RISE INSIGHTS / WORK & PRODUCTIVITY

AI that actually SAVES TIME

A practical guide to better work, clearer workflows and measurable value



Dr. Tony Bader

MAKE USEFUL AI A TEAM HABIT.

RISE ACADEMY

A NOTE TO THE READER

01

The real measure is finished work

A draft that appears in seconds can still take an hour to fix. A summary can look complete while leaving out the only decision that mattered. The useful question is how much better the whole task becomes, from the first input to an approved result.

This book is for professionals, managers and small teams who want a sensible starting point. It follows one idea: use AI where the purpose is clear, the source is available and a person can judge the result. Start with a repeatable task, then build the habits around it.

You will see how to select a first use, brief the tool, organise review and measure what changes. The methods are original practical recommendations. Worked examples and figures are illustrative; the research notes distinguish evidence from advice.

CONTENTS

03 Find the friction

04 Choose a first use

05 Give the task a usable brief

06 Build a source pack

07 Make review specific

08 Count the whole task

09 Turn a good result into a routine

10 Build team capability

11 Keep automation within its limits

12 Your first 30 days

13 Evidence & further reading

14 Work with Tony Bader

Make useful AI a team habit.

RISE ACADEMY / Training and research, Innovative Solutions
First edition, November 2026. Practical guidance by Dr. Tony Bader.

01 / FIND THE FRICTION

Start with the work that slows you down

The best first use is usually hiding inside an ordinary task.

Look at the handoffs

Follow a recurring task from request to completion. A weekly update might involve finding notes, checking figures, drafting, translating and waiting for approval. These steps require different kinds of effort. Treating them as one writing problem makes it difficult to see where AI might help.

Look for moments when people repeatedly reorganise information, reformat approved content or begin from a blank page. Also notice delays caused by missing facts or unclear ownership. A faster draft will have limited value when the person who must confirm the figures is unavailable.

Choose a small, visible improvement

A useful target is specific: turn approved project notes into the first draft of a weekly update. The source is known, the output is recognisable and a reviewer can compare the two. "Improve our communications" is too broad to test.

Set aside tasks whose quality you cannot judge. If nobody on the team can verify the answer, adding a confident draft creates another dependency. First strengthen access to expertise, reliable sources or a clearer process.

A TASK WORTH TESTING

The same update is written every Friday. Staff have the approved facts, but spend time grouping them into progress, decisions and next steps. AI could propose that structure; the project owner still confirms what happened.

Find one repeated transformation with a clear source and a clear finish.

02 / CHOOSE A FIRST USE

A useful first task has four qualities

Frequency matters. So do consequences, inputs and the ability to review.

Use four selection questions

Does the task happen often enough to learn from? Can you use approved inputs? Can a person check the result against something dependable? Would an error be caught before it affects another person? A good first use has credible answers to all four.

For a small team, drafting from a published service guide may be a reasonable trial. Deciding who qualifies for that service brings different consequences and requires a separate assessment. The same tool can sit inside very different levels of responsibility.

Define success before trying it

Write one sentence that names the output and the quality threshold. For example: "Prepare a weekly update that contains every approved decision, adds no unsupported claims and takes less total staff time to complete."

Choose a stop condition as well. Repeated invented facts, sensitive information appearing in the process or a reviewer who cannot keep up should trigger a pause. An attractive result on one occasion is a reason to investigate, rather than a basis for rolling it out.

ILLUSTRATIVE FIRST-USE COMPARISON

Public event notice: approved facts and easy checking. Internal budget summary: check data permissions and numerical accuracy first. Service eligibility decision: keep within the authorised decision process and assess separately.

Find your best first use

Ask about a facilitated AI workflow session. Bring one recurring task and leave the discussion with a clearer trial scope.

DISCUSS A WORKFLOW SESSION →

03 / GIVE THE TASK A BRIEF

Tell AI what a good result must do

A useful prompt carries the information a competent colleague would need.

State the job and its boundaries

Name the reader, the purpose, the permitted source and the format. Add the constraints that matter to your organisation: length, language, required sections and claims that need confirmation. Leave out background that has no bearing on the task.

Explain how gaps should appear. Asking the tool to list missing details separately makes uncertainty easier to spot. This instruction can still fail, so the reviewer must compare the output with the source. A prompt sets expectations; it cannot certify compliance.

Refine one dimension at a time

If the first result is too long, change the length. If the facts are wrong, investigate the input and source handling before adjusting the tone. Repeatedly asking for a more polished answer may disguise the original problem.

Keep the approved source and acceptance criteria stable while testing revisions. Otherwise, you will struggle to tell whether an improvement came from a better instruction, easier material or a different judgement about quality.

A WORKED BRIEF

Using only the approved notes below, draft a 180-word update for project partners. Use three headings: Progress, Decisions, Next steps. Keep dates and figures unchanged. Put missing owners or deadlines in a separate Questions list. Do not invent commitments.



RISE / WORK

The prompt should make the reviewer's job easier as well as the drafter's.

04 / BUILD A SOURCE PACK

Better inputs reduce avoidable repair

Decide what the tool may rely on before asking it to write.

Give the task a single reference point

A source pack is a small set of approved materials chosen for one job. For a project update, it might contain confirmed milestones, the current schedule and the decisions the project owner has signed off. Label the date and owner of each item.

Separate confirmed facts from proposals. A note saying "we could extend registration" should not become "registration has been extended". Put unresolved items under an explicit heading so both the tool and the reviewer can see the distinction.

Use only what the task needs

If a public update needs the number of participants, an aggregate count may be enough. The full attendance register adds information without improving that output. Ask the data owner which material can enter the approved tool and account.

Removing a name does not automatically make a record safe to use. A location, date and unusual event can still point to a person. Where sensitive records are involved, use the organisation's designated process and qualified support.

Keep the source available during review

Store the final approved reference in the normal workspace, with appropriate access. A reviewer should be able to find it without searching an entire chat history. If the source changes, record which version supported the published result.

KEEP THE PARTS VISIBLE

01	02	03
Approved facts	Open questions	Named owner

A useful source pack is current, permitted and small enough to check.

05 / MAKE REVIEW SPECIFIC

Assign a check, not just a reviewer

"Please review" leaves too much of the job undefined.

Break quality into observable checks

For a weekly update, ask the reviewer to check the dates and figures, confirm that every decision is represented and identify any claim without a source. Then ask whether the language suits the reader. These are distinct checks and may need different people.

Read the source independently before comparing it with the draft. This helps the reviewer form an expectation of what should be there. Checking only what the AI wrote can leave omissions invisible, especially when the draft sounds complete.

Give review enough authority

The reviewer needs time, access to evidence and permission to reject the output. If publication is automatic or the deadline makes changes impossible, the review step offers little practical control.

Decide what requires escalation. A wording change may sit with the editor. An eligibility statement may require the programme owner. A statement about an individual may require a different authorised route. Record who can make the final decision.

Save examples of failure

Keep a short, appropriately protected record of recurring mistakes. Use it to revise the brief and train colleagues. Remove sensitive details from learning examples and avoid copying restricted records into a general prompt library.

> THREE OUTCOMES
>
> Approve when the result meets the agreed checks. Revise when a correctable issue remains. Hold when the evidence, permission or decision owner is missing. Each outcome should be available in the real workflow.

Review works when someone can detect a problem and change what happens next.

06 / COUNT THE WHOLE TASK

Measure the whole task

A fair comparison includes preparation, correction and approval.

Compare similar work

Record the usual time for a few representative examples before the trial. Include source preparation, drafting, review and corrections. During the trial, track the same steps and note unusual cases. A short, easy assignment is a poor comparison with last month's hardest report.

Track quality alongside time: unsupported claims, missing decisions, corrections after release and the reviewer's judgement. Faster completion is useful only when the result still meets the agreed standard.

Distinguish capacity from cash

If a team saves staff time, decide how that time will be used. It might create room for follow-up calls or better checking. It becomes a cash saving only when spending actually changes. Keep software, training and setup costs visible in either case.

ILLUSTRATIVE TIME COMPARISON / MINUTES PER UPDATE

Stage	Usual process	AI trial
Prepare inputs	10	15
Draft	35	10
Review & corrections	15	20
Total	60	45

Illustration only: 15 minutes saved per update. At 20 comparable updates, that is 5 staff hours before setup, training and software costs.

Prove the value in your setting

Ask about designing a small AI pilot with meaningful measures for time, quality and reviewer effort.

DISCUSS AN AI PILOT →

07 / MAKE IT REPEATABLE

Turn a good result into a team routine

A successful chat becomes useful when another person can reproduce the process.

Save the method with the prompt

A prompt on its own does not explain which inputs are approved, who checks the answer or where the final version goes. Keep a short process note beside it: purpose, source requirements, expected output, review checks and escalation route.

Add an example of an acceptable result and an example of a failure. A colleague can learn more from seeing a missing decision restored than from a vague instruction to be accurate. Use fictional or appropriately cleared material.

Name the owner of the routine

Someone should update the source pack, review feedback and retire obsolete versions. Ownership can sit with the person who already manages the task. Without it, staff may continue using an old policy, date or service description because it is embedded in a saved example.

Use the same approved account and workflow conditions when comparing results. Record important changes to the tool, settings or process. A routine that worked before a change deserves another representative check.

Allow an ordinary fallback

Keep the normal template and source material accessible. Staff should be able to complete the task when the AI service is unavailable, produces an unusable result or cannot handle the approved inputs.

THE ROUTINE IN ONE PARAGRAPH

Each Friday, the project officer uses the approved milestone note to draft the update in the authorised tool. The project owner checks decisions, dates and commitments. The officer saves the approved copy in the team drive and records any recurring error.



RISE / WORK

Standardise the conditions that make the result dependable.

08 / BUILD TEAM CAPABILITY

Teach people to judge the result

Confidence with a tool and competence in a task are different learning goals.

Use work people recognise

Training becomes easier to apply when participants can see their own job in the example. Use an approved notice, a fictional meeting summary or a familiar reporting format. Ask people to explain what makes the output acceptable before they start prompting.

Let participants compare a fluent but flawed answer with its source. Include an omitted condition, an invented date and a statement that changes a proposal into a promise. Discussion should identify the consequence of each change and the person who can correct it.

Make uncertainty safe to report

A colleague who says "I cannot verify this" provides useful information. Give staff a place to ask questions and a clear route for difficult cases. If success is defined only as producing more output, people may hide the checking time or overlook errors.

Protect learning time for experienced staff as well as new users. Familiarity with the work can help detect subtle mistakes, while familiarity with the tool can help structure the task. Pairing these strengths can improve the discussion.

Assess the behaviour you need

A short scenario can test whether someone chooses permitted inputs, preserves a source fact and escalates an unsupported claim. These behaviours tell you more about readiness than asking whether participants enjoyed the demonstration.

A USEFUL TEAM DISCUSSION

"The draft reads well, but the source says the date is provisional. Who confirms it, and what should we publish until then?" The value lies in recognising the missing decision and handling it correctly.

Build a shared standard for acceptable work, then practise applying it.

09 / KEEP AUTOMATION BOUNDED

Drafting permission is not sending permission

Connecting AI to other systems changes what an error can do.

Separate suggestions from actions

A tool that drafts an email gives a person a chance to inspect the text. A connected workflow that sends the email, changes a record or approves a payment has additional authority. Assess each permission separately before enabling it.

Begin with the narrowest access the task requires. A system preparing a public update may only need approved material in a selected folder. Access to the whole shared drive adds exposure and makes it harder to know which source supported the result.

Treat incoming material as material

Documents and web pages can contain instructions intended to redirect an AI system. Your team should treat external content as information to assess. Where tools can take actions, involve the person responsible for security and integration design before relying on a prompt to control access.

Use explicit approval for consequential actions, practical activity records and a way to stop the connection. Test the fallback. A button labelled "human approval" is useful only when the person can understand the proposed action and prevent it.

Expand after evidence

A reliable drafting trial is one piece of evidence. Before adding a live connection, test the new permissions and failure cases. The workload may look similar to staff while the consequences of an error have changed substantially.

> A BOUNDARY TO MAKE EXPLICIT
>
> "The system may prepare a reply from the approved FAQ. Only an authorised staff member may send it. Requests outside the FAQ go to the service owner." This boundary must be supported by the actual access settings.

Every new permission deserves its own reason, owner and test.

10 / YOUR FIRST 30 DAYS

Build evidence before expanding

A modest trial can produce a useful decision in one month.

Week 1 / Understand the task

Choose a recurring use, map the existing steps and record a baseline. Confirm the permitted inputs, review owner and normal quality standard. State what would make you pause the trial.

Week 2 / Test away from live release

Use approved or fictional material to try the brief on representative cases. Include incomplete notes and an example where the correct answer is that information is missing. Compare every result with its source.

Week 3 / Run a controlled trial

Use the method on a small amount of suitable work, with review before release. Record total time, corrections and exceptions. Invite staff to describe where the process adds effort as well as where it helps.

Week 4 / Decide and document

Compare the evidence with the original goal. Continue when the benefit and controls hold. Revise when a specific problem can be addressed. Stop when the task is a poor fit or the team cannot support reliable checking.

If you expand, change one major condition at a time. A new audience, source type or system permission creates a fresh question to test. Keep the owner, review process and fallback visible as the routine develops.

A good first month ends with a justified decision and a named owner.

Build the next step around your work

Work with Tony Bader on a tailored session for your team: select the use, practise the review and define a manageable pilot.

PLAN A TEAM SESSION →

SOURCES & READING NOTES

Evidence & foundations



Research can help frame a trial. Results from a particular task or workplace do not establish what your team will achieve.

01

BRYNJOLFSSON, LI & RAYMOND / 2025

Generative AI at Work

The Quarterly Journal of Economics, 140(2), 889-942. A customer-support study found productivity benefits that varied across workers. It supports testing the actual task and user group.

READ THE ORIGINAL SOURCE →

02

DELL'ACQUA AND COLLEAGUES / WORKING PAPER, 2023

Navigating the Jagged Technological Frontier

The study examines uneven AI performance across knowledge-work tasks. Its practical implication here is to test specific uses rather than assume one success transfers to another.

READ THE ORIGINAL SOURCE →

03

NIST / 2024

Generative Artificial Intelligence Profile

NIST AI 600-1. Further reading on fabricated outputs, information security and human interaction with generative AI.

READ THE ORIGINAL SOURCE →

Published by Rise Academy, the training and research division of Innovative Solutions. This book offers practical learning and discussion. Adapt the guidance to your organisation and operating context.



TAKE THE NEXT STEP

Make the next task a better process.

Bring one recurring task, its usual output and the friction your team experiences. Together, we can explore where AI could help, what people should keep control of and how to tell whether the change works.

Team training

Build practical skills around the tasks your colleagues actually perform.

Workflow design

Clarify inputs, review roles and the path from draft to approved work.

Pilot planning

Agree a focused trial with meaningful measures and a review point.

Dr. Tony Bader

AI governance and digital transformation expert, researcher and speaker. His work connects responsible AI adoption, communication and organisational capacity building.

START A CONVERSATION

badertony.com

Mention this book and the task you want to improve.