04

# Where judgement still belongs.

CTRL isn't about letting agents run wild. It's about knowing where human review belongs. Decide before the agent starts.

<p><b>REVIEW BEFORE ACTION</b></p> <p><b>Human approves first.</b></p> <ul style="list-style-type: none"> <li>Sending emails</li> <li>Publishing content</li> <li>Contacting clients</li> <li>Making purchases</li> <li>Changing or deleting records</li> <li>Scheduling external meetings</li> <li>Submitting forms</li> <li>Claims about law, finance, health, safety, employment</li> </ul>	<p><b>REVIEW AFTER DRAFTING</b></p> <p><b>Human reviews output.</b></p> <ul style="list-style-type: none"> <li>Briefings</li> <li>Reports</li> <li>Content drafts</li> <li>Research summaries</li> <li>Meeting prep</li> <li>Internal notes</li> <li>Training materials</li> <li>Proposal drafts</li> </ul>	<p><b>REVIEW BY EXCEPTION</b></p> <p><b>Agent can run alone.</b></p> <ul style="list-style-type: none"> <li>Formatting</li> <li>Sorting</li> <li>Summarising</li> <li>Deduplicating</li> <li>Renaming</li> <li>Organising</li> <li>Extracting data</li> <li>Preparing first-pass options</li> </ul>
--	---	---

*“ The agent can prepare. The human approves impact.*

PROTECT THE USER.  
 PROTECT THE BUSINESS.  
 PROTECT THE BRAND.

04 THE IMPACT

# Agentic AI fails differently.

With prompting, the most common failure is a poor answer. With agents, it's poor execution — and that raises the stakes. CTRL reduces those failures by making the handoff explicit.

<h2>4<sup>x</sup></h2> <p>more impact per task when context, task, rules, and loop are all set before the agent starts.</p>	<h2>~30<sup>s</sup></h2> <p>average added prep time per handoff. Pays back in fewer revisions and less cleanup.</p>	<h2>0</h2> <p>agent autonomy without human review on impact actions. The agent prepares; the human approves.</p>
---	---	--

CONTRIBUTION TO AGENT EFFECTIVENESS

## How much each CTRL control matters.



Context carries the most weight because the agent acts inside whatever reality it's given. Task and Rules define the finish line and the boundaries. Loop controls execution quality once the work begins.

THE FOUR IN ONE BREATH

- C Context.** Set the world.
- T Task.** Define done.
- R Rules.** Keep control.
- L Loop.** Check, correct, stop.

WHAT CTRL IMPROVES

- Less drift.** The agent works inside reality.
- Less busywork.** The task defines completion.
- Less risk.** Rules prevent overreach.
- Less rework.** The loop self-corrects.

## 05 NEXT STEPS

# Where do you go from here?

Agentic AI isn't just a new tool category — it's a new work relationship. You're no longer only a question-asker. You're a manager of machine work. Three principles keep CTRL useful as the technology keeps moving.

01

## Handoff over prompt.

Don't start by asking "what should I type?" Start by asking "what work am I handing off?" The question reframes the whole interaction. You stop chatting and start delegating.

---

02

## Control over autonomy.

Don't give the agent freedom because the software can handle it. Give the agent freedom only where the rules, review points, and stopping conditions are clear. Autonomy without a leash is just risk wearing a productivity costume.

---

03

## Judgement over automation.

The agent can move the work forward. The human still owns the decision. CTRL keeps that relationship intact — and that's the relationship clients, teams, and regulators will all eventually ask you about.

*Set the context. Define the task. Write the rules. Run the loop.*

● END OF GUIDE

CLOSING

# AI mastery now requires better delegation.

Thank you for being part of this next stage of the work. PROMPT helped define how leaders speak to AI when they need better answers. CTRL extends that discipline into the agentic era — where AI doesn't merely respond but assists, executes, and moves through work.

CONTINUE THE WORK

For coaching, frameworks, and tooling that turn AI into a leadership instrument, visit [aibrain.coach](https://aibrain.coach) or reach the team at [hello@aibrain.coach](mailto:hello@aibrain.coach).

---

CONTRIBUTORS

CONCEPT & FRAMEWORK

**Ray Tan**

Built as a companion to The PROMPT Framework.

---

RESEARCH & STRATEGY

**Ray Tan**

---

CONTENT & WRITING

**Ray Tan**

---

CONTACT

STUDIO

AI Brain Coach

[aibrain.coach](https://aibrain.coach)

REACH OUT

[aibrain.coach](https://aibrain.coach)

[hello@aibrain.coach](mailto:hello@aibrain.coach)

[@aibraincoach](https://aibrain.coach)