

Automate Your Admin

How to spot and remove the repetitive work slowing everyone down



THE PRODUCTIVITY OPPORTUNITY

Across 18 realistic professional tasks, people using AI completed 12.2% more work, finished 25.1% faster and produced outputs rated more than 30% higher in quality.

Dell'Acqua et al., 2026 | Controlled experiment with 758 consultants

The Cost That Never Appears in Accounts

No budget line reads “retyping”. Nothing in the management accounts says “chasing an update” or “copying figures between systems”. Yet in almost every organisation, skilled people spend a substantial part of their week doing precisely that, and because it's paid in hours rather than invoices it never gets challenged.

The work is worth challenging. Not because the people doing it are inefficient, but because it accumulated by accident, one reasonable decision at a time, and nobody has ever been given the job of looking at it as a whole.

This guide is a method for doing that. How to find the work, how to judge which parts of it can safely be handed over, how to sequence the changes so nothing breaks and how to make sure the hours you free up turn into something worth having.

72%

of UK organisations using AI apply it to administration, level with marketing as the most common area (DSIT, 2026)

25 to 40%

of the time spent on content, planning and reporting saved at one consumer products company (BCG, 2025)

30%


of organisations check 20% or less of AI output before it's used (McKinsey, 2025)

What's Inside

- 01 Where the Work Hides
- 02 Finding It: A One-Week Audit
- 03 The One-Page Test
- 04 Sequencing: What to Do First
- 05 Building In the Check
- 06 Where the Hours Go

Where the Work Hides

Before automating anything, it helps to understand why this work is invisible in the first place. That's what explains why it's still there.



Nobody Owns It, So Nobody Reviews It

Ask who's responsible for your organisation's admin and you'll rarely get a clean answer. It isn't a department. It's a layer that formed over years: a funder asked for a return, so a spreadsheet appeared. Two systems didn't talk, so somebody bridged them by hand. A report was useful once, so it's produced monthly a decade later.

Every one of those decisions made sense on the day it was made. Together they've built a quiet second job sitting on top of everyone's real one, and because no single person owns the whole of it, no single person has ever questioned it.

Your Most Capable People Are Absorbing It

The uncomfortable part is who ends up doing this work. It's your finance lead formatting a board pack. Your best manager writing up minutes. You, at nine in the evening, moving numbers into a spreadsheet because it's quicker than explaining it to somebody else.

The cost isn't really the task, then. It's that the people best placed to improve the organisation are the ones with the least time to do it, and the thing consuming that time doesn't appear anywhere you'd think to look.

This Is Where the Technology Is Genuinely Strong

It's worth being clear that this isn't speculative. Administration is now the joint most common place UK organisations put AI, named by 72% alongside marketing, and 70% describe what they use it for as administrative and support work (DSIT, 2026).

The results follow. Across functions, a majority of organisations using AI report cost reductions where it's deployed, including 61% in supply chain and 58% in service operations (McKinsey,

2025). Repetitive, rule-following, document-heavy work is precisely what this technology handles well. The hard question was never whether it could, but which of your tasks qualify, which is what the next two sections settle.

One Warning Before You Start

Admin is the right place to begin and the wrong place to stop. Around 70% of AI's potential value sits in core operating functions rather than back-office support, up from 62% the year before (BCG, 2025).

So treat the admin layer as the training ground it is. You'll learn how this works, build the review habits and free the capacity, and the capacity is what lets you go after the bigger prize in the operation itself.

Signs You're Looking at Automatable Work

If several of these describe a task, it belongs on your shortlist.


- The same steps happen every week, with only the content changing
- Someone moves information from one place to another without changing it
- The output is a document that follows a familiar shape
- People describe the task with the phrase "it's just"
- It exists because two systems don't talk to each other
- Nobody can remember who decided it should be done this way
- It's done in the evening or at the weekend because there's no room in the day

WHAT HAPPENS TO THE TIME SAVED

When organisations track what happens to the time AI saves, it is most often spent on entirely new activities. The first return is capacity for work that wasn't being done before.

McKinsey & Company, 2025 | The State of AI

Finding It: A One-Week Audit



You can map this yourself, in a week, without buying anything or commissioning a study. The point is not precision. It's getting the invisible written down.

Ask People, Don't Guess

The fastest route to an accurate picture is asking the people doing the work where their week disappears. They know. They've usually known for years, and nobody's asked in a way that suggested anything might change.

Frame it carefully. If people suspect the exercise is about headcount, you'll get a carefully curated answer. Say plainly that you're looking for work worth removing so they can get on with the rest, and mean it.

Capture Frequency and Time, Roughly

For each task you want three things: how often it happens, roughly how long it takes and who does it. Rough is fine. A task done daily for twenty minutes by three people is a meaningful number even if every figure in it is an estimate.

Frequency matters more than duration when you're choosing where to start. A four-hour monthly task looks bigger than a fifteen-minute daily one, but the daily task consumes more of the year and the saving arrives more often.

Look at the Joins Between Systems

Some of the richest material sits in the gaps: the moment information leaves one system and gets manually re-entered into another, or where a person acts as a translator between two teams who each keep their own records.

These joins are almost never anyone's job officially, which is exactly why they never get examined. They're also frequently the single largest time cost in a process and often the place where errors originate.

The Week

Day	What to do	What you end up with
Monday	Explain the exercise to everyone involved and say clearly what it's for	People willing to give you honest answers rather than defensive ones
Tuesday	Ask each person to note tasks that repeat, with rough frequency and duration. No formal template needed	A raw list, longer than you expected
Wednesday	Walk one or two processes end to end, following the information rather than the organisation chart	The joins between systems, which people rarely list themselves
Thursday	Group the list: document handling, data movement, summarising, record keeping, correspondence, reporting	A shortlist of eight to twelve candidates
Friday	Score each against the one-page test in the next section. Pick three to look at properly	Three candidates and a clear reason for choosing each

SECTION 03

The One-Page Test

There's a simple test for whether a task is a good candidate, and it takes about a minute to apply.

WHAT THE SAVING CAN LOOK LIKE

A consumer-products company cut the time spent on content creation, planning and reporting by 25% to 40%, while doubling its speed to market.

Boston Consulting Group, 2025 | Organisational example reported in the study

Could You Write the Instructions on One Page?

If you could hand a competent new starter a single page of instructions and expect the task done correctly, it's a strong candidate. If the honest answer is that they'd need to sit with someone for a fortnight and absorb how things work around here, it isn't. Not yet.

The test works because it's really asking whether the task depends on context that exists only in someone's head. Where it does, automation reproduces the visible steps and quietly loses the invisible judgement that made the output correct.

What Reliably Qualifies

In most organisations the same categories come up. Pulling specific information out of documents. Summarising what's arrived, whether that's applications, enquiries, reports or correspondence. Keeping trackers and records current. Producing meeting notes and the follow-up actions from them. Drafting the routine correspondence that fills the gaps between everything else.

The evidence supports starting here. One consumer products company saved 25% to 40% of the time spent on content creation, brand planning and reporting. It also doubled its speed to market for campaign activation (BCG, 2025). The pattern repeats because the work repeats, which is the entire point.

What Should Stay With People

Anything requiring judgement across conflicting information. Anything where being wrong carries real consequence. Anything depending on relationships, or on knowing why a decision was made three years ago.

There's hard evidence for this boundary. In a controlled experiment with 758 consultants, AI improved output on suitable tasks by more than 30%, but on a task requiring reconciliation of conflicting sources it caused a 19 percentage point drop in correct answers (Dell'Acqua et al., 2026). The aim is never to automate a job. It's to strip the mechanical layer off it so the judgement gets the time it deserves.

Good Candidate Against Poor Candidate

Strong candidate	Leave it alone for now
Instructions fit on one page	Requires a fortnight of shadowing to learn
Same shape every time, only the content varies	Every instance is different in kind
The information needed is written down somewhere	It depends on context in someone's head
A wrong output is annoying and easily spotted	A wrong output is costly or hard to notice
It happens weekly or daily	It happens twice a year
Internal, with a person reviewing before it goes out	It goes straight to a customer or regulator unchecked
Judgement sits before or after the task	Judgement is the task

Sequencing: What to Do First

Having found several candidates, the order you tackle them in matters more than which ones you choose. Two factors decide it.

High Frequency, Low Risk, First

Rank your candidates on how often the task happens and how bad it would be to get it wrong. Start in the corner where frequency is high and risk is low. That combination gives you the fastest visible saving with the smallest downside if the first version needs adjusting.

It also builds the thing you'll need for everything after: a story people in your organisation have seen work with their own eyes.

One Task, Properly, Rather Than Six Halfway

Organisations achieving real returns achieve a 76% higher match between where they deploy AI and where it delivers impact, and 62% of their initiatives are actually live against 12% among those struggling (BCG, 2025).

Running six small experiments doesn't spread your risk. It divides your attention six ways and gives none of them the support needed to survive contact with real work. Finish one. The second is always cheaper than the first.

Remove the Steps That Are Now Unnecessary

This is where most of the benefit is won or lost. If a report now assembles itself, the meeting that existed to assemble it needs to change. If a tracker updates automatically, the weekly chase email should stop.

Organisations routinely add AI to a process while keeping every step AI made redundant, then conclude the technology didn't work. The tool isn't the change. Removing what it made unnecessary is the change.

Buy, Configure or Have It Made

Most admin automation starts with something you can buy, though few organisations expect it to fit as it comes. 85% expect to customise AI to their own needs rather than run it off the shelf (Deloitte, 2026).

The choice usually comes down to how particular your process is. Where the workflow is standard, buy it and move on. Where the work is genuinely yours, or the data can't leave, something made to fit will earn its keep, and it tends to come in below what people brace themselves for.

A Worked Example

One task, followed from discovery through to live, at a realistic pace.

Stage	What happens
The task	Monthly performance report. Someone pulls figures from three systems, formats them into a template, writes a short commentary and circulates it. Around six hours, every month, done by a manager
Why it qualifies	Same shape every time. Instructions fit on a page. The information is already recorded. A wrong figure would be spotted at review. It happens reliably
What changes	Figures are gathered and assembled automatically. A first-draft commentary is generated from the data. The manager reviews it, corrects the interpretation and adds the context only they have
What stays human	The commentary's judgement. What matters this month, what's worth flagging and what the numbers don't show
What gets removed	The meeting that existed to assemble the figures. The chase emails for late data. The reformatting
The result	Around an hour instead of six, and the report arrives earlier in the month, which is what people actually noticed

SECTION 05

Building In the Check

Automation without a review step is where organisations get hurt. The good news is that a proper check costs minutes and catches most of what would otherwise reach somebody else's desk.

THE HUMAN REVIEW BOUNDARY

On a task that required people to reconcile spreadsheet data with conflicting interview evidence, AI users were faster but 19 percentage points less likely to reach the correct answer.

Dell'Acqua et al., 2026 | Controlled experiment with 758 consultants

Polished Output Is Not Correct Output

The most useful finding in the research is also the most counterintuitive. In the controlled experiment, AI-assisted work scored higher on persuasiveness and internal coherence than unassisted work whether the underlying answer was right or wrong, with coherence scores rising between 17.9% and 25.1% (Dell'Acqua et al., 2026).

In other words, reading something and judging that it reads well is no longer a check. It never really was, but AI removes the rough edges that used to signal a problem. Review the reasoning and the figures, not the prose.

Decide What Gets Checked, and by Whom

30% of organisations check 20% or less of AI-produced content before it's used, and 47% have already experienced at least one negative consequence from using it (McKinsey, 2025). The survey reports those separately and doesn't link them. They're still the two numbers to put side by side when you decide what gets checked.

You don't need to check everything. You need to decide, in advance and in writing, what must be checked: anything leaving the organisation, anything informing a decision that carries consequences, anything a regulator or auditor might ask about. Everything else can run with sampling.

Most Organisations Check, Few Write It Down

In the UK, 84% of organisations using AI report at least some human input or checking and 67% report significant checking, while 2% report none at all (DSIT, 2026). In the interview strand of the same research, only a few businesses described having a specific procedure for checking outputs, usually a checklist.

That's the gap worth closing. An unwritten habit depends on the conscientiousness of whoever happens to be doing the task, and it disappears the week they're on leave.

Sample the Rest, and Log What Breaks

For routine internal output, pull a handful at random each week and look at them properly. You're watching for drift: cases the workflow was never designed for, or a pattern of small errors suggesting something upstream has changed.

Keep a simple log of every case the workflow couldn't handle. That log is genuinely valuable. It tells you what to fix next, and it's usually where your second automation project comes from.

Review Checklist


Use this on anything before it leaves the desk.

- Do the figures match the source, spot-checked against the original?
- Is anything asserted that the underlying data doesn't actually support?
- Is anything important missing rather than wrong? Omissions are harder to see than errors
- Would this reasoning survive being questioned by someone who knows the subject?
- Has anything sensitive been included that shouldn't leave the organisation?
- If this is wrong, who finds out and how quickly?

SECTION 06

Where the Hours Go

The last question is the one people in your organisation will be asking privately from the moment you start. It deserves a straight answer.



The Time Mostly Becomes New Work

Time saved is most often spent on entirely new activities, then on the work that hasn't been automated (McKinsey, 2025). The person freed from maintaining a tracker starts improving the process behind it. The manager who stops writing up minutes starts running better meetings.

Don't promise more than that. The same research finds headcount reduction is one of the attributes most strongly associated with value realised, and expectations of reduction concentrate in service operations and supply chain. Say what's true where you are, and say it

early, because an assumption left unaddressed becomes quiet resistance that's much harder to shift later.

The Second Benefit Is the One Nobody Forecasts

Three quarters of UK organisations using AI report improved workforce productivity and 57% report new or improved processes, but 19% also report improved staff retention or wellbeing (DSIT, 2026). That last figure rarely appears in a business case.

It should. Nobody joined your organisation to reformat a spreadsheet. Taking that work off a capable person changes how they feel about the job, and that shows up in who stays.

Be Honest About Entry-Level Work

There's one area that deserves care. The tasks being automated first, data entry, reconciliation and first-line queries, are often where careers start (Deloitte, 2026). Remove them without thinking and you've closed a route into the profession without meaning to.

The answer isn't to preserve the work. It's to be deliberate about what a first year looks like instead, and to give junior people the reviewing and judgement work sooner than you otherwise would. Handled well, that's an improvement rather than a loss.

Speed Changes Behaviour More Than Savings Do

One effect consistently surprises people. When a report takes minutes rather than days, people start asking for it when they actually need it rather than waiting for the monthly cycle. Decisions get made on current information instead of last month's.

That shift is usually worth more than the hours saved, and it never appears in the business case, because nobody thinks to put it there.

Questions People Will Ask, and Straight Answers

What they'll ask	What's true
Is this about cutting jobs?	Say what's actually true in your case, plainly. If it isn't, say so and mean it. If some roles will change, say that too. Evasion costs more trust than bad news
What if it gets something wrong?	It will, occasionally. That's why there's a review step, and why we're starting with work where a mistake is annoying rather than serious
Am I expected to check everything now?	No. Here's exactly what gets checked and what gets sampled, written down so nobody has to guess
What do I do with the time?	The work you've been meaning to get to. Name something specific, because a vague answer sounds like a euphemism

Give the Work Back Its Shape

None of this requires a technology strategy. It requires a week of looking properly at work everyone already knows is there, a test you can apply in a minute and the discipline to remove the steps that stopped being necessary.

The benefit compounds in a way one-off savings don't. A task done weekly by four people is a task automated once and saved every week after, for everyone, without anyone needing to remember.

You'll know it's working long before it reaches a report, because somebody will tell you they got their Friday afternoon back, and they'll be doing something better with it.

Curious how much of the week is automatable? Ask us, it's a shorter conversation than you'd expect.

Jacobs & Xu

Applied AI. Results You Can Measure.

contact@jacobsandxu.com | +44 (0)203 105 7544

Sources

References are given in the text by author and year. Full details are below, in Harvard format.

Boston Consulting Group (BCG) (2025) *The Widening AI Value Gap: Build for the Future 2025*. Boston: Boston Consulting Group. Survey of 1,250 senior executives and AI decision makers across 68 countries and 25 or more sectors. Published September 2025.

Deloitte AI Institute (2026) *State of AI in the Enterprise: The Untapped Edge*. Deloitte Insights. Survey of 3,235 board, C-suite and director-level respondents across 24 countries, fielded August to September 2025. Published January 2026.

Dell'Acqua, F., McFowland, E., Mollick, E., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Kraymer, L., Candelon, F. and Lakhani, K. (2026) 'Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of Artificial Intelligence on Knowledge Worker Productivity and Quality', *Organization Science*, 37(2), pp. 403 to 423. Preregistered randomised controlled experiment with 758 consultants.

Department for Science, Innovation and Technology (DSIT) (2026) *AI Adoption Research*. London: DSIT. Research conducted by IFF Research and Technopolis Group. Telephone survey of 3,500 UK businesses with five or more employees, fielded 12 February to 2 May 2025, with 100 follow-up interviews. Published January 2026, updated February 2026. Available at: <https://www.gov.uk/government/publications/ai-adoption-research> (Accessed: 12 August 2026).

McKinsey & Company (2025) *The State of AI: How Organizations Are Rewiring to Capture Value*. Survey of 1,491 participants across 101 nations, fielded July 2024. Published March 2025.

Figures drawn from executive surveys are self-reported by the organisations taking part. Where this guide describes what high-performing organisations do differently, it describes a pattern found among them rather than a guaranteed result. Four points of context are worth carrying with the numbers. Every Deloitte respondent already runs at least one AI implementation, so its access and usage figures are an upper bound rather than an average. BCG's sample is weighted towards large enterprises, so its findings read better as principles than as benchmarks. McKinsey's fieldwork was carried out in July 2024 and adoption has risen since. DSIT figures cover UK businesses with five or more employees and measure deliberate deployment of a recognised AI technology, which is a narrower definition than most commercial surveys use.