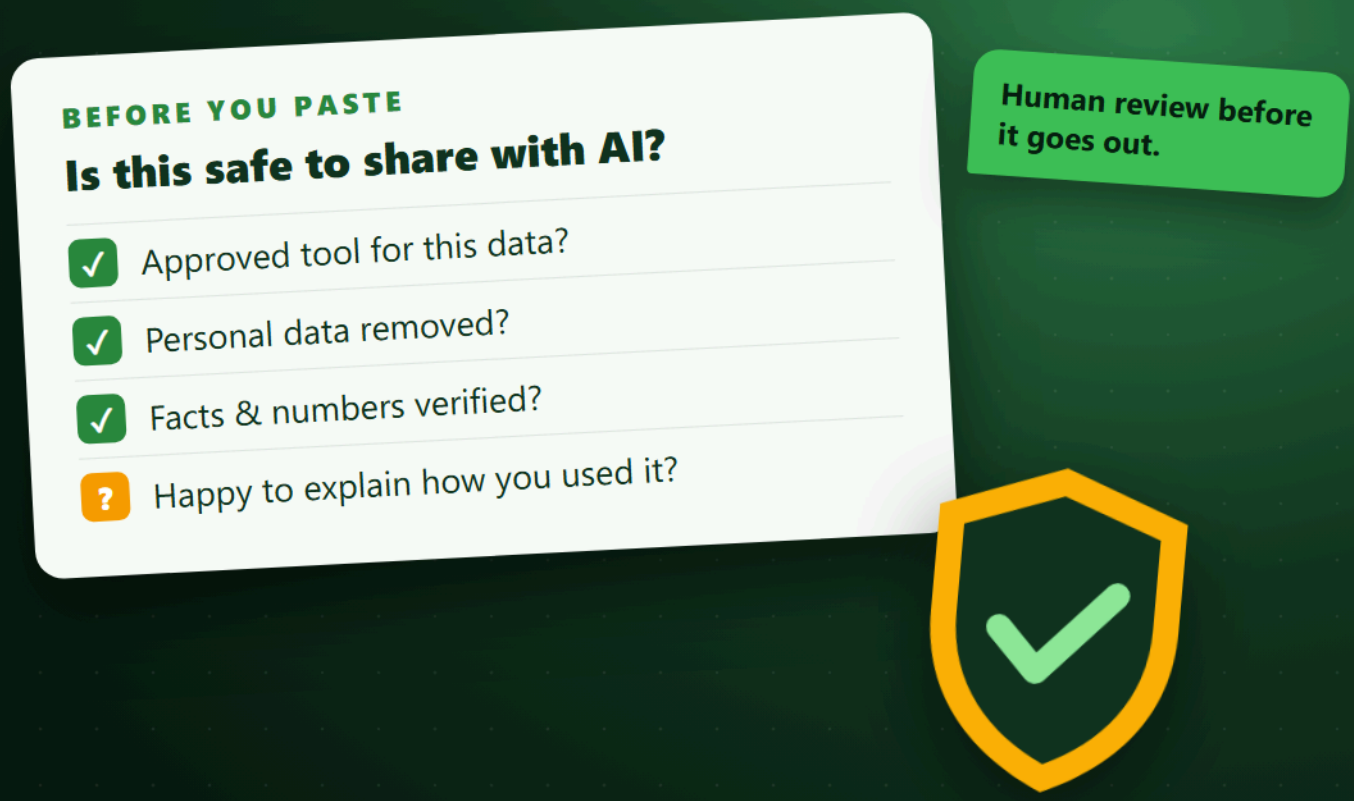


THE PRACTICAL GUIDE

Responsible AI for Everyday Work

Use AI Assistants Safely, Fairly, and Wisely on the Job

Privacy · Accuracy · Bias · Disclosure · Security
Checklists · Templates · Real workplace scenarios

PROTECT · VERIFY · DISCLOSE · INCLUDE · OWN IT

THENEXBOOKS

Responsible AI for Everyday Work

A Practical Guide to Using AI Assistants Safely, Fairly, and Wisely on the Job

First Edition

THENEXBOOKS

Responsible AI for Everyday Work

A Practical Guide to Using AI Assistants Safely, Fairly, and Wisely on the Job

First Edition, 2026

Copyright © 2026 thenexbooks (thenexbooks.gumroad.com). All rights reserved.

No part of this publication may be reproduced, distributed, or transmitted in any form or by any means without the prior written permission of the publisher, except for brief quotations in reviews and certain other non-commercial uses permitted by copyright law. The checklists, templates, and prompts in this book may be copied and adapted for use inside your own organization.

This book provides general information and practical guidance only. It is not legal, compliance, security, HR, medical, or financial advice. Laws, regulations, and AI products change often and differ by country and sector; always follow your organization's policies and consult qualified professionals for decisions in your situation. The information is provided "as is" without warranty of any kind, and the publisher assumes no responsibility for errors or omissions, or for damages resulting from its use.

Example AI responses are illustrative and were written for this book; they are not output captured from any specific product. Wrenfield and all people, companies, customers, and data associated with it are fictional. ChatGPT, Claude, Gemini, Microsoft Copilot, and other product names are trademarks of their respective owners and are used for identification only.

Preface

AI assistants arrived at work faster than almost any tool before them. One week they were a curiosity; the next, colleagues were using them to draft emails, summarize reports, answer customers, plan projects, and analyze spreadsheets. Many organizations are still catching up with policies and training. In the meantime, millions of people make small decisions every day: *Can I paste this? Can I trust this answer? Should I tell my client that AI helped? Is it fair to use this for hiring?*

Those small decisions are what responsible AI really is. It is not only a topic for engineers, lawyers, or executives. It is the habit of the person who pauses before pasting a customer list into a chatbot, who checks the number before it goes into the board report, who notices that an AI-written job ad quietly favors one kind of candidate, and who reports a mistake early so the whole team learns from it.

This book teaches those habits in plain English. There is no code and no jargon without an explanation. Instead you will find checklists, decision tables, "risky vs. responsible" comparisons, ready-to-use prompts, and realistic workplace stories. The goal is not to make you afraid of AI. It is to help you use it confidently — getting the speed and creativity it offers while protecting the people, information, and trust your work depends on.

Who This Book Is For

- **Everyday AI users** in any role — office, sales, marketing, customer service, finance, HR, operations — who want to use AI well and avoid costly mistakes.
- **Team leads and managers** who need to set expectations, answer questions, and model good practice for their teams.
- **HR, legal, compliance, and IT colleagues** who are writing AI policies and training and want practical material people will actually read.
- **Freelancers, small-business owners, public-sector and nonprofit staff** who have no AI policy to lean on and want sensible rules of their own.

What You Will Learn

- How AI assistants work and why they make confident mistakes — explained without technical jargon.
- The core principles of responsible AI and how to turn them into daily habits.
- How much to trust an answer, and how to scale your review to the stakes.

- What you can and cannot paste, how to protect confidential and personal data, and how to set up tools safely.
- How to spot AI-powered scams, deepfakes, and hidden instructions in documents and web pages.
- Verification habits for facts, numbers, sources, and summaries.
- When and how to disclose AI help, and how to respect copyright and other people's work.
- How to recognize bias, keep people decisions human, and communicate inclusively.
- Responsible practice for email, research, meetings, marketing, customer service, finance, legal, and safety work.
- How to handle incidents, what the main laws and standards mean for you, and how to build a responsible AI culture.

How This Book Is Organized

Part I — Foundations of Responsible AI at Work (Chapters 1–5) explains what responsible AI means, how assistants work and fail, the core principles, the human-in-the-loop mindset, and workplace AI policies. **Part II — Protecting People and Information** (Chapters 6–10) covers data privacy, confidential information, security threats, safe tool setup, and personal data. **Part III — Accuracy, Quality, and Honesty** (Chapters 11–15) covers hallucinations, verification, disclosure, intellectual property, and keeping your own judgment sharp. **Part IV — Fairness and People** (Chapters 16–20) covers bias, people decisions, inclusive communication, customer-facing AI, and teamwork. **Part V — Responsible AI in Everyday Tasks** (Chapters 21–25) applies everything to common kinds of work. **Part VI — Building a Responsible AI Culture** (Chapters 26–28) covers incidents, laws and standards, and leadership. The appendices provide a quick reference, checklists, templates, and a glossary.

You can read the book from start to finish or jump to the chapter you need today. Each chapter ends with a **Try It Yourself** section and a **Key Takeaways** box.

A Note on Tools and Laws

This book is tool-neutral. The habits it teaches apply to ChatGPT, Claude, Gemini, Microsoft Copilot, and the AI features built into email, office suites, meeting tools, and business software. Features, plans, and data settings differ between products and change often, so we describe them in general terms; always check your product's current documentation and your organization's settings. Where we mention laws and standards, we do so at a high level to help you ask better questions. We are not giving legal advice; your legal and compliance colleagues are the right people to interpret the rules for your situation.

The Running Example: Wrenfield

Throughout the book we follow **Wrenfield**, a fictional company of about 320 people that designs and sells outdoor furniture and garden products from its headquarters in Ashford Bay and its factory in Millbrook. Wrenfield has recently rolled out an approved company assistant, **Wrenfield Assist**, and a short AI use policy. You will meet Dana, the COO who sponsors the program; Omar from IT security; Grace from HR; Tomás and his customer service team; Aisha in marketing; Ben in finance; Priya, the in-house lawyer; Leo at the factory; and Nina, a new sales coordinator who loves AI and learns — sometimes the hard way — how to use it responsibly. Their successes and mistakes are the kind that happen in real organizations every day.

Conventions Used in This Book

> **TIP**
>
> A recommended habit or a technique that makes responsible use easier.

> **NOTE**
>
> Background information, a definition, or a useful aside.

> **WARNING**
>
> A common mistake that can harm people, leak information, or damage trust.

`Prompt`

> Prompts you can type into an AI assistant appear in boxes like this one.
> Words in [square brackets] are blanks for you to fill in.

`Example response (illustrative)`

> Example responses appear in dashed boxes. They were written for this book to show the typical shape of an answer; real assistants will word things differently.

Contents

Preface

PART I — FOUNDATIONS OF RESPONSIBLE AI AT WORK

1 Why Responsible AI Is Everyone's Job

2 How AI Assistants Work — and Why They Fail

3 The Core Principles in Plain English

4 The Human-in-the-Loop Mindset: How Much to Trust AI

5 Your Workplace AI Policy: Reading It, Following It, Filling the Gaps

PART II — PROTECTING PEOPLE AND INFORMATION

6 What You Can and Cannot Paste: Data Privacy Basics

7 Confidential Business Information and Client Data

8 Security Risks: Phishing, Deepfakes, Prompt Injection, and Shadow AI

9 Choosing and Setting Up AI Tools Safely

10 Personal Data About Customers and Colleagues

PART III — ACCURACY, QUALITY, AND HONESTY

11 Hallucinations and Everyday Verification Habits

12 Checking Sources, Numbers, and Summaries

13 Transparency and Disclosure: When to Say AI Helped

14 Copyright, Ownership, and Intellectual Property in Everyday Use

15 Keeping Your Judgment and Skills Sharp

PART IV — FAIRNESS AND PEOPLE

16 Understanding Bias in AI Outputs

17 AI in Hiring, Performance, and Other People Decisions

18 Inclusive and Accessible Communication with AI

19 Customer-Facing AI: Honesty, Empathy, and Escalation

20 AI and Your Colleagues: Trust, Workload, and Change

PART V — RESPONSIBLE AI IN EVERYDAY TASKS

21 Writing, Email, and Everyday Communication

22 Research, Analysis, and Reports

23 Meetings, Notes, and Recording Tools

24 Marketing, Sales, and AI-Generated Images

25 High-Stakes Work: Finance, Legal, Health and Safety

PART VI — BUILDING A RESPONSIBLE AI CULTURE

26 When AI Goes Wrong: Mistakes, Incidents, and Reporting

27 Laws, Regulations, and Standards in Plain English

28 Leading by Example: Champions, Team Practices, and Your Personal Plan

APPENDICES

A Responsible AI Quick Reference

B Checklists

C Templates

D Glossary

PART I

Foundations of Responsible AI at Work

What responsible AI means for someone who is not an engineer, how AI assistants really work and fail, the principles behind good use, how much to trust an answer, and how to work with your organization's AI policy.

1 · Why Responsible AI Is Everyone's Job
2 · How AI Assistants Work — and Why They Fail
3 · The Core Principles in Plain English
4 · The Human-in-the-Loop Mindset
5 · Your Workplace AI Policy

CHAPTER 1 · PART I — FOUNDATIONS OF RESPONSIBLE AI AT WORK

Why Responsible AI Is Everyone's Job

AI assistants have moved from novelty to normal work tool, and the people who use them every day now make most of the decisions that matter. This chapter explains what "responsible AI" means for someone who is not a data scientist, introduces the company we will follow through the book, and gives you five simple questions to ask before you rely on any AI tool.

A New Job, a New Tool

On her first Monday at Wrenfield, Nina Park arrived early, found her desk in the sales area of the Ashford Bay headquarters, and opened her new laptop to a welcome message. Underneath the usual links — the staff handbook, the holiday calendar, the expense system — was a new icon labeled "Wrenfield Assist". A short note explained that this was the company's approved AI assistant and that everyone could use it for everyday work, as long as they followed the Wrenfield AI Use Policy.

Nina smiled. She had used AI chatbots in her previous job and at home for everything from planning trips to rewriting awkward emails. She thought of herself as a confident user. If anything, she expected to be the person teaching her new colleagues a few tricks.

By the end of the week she had learned something more useful than a trick. She had learned that being good at *getting answers* from an AI assistant is not the same as being good at *using* one at work. The second skill involves knowing what information you are allowed to share, how much to trust what comes back, who might be affected by the result, and who is responsible when something goes wrong. That second skill is what this book is about.

You may recognize yourself in Nina. Perhaps you use an AI assistant daily, or your organization has just switched on AI features in your email and meeting software, or you run a small business and must set your own rules. Wherever you are starting from, the habits in this book will help you get real value from AI while protecting the people, information, and trust that your work depends on.

What "Responsible AI" Means for a Normal Worker

For technology companies, governments, and researchers, "responsible AI" covers big questions: how models are trained, tested, and regulated. Those questions matter, but other people decide them. You do not choose how a large AI model is built.

What you *do* choose, every working day, is how you use it: what you type into the box, whether to check the answer, whether to tell a client that AI helped, whether an AI tool should play any part in a decision about an applicant, a complaint, or a safety procedure. In most organizations, those small, frequent choices affect real people more than any policy document.

So, in this book, we use a simple working definition:

> **NOTE**
>
> **Responsible AI use at work** means using AI tools in ways that are useful, safe, fair, honest, and accountable — protecting the people and information involved, checking the output in proportion to what is at stake, and keeping a human responsible for the result.

Notice that this definition does not say "avoid AI", and it does not require you to understand the mathematics inside a model. It asks you to bring the same good judgment you already bring to other parts of your job — handling money, speaking to customers, keeping personal files private — to a new kind of tool.

Think of email: nobody needs to understand networking to send one, but everyone is expected not to forward confidential files to the wrong person. Those habits are professional, not technical. The same is true of AI.

Why It Is Everyone's Job, Not Just IT's

It is tempting to think that responsible AI belongs to the IT department, the legal team, or a special committee. Those groups play an important role. They choose which tools are approved, set up the data settings, write the policy, and handle serious incidents. But they cannot see what you paste into a chat window at four o'clock on a busy Thursday. They cannot check every email you draft or every summary you send to your manager.

There are three reasons why responsibility has spread to everyone.

1. AI is now inside ordinary tools

AI features now sit inside tools people already use: a "summarize this thread" button in email, a "draft with AI" option in a word processor, automatic meeting notes, suggested replies in a customer-service system. Many people use AI without thinking of it as "using AI" at all, which makes individual awareness more important, not less.

2. The user controls the inputs and the outputs

If you paste in a customer list, that information has left your hands. If you copy the answer into a proposal without reading it, any error is now your error. The person at the keyboard stands at both doors, and no technical control fully replaces their judgment.

3. The consequences land on real people

When AI output is wrong or unfair, the people affected are usually not the ones who typed the prompt. They are customers given a wrong answer, applicants screened unfairly, or colleagues whose private comments end up in meeting notes. Responsible use keeps those people in mind, even when they are not in the room.

> **TIP**
>
> A useful habit is to picture the person at the other end of your AI-assisted work — the customer reading the reply, the colleague reading the summary, the applicant whose CV is being assessed. Ask yourself whether they would be comfortable with how the work was done.

The Real Benefits — and the Real Risks

A book about responsible AI can easily sound like a list of dangers. That would be misleading. AI assistants are genuinely useful, and people who avoid approved tools out of fear often end up quietly using unapproved ones with no guidance at all. Here is a realistic view of both sides.

Where AI often helps	Where things often go wrong
Drafting emails, letters, and first versions of documents	Confidential or personal data shared with the wrong tool
Summarizing long material you can check	Confident answers that are simply wrong
Rewording for tone, clarity, or plain language	Summaries that drop or distort key points
Brainstorming ideas, headlines, and options	Biased or stereotyped wording and images
Explaining unfamiliar terms and concepts	Numbers and calculations that do not add up
Turning notes into structured lists or tables	People trusting output because it sounds fluent
Preparing questions before a meeting	Decisions about people handed to a machine
Practicing a difficult conversation	Scams using fake voices, images, and messages

Most of the risks in the right-hand column are not about technology "going rogue". They are ordinary mistakes — sharing the wrong thing, trusting too quickly, letting a tool make a call a person should make — and ordinary habits prevent them.

It is also worth being honest about pressure. People rarely skip a check out of carelessness; they skip it because twelve things are due by five o'clock and the answer looks fine. So this book favors quick checks and clear lines that are easy to remember when you are busy.

Meet Wrenfield

To make the ideas in this book concrete, we will follow one organization from chapter to chapter. **Wrenfield** is a fictional mid-size company that designs and sells outdoor furniture and garden products: planters, benches, patio sets, and garden tools. It has about 320 employees. Its headquarters and design studio are in the fictional city of Ashford Bay, and it runs one factory in Millbrook. Wrenfield sells through its own online store, through garden centers, and to trade customers such as hotels, landscapers, and local councils. Everyone and everything about Wrenfield — the people, the customers, the data — is invented for this book.

Wrenfield is ordinary in the best sense: HR, finance, marketing, customer service, sales, legal, IT, and a factory floor. Whatever your job, someone there does something similar.

The people you will meet

- **Dana Whitfield**, Chief Operating Officer, sponsors Wrenfield's responsible-AI program and chairs its AI Working Group.
- **Omar Haddad**, IT security lead, owns the list of approved tools and worries, professionally, about where data goes.
- **Grace Liu**, HR business partner, handles recruiting, performance reviews, and employee questions.
- **Tomás Rivera** leads the six-person customer service team that handles the website chat and the email queue.
- **Aisha Bello**, marketing manager, is responsible for product copy, social media, and catalog images.
- **Ben Kowalski**, finance analyst, prepares monthly reports, checks supplier invoices, and builds forecasts.
- **Priya Shah**, in-house legal counsel, deals with contracts, privacy, and compliance.
- **Nina Park**, new sales coordinator, joined in 2026 and is an enthusiastic AI user who learns responsible habits as the book goes on.
- **Leo Martins**, factory operations supervisor in Millbrook, writes safety reports and manages shift schedules.

Wrenfield's **AI Working Group** is chaired by Dana and includes Omar, Priya, Grace, and Tomás, plus rotating "AI champions" from each team. The group does not try to control every use of AI. Its job is to choose the tools, set a few clear rules, answer questions, and learn from mistakes.

Wrenfield Assist and the policy

In early 2026, Wrenfield rolled out **Wrenfield Assist**, an approved company AI assistant. It is an enterprise plan of a mainstream assistant, configured so that company data is not used to train the vendor's models. We will not name the vendor, because the lessons in this book apply to any mainstream assistant — and because the details of plans and settings vary between products and change over time.

Alongside the tool, the Working Group published the **Wrenfield AI Use Policy**. It is only two pages long, deliberately. Dana's view was simple: a policy nobody reads protects nobody. We will walk through it in detail in Chapter 5. For now, the key points are that staff should use approved tools for work, follow the data classification rules below, check AI output before relying on it, and report problems to `ai-help@wrenfield.example` without fear of blame.

The data classification

Like many organizations, Wrenfield sorts its information into levels. The levels existed long before AI arrived — they already guided what could be emailed externally or stored on a shared drive. The AI Use Policy simply connects each level to a clear rule about AI tools.

Level	Examples at Wrenfield	AI rule of thumb
Public	Published catalog, website text, press releases	Any tool, including personal accounts
Internal	Meeting agendas, internal guides, org charts	Wrenfield Assist
Confidential	Pricing plans, supplier terms, unreleased designs	Wrenfield Assist only, and only when needed
Restricted	Customer and employee personal data, health info, payroll, passwords and keys, unreleased financials, legal matters	Never in any AI tool, unless a specific approved process says so

One more line in the policy matters a great deal: **consumer or personal AI accounts may be used for Public data only.** A personal chatbot account on your phone is not the same as Wrenfield Assist, even if the two look similar on screen. The company has no agreement with the provider of your personal account, cannot see its settings, and cannot help if something goes wrong.

Clear levels turn most decisions into quick ones, instead of a fresh judgment call under pressure every time. We will return to classification in depth in Chapter 6, including how to remove personal details from material so that it can be used safely.

> **WARNING**
>
> Classification labels are only as good as the person applying them. A spreadsheet is not "Internal" just because it lives on an internal drive. If it contains customer names, email addresses, or order histories, it contains personal data — and that makes it Restricted, whatever the folder is called.

The Five Questions

No policy can anticipate every situation. What people need most is a short, memorable way to pause before relying on AI. Wrenfield's Working Group settled on five questions, which we use throughout the book. Once habitual, they take about ten seconds.

1. **Is this the right tool for this job?** Is the tool approved for this kind of work? Is AI a sensible choice at all, or is this something better done yourself, or by a specialist?
2. **What am I sharing?** What information will go into the tool — text, files, screenshots, recordings? What level is it? Does it include anyone's personal data?
3. **What happens if it is wrong?** If the output contains an error, who would notice, and what would the damage be? A clumsy sentence in an internal note is very different from a wrong figure in a customer quote.
4. **How will I check it?** What will you compare the output against — the original document, a trusted source, a colleague's review, your own knowledge? Is the checking proportional to the stakes?
5. **Who is affected, and could I explain how I used AI?** Would you be comfortable telling your manager, the customer, or the person concerned exactly how AI was involved? If not, why not?

The first two questions protect *information*. The third and fourth protect *quality*. The fifth protects *people and trust*. You will find that almost every chapter in this book is really an expansion of one of these questions.

Here is a quick way to keep them in view:

Question	It protects	Quick version
1. Right tool?	Data and quality	"Approved and suitable?"
2. What am I sharing?	Privacy and confidentiality	"What level is this?"
3. What if it is wrong?	Quality and safety	"How bad is a mistake?"
4. How will I check it?	Accuracy	"Against what?"
5. Who is affected?	Fairness and trust	"Could I explain this?"

> **TIP**
>
> You do not need to answer the five questions in writing. For everyday tasks, they are a mental checklist. For higher-stakes work — anything involving money, people decisions, legal matters, health, or safety — it is worth jotting down your answers, even in one line each.

Nina's First Week

Let us return to Nina and watch how the five questions played out in her first week. Nothing dramatic happened. That is exactly the point: most responsible AI use is made of small, sensible decisions.

Monday: onboarding with Omar

On Monday afternoon, Omar Haddad from IT security ran the AI section of the new-starter onboarding. He opened not with warnings but by asking what everyone had used AI for recently — holiday plans, homework help, a wedding speech.

"Good," Omar said. "Keep using it. What I want you to take away today is one sentence." He put the classification table on screen and pointed to the bottom row. "Restricted data never goes into any AI tool unless there is a specific approved process. Everything else, we can talk about."

Nina asked the question half the room was thinking: "What about the chatbot on my phone? It is the same kind of thing." Omar nodded. "It might be the same technology. But the agreement and settings are different, and we cannot control either. Personal accounts are for Public data only."

Nina wrote that down. It seemed obvious in the room; it is less obvious at a busy desk.

Tuesday: a follow-up email

On Tuesday, Nina's manager asked her to follow up with a garden center that had asked about stocking a new range of planters. Nina had the notes from the phone call and the public product page. She opened Wrenfield Assist and started to type.

Her first instinct was to paste everything: call notes, the garden center's order history, and the internal price list. Then she ran through the questions. The product page was Public; her call notes were Internal. The order history included the store manager's name, direct phone number, and past spending — data she did not need for a follow-up email. The price list was Confidential, and pricing was to be sent separately anyway.

So she shared only what the task required.

Prompt

```
You are helping a sales coordinator at a garden products company.
Draft a short, friendly follow-up email to a garden center manager
who asked about stocking our new planter range.

Key points from the call:
- They liked the large terracotta-style planters
- They asked about delivery lead times for spring
- They want a sample before ordering

Do not include prices. Say pricing will follow separately from our sales team.
Keep it under 150 words. Use [manager name] and [store name] as placeholders.
```

Example response (illustrative)

```
Subject: Following up on our planter range

Dear [manager name],

Thank you for taking the time to speak with me about stocking our new
planter range at [store name]. I'm glad the terracotta-style planters
caught your eye.

I'll confirm spring delivery lead times with our logistics team and
send them to you this week. We'd be happy to arrange a sample so you
can see the finish and quality for yourself...
```

The draft was good. Nina changed two phrases to sound more like herself, filled in the real names in her email program — not in the AI tool — and checked one thing: had she actually promised a sample? Her notes said the manager *wanted* one; she had not yet confirmed Wrenfield could provide one. She softened the line to "I'll check how quickly we can arrange a sample" and sent it.

The tool saved Nina time, but the two decisions that mattered — what to share and what to promise — were hers.

Wednesday: the contract question

On Wednesday, a hotel group sent over its standard supplier agreement ahead of a possible order for patio sets. Nina's manager asked her to "get a sense of anything unusual" before passing it to legal. The document was fifteen pages long.

Nina's first thought was to upload it to Wrenfield Assist and ask for a summary of risky clauses. Then she hit the third question: what happens if it is wrong? If the summary missed a clause about penalties or liability, Wrenfield could sign up to something costly. And the agreement itself — a customer's contract terms — was at least Confidential. She was not sure it was not a legal matter, which would make it Restricted.

She sent a short message to Priya Shah in legal.

Prompt

> Hi Priya, a hotel group has sent their supplier agreement before a possible patio-set order. My manager asked me to flag anything unusual before it goes to you. Is it OK to use Wrenfield Assist to summarize it first, or should I just send it straight over? Happy to do whichever is more useful.

Priya's reply was quick and friendly. Contracts under negotiation were treated as legal matters, so they were Restricted under the policy. Legal had its own approved process for AI-assisted contract review, which Priya ran herself. "Send it straight to me," she wrote. "But thank you for asking — that is exactly the right instinct." (We will see how Priya uses AI carefully in her own contract work in Chapter 25.)

Nina felt slightly embarrassed, but she need not have. Most workplace AI mistakes come from people who assume, not from people who ask.

Thursday: a confident wrong answer

On Thursday, a landscaping firm emailed asking how heavy Wrenfield's largest bench was, because they needed to plan how many people would lift it into a rooftop garden. Nina was in a hurry and typed the question straight into Wrenfield Assist.

Prompt

> How much does the Wrenfield large hardwood garden bench weigh?

Example response (illustrative)

> The Wrenfield large hardwood garden bench typically weighs around 28 kg (about 62 lb), making it manageable for two people to lift...

The answer sounded precise and helpful. Something about it made Nina hesitate: Wrenfield Assist had not been connected to any product database, and she had not given it any product information. Where had the number come from? She opened the official product specification sheet. The bench weighed considerably more than the answer suggested — enough that the difference would matter for a team carrying it up stairs.

The assistant had not "lied". As Chapter 2 explains, it produced a plausible answer based on what garden benches in general tend to weigh, rather than looking up a fact it did not have. For a customer planning a rooftop lift, that could have led to an injury.

Nina answered from the specification sheet, then tried a better approach for future questions.

Prompt

> Below is the official product specification for one of our benches.
> Using ONLY this text, write a short reply to a customer who asked
> how heavy it is and how many people should lift it.
> If the specification does not say something, say it is not stated.
>
> [paste the public specification text here]

Giving the assistant the source material, and telling it to stick to that material, reduces the risk of invented details. It does not eliminate it — she still compared the reply against the sheet before sending — but it moves the assistant from guessing to summarizing, which is much safer.

> **WARNING**
>
> A fluent, specific answer is not evidence of accuracy. AI assistants can state invented figures, names, and facts in exactly the same confident tone as correct ones. Whenever an answer includes a number, a date, a name, or a safety-relevant detail, check it against a source you trust.

Friday: looking back

On Friday afternoon, Nina had a short check-in with Tomás Rivera, who sits on the AI Working Group and had offered to mentor new starters in the customer-facing teams. He asked how her first week with Wrenfield Assist had gone.

"Honestly? I thought I knew all this already," Nina said. "I know how to get good answers out of it. What I did not really think about was everything around the answer — what I put in, what I promise, who checks it."

Tomás laughed. "That is most of the job. My team drafts chat replies with Assist all day. The skill is knowing which replies need rewriting, which need a manager, and which should never be automated."

Before logging off, Nina turned her scribbled notes from the week — her own Internal notes, nothing about customers — into a personal checklist.

Prompt

> Turn these rough notes into a short personal checklist I can keep
> next to my screen, with no more than 7 items. Plain language.
>
> Notes: only approved tools for work stuff / personal chatbot = public only /
> don't paste customer details, use placeholders / contracts go to Priya /
> check any numbers against the spec sheet / don't promise things AI suggests /
> ask if unsure - ai-help mailbox

Example response (illustrative)

My AI checklist
1. Use Wrenfield Assist for work; personal chatbots for public info only.
2. Leave out customer names and details - use placeholders instead.
3. Send contracts and legal documents to Priya, not to AI.
4. Check every number, date, and safety detail against the source.
5. Don't promise anything just because the AI draft says it.
6. Read every draft before sending - it goes out in my name.
7. Unsure? Ask, or email ai-help@wrenfield.example.

It was not a perfect list — as we will see in later chapters, Nina still had lessons ahead, some of them uncomfortable. But it was a good start, and it was hers.

Beyond Wrenfield: The Same Questions Everywhere

You do not need a company-approved assistant, a Working Group, or a classification table to use AI responsibly. The five questions work anywhere. Consider two very different people.

A freelance bookkeeper

Marisol runs a one-person bookkeeping business for about a dozen small clients: a café, a plumber, two hair salons, and a handful of sole traders. She uses a general AI assistant on a paid personal plan to draft client newsletters and explain tax deadlines in plain language.

One evening she was tempted to upload a client's full bank export and ask the assistant to categorize the transactions. She ran through the questions. Was this the right tool? She had never checked the plan's data settings, and her client agreements promised confidentiality. What was she sharing? Names of her client's customers, amounts, and account details — personal and financial data belonging to people who had not agreed to it. What if it was wrong? A miscategorized expense could flow into a tax return with her name on it.

Instead, she used the assistant to write clear category rules from made-up example transactions, then applied them herself in her accounting software. She got much of the time saving, none of the data risk, and kept full control of the numbers.

Prompt

> I'm a bookkeeper. Here are 10 made-up example transactions for a small café.
> Suggest clear rules for sorting them into categories such as ingredients,
> equipment, utilities, and wages. Explain any that are ambiguous and what
> question I should ask the client to resolve them.
>
> 1. "Fresh Dairy Co - 184.20"
> 2. "City Power - 96.50"
> ...