

FROM BYOAI TO BUSINESS VALUE

Channelling Bottom-Up Adoption into Enterprise Impact

31 MAY 2024 • ENTERPRISE AI SERIES



About this paper

This is the seventh of nine papers on the enterprise adoption of generative artificial intelligence.

The previous paper concerned accountability at the top of an organisation. This one concerns what has been happening at the bottom of it, largely unmeasured, for about 18 months.

One large survey carries most of the evidence and I have been careful about what it can and cannot support. The channelling sequence, the value tree and the dashboard are my own work and are tagged accordingly.

HOW TO CITE

Forrest, P. (2024). *From BYOAI to Business Value: Channelling Bottom-Up Adoption into Enterprise Impact*. Enterprise AI, 2023 to 2024, Paper Seven. Version 1.0, 31 May 2024.

HOW TO READ THE EVIDENCE TAGS

Every substantive claim carries a tag stating what kind of evidence stands behind it. A claim resting on two kinds carries both.

SURVEY Survey or polling data, cited with its sample and fieldwork dates.

LEGAL A statute, regulation, executive order, legislative proposal or court decision.

AUTHOR My own framework or judgement, proposed and not yet proven.

A NOTE ON DATES

Research for this paper closed on 28 May 2024. Its central source is a vendor-sponsored survey, and the text states the sponsor, the sample, the fieldwork window and the denominator of every figure taken from it, because the most quoted statistic in this area is routinely reported against the wrong denominator.

Contents

FRONT MATTER

Abstract and summary of the argument	4
--------------------------------------	---

THE PAPER

1	What the May survey says, and how to quote it	5
2	Five reasons people bring their own, only one of which is defiance	6
3	Channel rather than chase	8
4	Redesign the workflow, not just the tool access	10
5	An adoption dashboard that survives a sponsor	12
6	What changed in Europe this month	14
7	Reciprocity, and what the organisation owes back	15

END MATTER

Method, evidence and limits	16
Sources considered and set aside	17
About the author	18
References	19
Further sources consulted	20

ABSTRACT AND SUMMARY OF THE ARGUMENT

Three quarters of surveyed knowledge workers, and the question of whose tools they are using

A survey published on 8 May 2024 reports that 75 per cent of the global knowledge workers it sampled use artificial intelligence at work, and that use of generative tools has nearly doubled in six months.¹ **SURVEY** The same study reports that 78 per cent of those users bring their own tools, rising to 80 per cent at small and medium-sized companies.¹ **SURVEY** The sponsor sells products in this market, and the sampled population is knowledge workers only. Both facts belong beside the numbers.

Read carefully, the finding is not that employees are behaving badly. It is that demand has outrun supply for about 18 months, and that people under delivery pressure have solved their own problem in the way people always do. The organisation's position is that it now has the risk without the visibility and without the benefit.

This paper sets out how to convert that into something governed. It proposes a discovery method that does not rely on self-report alone, a channelling strategy that provides an approved route faster than it closes the unapproved one, a workflow redesign method, a value tree and an adoption dashboard that distinguishes logins from meaningful use.

WHAT THIS PAPER ARGUES

Bottom-up adoption becomes enterprise value when the organisation channels it. Attempting to stop it removes the visibility without removing the behaviour, and the evidence that it is happening anyway is now hard to argue with.

- The headline needs reading against the right denominator. The bring-your-own figure of 78 per cent is 78 per cent of artificial intelligence users rather than of knowledge workers generally, and reported the second way it overstates the position considerably.¹ **SURVEY**
- The doubling is about generative tools specifically. Use of generative artificial intelligence at work nearly doubled in six months, and 46 per cent of users began within the last six.¹ **SURVEY**
- The approved route should exist before the unapproved one is closed. An organisation that blocks first and provides later gets six months of invisible use and a workforce that has learned not to mention it. **AUTHOR**
- Adoption should be measured as repeat use on real work. Licence activation measures procurement, weekly repeat use on a named task class measures adoption, and the two differ by more than most sponsors expect. **AUTHOR**
- The European position changed on 21 May and is still not law. The Council adopted the regulation, completing the legislative procedure, and neither signature nor publication had happened when this paper closed.² **LEGAL**

SECTION 1

What the May survey says, and how to quote it

A large sponsored study with real findings and a denominator that is routinely misreported.

Microsoft and LinkedIn published their annual work trend report on 8 May 2024, based on a survey of 31,000 full-time employed or self-employed knowledge workers across 31 markets, conducted by an independent research firm between 15 February and 28 March 2024, alongside labour market data, product telemetry and research with large corporate customers.¹ **SURVEY**

Its headline findings are that 75 per cent of those surveyed use artificial intelligence at work, that use of generative tools has nearly doubled in the preceding six months, and that 46 per cent of users started within that period.¹ **SURVEY** Taken together those describe a fast and recent shift in behaviour.

The denominator that matters

The bring-your-own finding is reported in the study as 78 per cent of artificial intelligence users bringing their own tools to work, rising to 80 per cent at small and medium-sized companies.¹ **SURVEY** It is 78 per cent of users, not 78 per cent of knowledge workers, and the difference is the population who use nothing at all. Reported the second way it becomes a considerably larger claim than the study makes.

Two qualifications belong permanently attached to these numbers. The population is knowledge workers rather than all employees, which excludes most of the workforce in manufacturing, logistics, retail and care. And the study is published by a company selling the products it is measuring demand for, which does not invalidate it and does belong in the citation.

Why this is a supply problem in the shape of a compliance problem

The instinctive reading in a risk function is circumvention of controls. The more useful reading is that for 18 months demand has exceeded approved supply, and that people with work due have resolved the difference themselves. The behaviour is what any competent employee does when the sanctioned route is slower than the deadline.

That reframing changes the response. A compliance problem is addressed with prohibition and monitoring. A supply problem is addressed by supplying something, and the prohibition only works afterwards. Organisations that have taken the first route for the past year have generally produced invisible use, and I have yet to see the volume fall. **AUTHOR**

SECTION 2

Five reasons people bring their own, only one of which is defiance

The category that attracts the policy attention is the smallest one.

Before designing a response it is worth knowing which problem is being solved. The five causes below call for entirely different remedies, and a single policy response addresses one of them.

WHY PEOPLE BRING THEIR OWN, AND WHAT EACH REASON REQUIRES

REASON	WHAT IT LOOKS LIKE	WHAT ACTUALLY ADDRESSES IT
No approved tool exists	The organisation is still evaluating, 18 months in, and people have work due on Friday.	Provide something approved, even if limited. A constrained approved tool beats an unconstrained personal account.
The approved tool is worse	A restricted internal deployment that is slower, older or refuses more often than the free consumer product.	Close the capability gap or accept that the policy will be ignored. This is a product problem rather than a compliance problem.
Access friction	Approval takes three weeks and the task takes an afternoon.	A fast lane for low-risk use, with registration rather than assessment. Friction is the whole of the cause here.
Nobody said no clearly	Guidance exists, is long, sits on an intranet, and does not address the specific case.	One page, positively phrased, naming what may be entered and what may not. Length is inversely related to compliance.
Deliberate circumvention	A known prohibition, knowingly breached, usually under deadline pressure.	A management conversation rather than a technical control. This is the smallest category and it attracts most of the policy attention.

The second row is unwelcome

Where an organisation has deployed an approved assistant that is slower, older or more locked down than what an employee can obtain free in a browser, no amount of policy will hold the line. That is a product quality problem and it should be escalated as one, to whoever owns the internal deployment, with the capability gap documented. Calling it user resistance closes the file without fixing anything. AUTHOR

Discovery that does not depend on people volunteering

Self-report understates, because admitting unapproved use has a cost and denying it has none. Four methods together produce a usable picture, and none of them alone does.

1. An anonymous pulse survey with an explicit statement that responses carry no consequence, which is worth stating twice because nobody believes it the first time.
2. Structured interviews with a dozen people across functions, asking what they do. The task-level question surfaces use that the tool-level question does not.
3. Task diaries for two weeks in one function, which is intrusive and accurate and should be used sparingly.

4. Network telemetry with prior notice, because notice is not a courtesy here. Monitoring introduced without it converts a supply problem into an employee relations problem, and the second is considerably harder to fix.

Expect the picture to be larger than the telemetry suggests and smaller than the survey headline implies. Personal devices are invisible to the first, and the survey population is not the organisation's population.

SECTION 3

Channel rather than chase

The approved route has to be quicker to use than the unapproved one, and the rest of this section follows from that.

Channelling means making the sanctioned path the path of least resistance. It is a product design exercise rather than a control exercise, and it fails whenever the governance function owns it alone.

CHANNELLING RATHER THAN CHASING



FIGURE 1 The approved route has to be faster to use than the unapproved one, or the policy is a statement of preference. Each step here is a reduction in friction rather than an increase in control.

My framework. **AUTHOR**

1. Provide something approved quickly, even if it is capability-limited. A restricted approved tool available this month beats an excellent one available in the autumn.
2. A fast lane for low-risk use, meaning registration, decided in days. Reserve the full assessment for the high-consequence tier where it belongs.
3. Data zones stated positively, in three categories at most, covering what may go into an approved tool, what may go into an approved tool under conditions, and what may not leave the organisation at all.
4. Templates and prompt libraries for the tasks the discovery work found people actually doing, which removes the main practical advantage of the consumer product.

5. Office hours and champions, staffed with allocated time. Goodwill runs out in a quarter.
6. An amnesty, stated clearly and meant, because the first person disciplined under the new policy for something they did in March ends the flow of information permanently.

Sequencing, which is where this usually goes wrong

The common sequence is to announce a prohibition, begin an evaluation, and provide an approved tool some months later. That order produces the worst available outcome. Use continues, becomes invisible, and the eventual approved tool arrives to a workforce that has already built habits around something else and has learned that the official channel is slow. Reverse it wherever the risk tolerance permits. **AUTHOR**

SECTION 4

Redesign the workflow, not just the tool access

Providing a tool changes what people can do. Redesigning the task changes what the organisation gets.

Bottom-up adoption produces individual time savings that rarely aggregate into anything visible at the level of a team or a process. The reason is structural. An individual who saves 20 minutes on a draft returns the time to a queue whose throughput is set elsewhere, and the saving disperses.

SPECIFICATION

What a workflow redesign actually requires

- A baseline for the whole workflow, because a step-level saving in front of an unchanged bottleneck produces nothing.
- An explicit division between what the system does and what the person does, agreed with the people who do it now.
- The verification rung named, with the time for it added to the schedule. It will not come out of the saving.
- An exception path for the cases the system handles badly, and a record of how often that path is used, because the frequency is the finding.
- A job quality check. Whether the work became better or worse to do, asked of the people doing it, more than once and after the novelty has gone.

The fourth item is most often left out. A redesigned workflow that works on 85 per cent of cases and has no route for the remaining 15 will either fail visibly or, more commonly, produce quiet workarounds that reintroduce every risk the redesign was meant to remove. **AUTHOR**

THE VALUE TREE

FROM BYOAI TO BUSINESS VALUE

From task change to business outcome

A task got faster

▼ only if the task sat on the critical path

The workflow got faster

▼ only if the workflow was the binding constraint

The team produced more

▼ only if the output was wanted at that volume

The business gained

▼ and the figure reaches a board net of the risk adjustment

The number a board should see

Most published value claims stop two levels below where they are reported.

FIGURE 2 From task-level change to business outcome, with the risk adjustment applied before the number reaches a board. Most published value claims stop at the second level and are reported as though they reached the fourth.

My framework. **AUTHOR**

From task change to business outcome

The value tree exists to make the inference chain explicit, because most published value claims stop two levels below where they are reported. A task got faster. The workflow got faster only if the task was on the critical path. The team produced more only if the workflow was the constraint. The business gained only if what the team produces more of is something the business is short of. Each link is an empirical question and each is commonly assumed.

SECTION 5

An adoption dashboard that survives a sponsor

A dashboard needs seven measures, and the two that vendors supply are the two that predict least.

Adoption reporting in this field has converged on activation and monthly active users, both of which are available from the supplier and neither of which indicates that any work has changed. A dashboard built on them will show a strong first quarter and a confusing second.

AN ADOPTION DASHBOARD THAT DISTINGUISHES USE FROM LOGINS




FIGURE 3 Seven measures. Activation and first use are the two that vendors report and the two that predict least. Repeat use, task depth and verified quality are what indicate that something has changed.

My framework. **AUTHOR**

SEVEN ADOPTION MEASURES, AND WHAT EACH IS WORTH

MEASURE	DEFINITION	WHAT IT TELLS YOU
Licence activation	Accounts provisioned and signed into at least once.	Almost nothing. It measures procurement and is the figure most often presented as adoption.
Weekly active use	Distinct users performing at least one task in a week.	Whether the tool has entered the working week. Watch the trend after the novelty period rather than the level.
Repeat use on a task class	Users returning to the same named task class across three consecutive weeks.	The first measure that indicates the workflow has changed rather than been sampled.
Task depth	Proportion of use on multi-step work rather than single-shot drafting.	Whether capability is being used or a search box is being used.
Verified quality	Sampled assisted outputs reviewed blind against the prior standard.	Whether the work got better, which no usage metric can tell you.
Skill progression	Movement across the levels of the capability rubric, assessed on observed work.	Whether the organisation is building something durable or renting a habit.
Incidents and near misses	Confidential material entered, unverified output used, complaint received.	Whether the channelling worked. A rate of zero alongside high use means the reporting route is not functioning.

The measure to watch after the novelty period

Repeat use on a named task class across three consecutive weeks separates a tool that has entered the work from one that was tried. It is also the measure that falls most sharply around the eighth to twelfth week. By then the initial enthusiasm has worn off and only genuine utility remains. Sponsors should be warned about that decline in advance, because it arrives reliably and is read as failure when it is actually the first honest reading. **AUTHOR**

Incidents at zero

An organisation reporting high adoption and no incidents or near misses has a reporting problem rather than an unusually careful workforce. Near misses arise from ordinary confusion and should appear within weeks. Their absence means the route to report one is unclear, unused, or believed to carry a cost, and it is worth finding out which.

SECTION 6

What changed in Europe this month

The Council adopted the regulation on 21 May, which completed the legislative procedure and did not make it applicable.

The vocabulary in circulation will mislead a board. On 21 May 2024 the Council of the European Union formally adopted the artificial intelligence regulation, completing the legislative procedure.² **LEGAL** The European Parliament had adopted its first-reading position on 13 March 2024.³ **LEGAL**

As at the date of this paper the measure awaits signature by the presidents of both institutions and publication in the Official Journal. It has not entered into force. No obligation under it applies to any organisation today. The entry into force date follows publication by a fixed period, and since publication has not occurred that date cannot yet be stated. **LEGAL**

A DRAFTING NOTE FOR ANYONE WRITING THIS QUARTER

The regulation has no official citation number yet. Numbers are assigned on publication in the Official Journal, and a board paper citing one this month is citing something that does not exist.

LEGAL

Equally, descriptions such as passed into law, came into effect, or ratified are each wrong and wrong in different ways. Adopted, signed, published, in force and applicable are five distinct states, and an organisation planning against the wrong one will either act too early or discover a deadline late.

For a programme channelling bottom-up adoption, the practical implication is modest and worth stating. Nothing has to be done differently this quarter for this reason. What changes is that the inventory and the evidence record built for governance purposes now have a second use, and organisations without one have less time to build it than they did a month ago.

SECTION 7

Reciprocity, and what the organisation owes back

The programme asks employees to disclose what they have been doing. It should be clear what they get in return.

Everything in this paper depends on people telling the organisation what they are actually doing, and none of it works if that disclosure is one-directional. The practical bargain has four parts and organisations tend to offer the first two.

1. Tools that work, provided quickly, competitive with what people can get themselves.
2. Training that is worth the time, on their own tasks.
3. Honesty about the purpose, so that where the objective includes reducing headcount it is said, and where it does not that is said too, with something specific about what happens to time that is saved.
4. A say in the redesign, including the ability to say that a proposed change makes the work worse, without that being treated as resistance.

The third determines whether any of this holds. Employees are being asked to reveal efficiencies they have found, in an environment where the public discussion of this technology is largely about job displacement. An organisation that will not answer the question directly should expect the discovery exercise to return a flattering and useless picture. **AUTHOR**

What a good position looks like at the end of the year

An approved route that most people use because it is easier than the alternative. A discovery picture that the organisation trusts, and refreshes. Two or three workflows redesigned properly with baselines and exception paths. A dashboard that distinguishes repeat use from activation. And an incident record with entries.

Whether the underlying demand keeps growing at the rate the May survey describes is not something this paper can know, and the six-month doubling is unlikely to repeat indefinitely. What seems safer to assume is that the gap between what employees do and what organisations can see will keep closing only where somebody deliberately closes it.

END MATTER

Method, evidence and limits

How the sources in this paper were chosen and dated, and what the reader should distrust.

Selection

One large sponsored survey carries most of the empirical weight, cited with sponsor, population, sample, fieldwork window and method. Legal status is taken from the institutions themselves. Where a figure has a denominator that is commonly misreported, the text states the denominator in the sentence.

On the central source

The work trend report is published by a company that sells artificial intelligence products to the organisations it surveys, and its telemetry strand draws on that company's own products. Neither fact disqualifies it, and no comparable independent measurement of workplace adoption at this scale exists. Both facts belong beside every figure taken from it.

Legal status

The European regulation is described as adopted by both institutions and awaiting signature and publication. It is not in force at this date. The paper avoids the official citation number because none has been assigned, and avoids stating an entry into force date because it depends on a publication that has not happened.

What nobody has measured

No independent measurement of bring-your-own adoption exists outside vendor-sponsored research. No study measures whether channelling strategies reduce unapproved use, which is the central empirical claim in section three and is offered as reasoning. And no evidence addresses the reciprocity argument in section seven.

Hypotheses, not findings

The five causes of bring-your-own behaviour, the channelling sequence, the workflow redesign requirements, the value tree and the adoption dashboard are my own. The claim that repeat use falls around the eighth to twelfth week is a practitioner observation across a small number of engagements and should be treated as a hypothesis.

A declaration of interest

I advise on technology adoption and could supply the discovery and redesign work described here. The paper argues for providing tools quickly, which shortens the evaluation engagements that pay better, so whatever bias I carry is not obviously commercial.

END MATTER

Sources considered and set aside

Surveys I read alongside the central one and kept out of the text, with the reason for each.

SOURCE	PUBLISHED	WHY IT WAS SET ASIDE
Pew Research Center, Americans' use of ChatGPT is ticking up, but few trust its election information	26 March 2024	A consumer survey, United States only. It reports that 23 per cent of adults had ever used the tool, which is a different population from the one this paper is about.
Deloitte, State of Generative AI in the Enterprise, second quarterly report	April 2024	Leaders reporting on their organisations. It is a survey of the people this paper says are the last to know.
Boston Consulting Group, AI at Work: What People Are Saying	June 2023	A consultancy panel of about 12,900 people, fielded a year before this paper. Useful for the shape of the sentiment and too old for the adoption figure.
Salesforce, Generative AI snapshot on the promises and pitfalls of AI at work	15 November 2023	Vendor polling reporting that more than half of generative users at work were using tools their employer had not approved. Consistent with the central source, and from a company with the same commercial interest, so it cannot corroborate it.

END MATTER

About the author

Paul Forrest advises boards and executive teams on the adoption of new technology in operating businesses, with a practice grounded in delivery rather than in advisory work alone.

His background is in business process reengineering, which he has practised since the 1990s, and in applied natural language processing, which he has worked on since 2014. The workflow argument in section four comes from the first of those. Reengineering work showed, repeatedly and expensively, that improving one step of a process yields nothing unless that step was the constraint, and the current wave of task-level productivity claims is making the same mistake with better tools.

He writes on the condition that a claim carries its evidence, and this series is an attempt to hold that line while a subject moves quickly.

END MATTER

References

Three works cited in the text, each with its original date of publication. None is later than 28 May 2024.

1. **Microsoft and LinkedIn.** *AI at Work Is Here. Now Comes the Hard Part.* Work Trend Index annual report. Survey of 31,000 full-time employed or self-employed knowledge workers across 31 markets, conducted by Edelman Data and Intelligence, fieldwork 15 February to 28 March 2024, with LinkedIn labour market data, Microsoft 365 productivity signals and research with large corporate customers.
PUBLISHED 8 MAY 2024
2. **Council of the European Union.** Formal adoption of the regulation laying down harmonised rules on artificial intelligence, completing the legislative procedure.
PUBLISHED 21 MAY 2024
3. **European Parliament.** Position adopted at first reading, 523 votes to 46 with 49 abstentions, Texts Adopted P9_TA(2024)0138.
PUBLISHED 13 MARCH 2024

END MATTER

Further sources consulted

Nine sources I read for this paper and did not cite. They are listed so that the reading behind it can be seen in full, and none is later than 28 May 2024.

- **United States Office of Management and Budget.** *M-24-10: Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence.*
PUBLISHED 28 MARCH 2024
- **Organisation for Economic Cooperation and Development.** *OECD Employment Outlook 2023: Artificial Intelligence and the Labour Market.* Cited for worker consultation practice, with its 2022 fieldwork window noted.
PUBLISHED 11 JULY 2023
- **Dell'Acqua, F., McFowland, E., Mollick, E., and others.** *Navigating the Jagged Technological Frontier.* Harvard Business School working paper 24-013.
PUBLISHED 18 SEPTEMBER 2023
- **Stanford Institute for Human-Centered Artificial Intelligence.** *Artificial Intelligence Index Report 2024.*
PUBLISHED 15 APRIL 2024
- **National Institute of Standards and Technology.** *Artificial Intelligence Risk Management Framework 1.0.* NIST AI 100-1.
PUBLISHED 26 JANUARY 2023
- **International Organization for Standardization and International Electrotechnical Commission.** *ISO/IEC 42001, Information technology, Artificial intelligence, Management system.*
PUBLISHED 18 DECEMBER 2023
- **The White House.** *Executive Order 14110.*
PUBLISHED 30 OCTOBER 2023
- **Colorado General Assembly.** Senate Bill 24-205, concerning consumer protections in interactions with artificial intelligence systems, signed 17 May 2024.
PUBLISHED 17 MAY 2024
- **Microsoft.** *Will AI Fix Work?* The previous year's edition, cited for comparison of the two fieldwork windows.
PUBLISHED 9 MAY 2023

The next paper takes regulation on its own terms

This paper touched the European position in one short section because it changed during the month. The eighth will take it properly, alongside the draft generative risk profile the United States standards institute published for comment in April, and build a practical crosswalk between the two. It is due at the end of August.

By then the regulation should have been signed and published, and the phased timetable can be stated from the text. Readers with European exposure should treat the inventory work in this paper as the input to that one.

ENTERPRISE AI SERIES

THE NEXT PAPER WILL ASSUME THESE EXIST

- A discovery picture built from more than self-report, with an amnesty that people believed.
- An approved route that is faster to use than the unapproved one.
- An adoption dashboard reporting repeat use on named task classes rather than activation.