Dear Erik, Avi, and Catherine,

The attached draft does not include any figures, data tables, or specific numbers in the text, as these elements are still under review at OpenAI. We had hoped to complete this review by the March 26 deadline, but unfortunately the process is taking slightly longer than anticipated. Nonetheless, we wanted to submit the most complete draft we could, and we would be grateful for the opportunity to update our submission in the coming days with more specific details from our analysis.

Thank you for your consideration.

Sincerely,
David Holtz (on behalf of Ronnie Chatterji, David Deming, Neel Rakholia, Prasanna Tambe, and Gawesha Weeratunga)

# How Generative AI Diffuses Inside Organizations: Firm-Level Evidence from ChatGPT Enterprise

Aaron Chatterji[1], David Deming[2], David Holtz[3,1],
Neel Rakholia[1], Prasanna Tambe[4,1], and Gawesha Weeratunga[1]

[1] OpenAI
[2] Harvard University
[3] Columbia Business School
[4] Wharton, U. Pennsylvania

March 27, 2026

DRAFT. RESULTS ARE PRELIMINARY. DO NOT POST OR CIRCULATE.

**Abstract**

We study how a general-purpose AI technology diffuses inside organizations using longitudinal usage data from ChatGPT Enterprise. Diffusion is staged: over tenure, cumulative adopters, weekly active users, total messages, and within-firm breadth all rise, while use expands across tasks and becomes increasingly embedded in workflows through broader use of tools, custom GPTs, and connectors. Later adoption cohorts move through these stages faster, reaching similar milestones in fewer weeks since adoption. To identify forces accelerating diffusion, we study two empirical settings. First, a matched-pair difference-in-differences design shows that formal enablement increases adoption, participation, and usage intensity while reducing concentration within firms. Second, using a temporal causal inference framework around the GPT-5 release in August 2025, we find that improvements in the model frontier also materially accelerate within-firm diffusion. Taken together, the results suggest that the diffusion of generative AI depends both on organizational capability-building and on continued technological progress.

# 1 Introduction

General-purpose technologies rarely generate immediate, economy-wide gains. Their impact unfolds through a slower process of co-invention in which firms discover use cases, invest in complements, and reorganize production so that a new capability becomes reliable in everyday work (Griliches, 1957; Mansfield, 1961; Hall and Khan, 2003; Jovanovic and Rousseau, 2005; Bresnahan and Trajtenberg, 1995). A growing literature argues that AI, and especially large language models, have the characteristics of a general-purpose technology (Goldfarb et al., 2023; Eloundou et al., 2024), suggesting that the same diffusion logic should apply here as well; although LLMs are flexible and broadly applicable, sustained organizational value will require workflow redesign, training, governance, and integration into existing tools and routines (Chatterji et al., 2025; Bresnahan, 2024; Bresnahan et al., 2002; Brynjolfsson et al., 2023).

In this paper, we study this co-invention and diffusion process using granular longitudinal data from organizations on ChatGPT Enterprise, OpenAI's enterprise workspace product for organizational use. These data allow us to observe how adoption unfolds inside organizations over time, rather than only whether a firm has adopted AI at all, and to distinguish experimentation, workflow integration, and broad diffusion within firms. This within-firm perspective stands in contrast to large-sample studies that measure AI adoption at the level of the firm (McElheran et al., 2024; Yotzov et al., 2026), and it connects naturally to the IT implementation and assimilation literature, which emphasizes that formal adoption and routine, organization-wide use are distinct stages of technological change (Cooper and Zmud, 1990; Fichman and Kemerer, 1997, 1999; Jasperson et al., 2005).

We first document several empirical regularities in how ChatGPT Enterprise diffuses within organizations. Over tenure, cumulative adopters, weekly active users, total messages, and message entropy all rise, indicating broader participation across employees. This pattern is consistent with classic evidence that new technologies spread gradually and unevenly within firms (Mansfield, 1963; Fuentelsaz et al., 2003). Task diversity also increases with tenure,

suggesting that firms expand the range of activities for which they use ChatGPT. At the same time, firms use more distinct tools, custom GPTs, and connectors, consistent with greater workflow embedding over time. We also find evidence consistent with peer diffusion within firms: internal sharing of ChatGPT artifacts, such as custom GPTs, is associated with greater persistence of use. Together, these patterns are consistent with a staged diffusion process.

We next examine what accelerates diffusion within firms. Later adoption cohorts move through the diffusion process more quickly than earlier cohorts, reaching similar milestones in fewer weeks since adoption. Two explanations are especially plausible in our setting: later adopters benefit from accumulated implementation knowledge—playbooks, templates, organizational practices, and a maturing ecosystem—and they also adopt against a stronger technology frontier as model capability improves over time. While these cohort patterns are informative, they are not by themselves causally identified. To isolate these mechanisms more directly, we exploit shocks to both organizational capability-building and the technology frontier. Specifically, we use the timing of formal training and enablement events in a matched-pair difference-in-differences design to estimate the effects of structured enablement on adoption, participation, and workflow embedding, and we use major model releases together with temporal causal forest methods (Wager and Athey, 2018; Turjeman and Feinberg, 2024) to estimate the effects of improvements in model capability. Both of these empirical designs identify concrete forces and practical levers that accelerate the transition from initial access to broader and more embedded organizational use.

The paper contributes to several related literatures. First, it provides new measurement on the diffusion of a general-purpose AI technology inside firms using high-frequency organizational data, complementing work on firm-level AI adoption and connecting to classic research on intrafirm diffusion and GPT adoption (McElheran et al., 2024; Yotzov et al., 2026; Mansfield, 1963; Fuentelsaz et al., 2003; Bresnahan and Trajtenberg, 1995; Jovanovic and Rousseau, 2005). Second, it contributes to the literature on human capital, organizational learning, peer effects, and technology assimilation by showing how adoption unfolds

after initial access—through broader participation, expanding task scope, and workflow embedding—and by using formal training events and peer-diffusion patterns to study how know-how spreads within firms and how structured support accelerates adoption (Aghion and Howitt, 2008; Nelson and Phelps, 1966; Foster and Rosenzweig, 1995; Rogers et al., 2014; Conley and Udry, 2010; Attewell, 1992; Fichman and Kemerer, 1997; Cooper and Zmud, 1990; Fichman and Kemerer, 1999; Jasperson et al., 2005). Third, it contributes to the broader GPT literature by showing that diffusion is shaped not only by firm-side learning but also by continuous improvement in the underlying technology: model advances materially accelerate within-firm adoption and embedding (Bresnahan and Trajtenberg, 1995; Jovanovic and Rousseau, 2005). Taken together, these results inform broader debates about spillovers, complements, task reallocation, and the pace of AI diffusion (Acemoglu and Restrepo, 2019).

# 2 Data, measurement, and sample

## 2.1 Enterprise usage data

Our data consist of longitudinal usage records for organizations subscribing to OpenAI's "ChatGPT Enterprise" plan, an admin-managed workspace product that provides organization-wide access to ChatGPT with enterprise controls and built-in capabilities. The sample covers global Enterprise customers from the beginning of 2024 through the last week in the panel, beginning January 26, 2026.

We construct an unbalanced panel at the firm-week level, where a "firm" corresponds to an Enterprise customer as defined in OpenAI's internal sales databases. Firms enter the panel beginning in Q1 2024 and are observed until churn or the end of the sample period. We link firms, where available, to country, two-digit NAICS sector, and employee size. These linked characteristics are not observed for all firms, so any analysis reported by country, industry, or firm size is estimated on the subset of firms for which the relevant characteristic is available. Using usage data generated by these enterprise customers, we construct measures

of within-firm diffusion, use-case breadth, and workflow embedding; Table 1 reports detailed definitions, and we summarize the main objects here.

a) ***Organizational diffusion and intensity*** captures how broadly and how intensively usage spreads within the firm. We track weekly active users, total messages, cumulative adopters, and message entropy over user message shares. Higher message entropy indicates more diffuse participation across users. We analyze these measures by time since adoption, and compare their evolution across industries, firm sizes, and adoption cohorts.

b) ***Use-case breadth and maturation*** captures the range of activities for which a firm uses ChatGPT and how concentrated that usage remains across tasks and artifacts. Because message-level task classification is available only from Q4 2025 onward, analyses of task scope are restricted to the later part of the sample. We track the number of distinct task categories used in a firm-week, the concentration of usage across task categories as measured by the Gini coefficient over task shares, and task-share composition in a recent cross section. Message-level task categories are assigned using an internal hierarchical OpenAI classifier that infers the user's underlying task from each interaction.[1]

c) ***Workflow embedding*** captures the extent to which usage becomes embedded in production workflows. At the firm level, we measure the number of distinct tools, distinct custom GPTs, and distinct connectors used over time, where tools are built-in ChatGPT capabilities, custom GPTs are specialized GPT artifacts created for recurring use, and connectors are integrations that link ChatGPT to external systems or data sources. For each of these three objects, we also construct entropy measures over usage shares across tools, custom GPTs, and connectors. The count measures capture the breadth of workflow components used within the firm, while the entropy measures capture whether usage is concentrated in a small number of artifacts or spread more

[1]These classifications are generated at run time. Consistent with OpenAI's zero data retention policy for Enterprise customers, the underlying conversation content is not stored for subsequent analysis.

broadly across them. Together, these measures are intended to capture movement from stand-alone prompting toward more diversified and embedded use.

We also construct a measure of internal sharing exposure, based on sharing of custom GPTs, ChatGPT Projects, and internal conversation links within the firm, and use this measure to study whether collaborative circulation of ChatGPT artifacts is associated with more persistent use.

In later sections, we study two empirical settings that help identify forces accelerating diffusion. First, we use the timing of trainings, workshops, and related enablement interventions in matched-pair difference-in-differences analyses to estimate whether structured support increases adoption, participation, and usage intensity while reducing concentration. Second, we use major model releases as technology-supply shocks and study event windows around those transitions to estimate whether improvements in model capability accelerate within-firm diffusion.

# 3 Descriptive diffusion patterns

We begin by documenting descriptive patterns in how ChatGPT Enterprise diffuses inside firms. We first show how adoption, active use, message volume, and within-firm concentration evolve over tenure and across industries, cohorts, and firm-size groups. We then examine how firms broaden the set of tasks for which they use ChatGPT and how usage becomes embedded in workflows through increasing use of tools, custom GPTs, and connectors. Finally, we provide descriptive evidence on within-firm sharing and persistence of use.

## 3.1 Baseline Within-Firm Diffusion by Tenure, Cohort, Industry, and Size

We find that cmulative adopters, weekly active users, and total messages all rise with weeks since adoption, indicating that both the stock of users and the intensity of use expand over

tenure. At the same time, message entropy rises steadily, indicating that usage becomes less concentrated across employees and more broadly distributed within the firm. Taken together, these patterns suggest that diffusion inside firms is gradual: use begins relatively concentrated and then broadens over time rather than arriving immediately at full organizational scale (Mansfield, 1963; Fuentelsaz et al., 2003).

We also find substantial heterogeneity by adoption cohort. Later cohorts generally move through these descriptive margins faster than earlier cohorts, reaching comparable levels of adoption, active use, message volume, and within-firm breadth in fewer weeks since adoption. This cohort acceleration is consistent with the idea that later adopters face a more mature implementation environment, though the descriptive evidence alone does not distinguish accumulated organizational know-how from contemporaneous improvements in the technology itself or other potential mechanisms (Battisti and Stoneman, 2003; Battisti et al., 2007; Bresnahan and Trajtenberg, 1995).

There is also substantial heterogeneity with respect to firm size. Larger size buckets exhibit higher levels of cumulative adopters, active users, message volume, and message entropy, likely in part due their larger potential user populations. At the same time, the overall shape of the tenure profile is similar across size groups: in each case, firms expand participation and reduce concentration over time.

## 3.2 Use-Case Breadth and Workflow Embedding

We next examine how firms broaden the range of activities for which they use ChatGPT and how that use becomes embedded in production workflows. We observe a concentrated core of recurring knowledge-work and technical tasks, with a long tail of smaller categories. The most common categories include documentation and technical writing, private personal messages, technical and digital work, sales and marketing, topic overviews, and nontechnical professional or academic assistance. This pattern suggests that enterprise use is not concentrated in a single task domain, but is instead distributed across a diverse set of business-relevant activities.

Over tenure, the concentration of use across tasks changes systematically. We observe that the task-share Gini increases with weeks since adoption, indicating that firms do not simply scale one initial use case. Instead, the distribution of usage across task categories shifts over time as organizations expand into a broader set of activities and refine where the technology is most useful. Importantly, the task-share Gini appears to level off on average, perhaps reflecting convergence on a core set of tasks for which ChatGPT is used within the firm. In this sense, the early post-adoption period appears to involve experimentation and search, while later weeks reflect a more mature task portfolio (Bresnahan and Trajtenberg, 1995; Helpman, 1998; David, 1990).

We also extend this picture from task mix to workflow embedding. We find that the number of distinct tools used within a firm rises early in tenure and then plateaus across industry, cohort, and firm-size groupings. We also find a corresponding rise and eventual flattening in entropy across tool usage shares. Within industries and firm-size groups, this pattern suggests that firms initially experiment with a broader set of tools and then converge on a relatively stable set of workflow components and use cases. Across adoption cohorts, however, the concentration patterns are less stable, consistent with the idea that later cohorts encounter a different tool environment as new capabilities are introduced over calendar time. Taken together, these patterns suggest movement away from stand-alone prompting and toward richer combinations of model interaction and supporting workflow components, consistent with the view that complementary investments are central to realizing the value of a general-purpose technology (Bresnahan and Greenstein, 1996; Brynjolfsson et al., 2021).

The broad pattern for custom GPTs is similar, though cleaner across cohorts. We find that the number of distinct custom GPTs used rises with tenure and then levels off, and that usage across custom GPTs also becomes more diffuse before stabilizing. Unlike the tool patterns, these trajectories do not exhibit the same cohort-specific lumpiness, suggesting a more uniform process of exploration followed by convergence over tenure. This pattern is consistent with firms increasingly building, adopting, and distributing specialized GPT

artifacts, but ultimately concentrating on a smaller set of reusable GPTs that support recurring organizational workflows.

Connectors exhibit a similar broad pattern, though their dynamics are shaped more clearly by calendar time. Because connectors were introduced later in the sample period, the early portions of many tenure profiles are mechanically flat, especially for earlier adoption cohorts. After launch, however, connector use rises quickly, particularly for more recent cohorts, and then begins to level off with tenure. We also observe a corresponding increase in connector-share entropy. Taken together, these results suggest that firms increasingly connect ChatGPT to external systems as adoption matures, but—as with tools and custom GPTs—usage eventually stabilizes around a more persistent set of integration points. More generally, the workflow results indicate that embedding is multidimensional: over tenure, firms use more tools, more custom GPTs, and more connectors, and activity becomes less concentrated within each of those artifact classes.

## 3.3 Within-Firm Sharing and Persistence

We next provide descriptive evidence on within-firm propagation through internal sharing. We compare short-run user retention between firm-weeks with and without internal sharing activity, where sharing includes circulation of custom GPTs, collaborative ChatGPT Projects, and internal ChatGPT conversation links within the firm. Retention is higher in sharing-active weeks across all size buckets, with estimated lifts in the single digit percentage points. Although this evidence is descriptive rather than causal, it is consistent with the view that the internal circulation of reusable AI artifacts reinforces continued use within organizations.

## 3.4 Summary

Taken together, the descriptive evidence in this section supports a staged view of diffusion inside firms. Adoption broadens over tenure as cumulative adopters, active users, message volume, and user-share entropy all rise, consistent with classic evidence that new technologies

spread gradually and unevenly within organizations rather than arriving immediately at full scale (Mansfield, 1963; Fuentelsaz et al., 2003). At the same time, firms expand the range of tasks for which they use ChatGPT and increasingly embed it in production workflows through broader use of tools, custom GPTs, and connectors. These workflow measures generally exhibit early expansion followed by stabilization, suggesting that firms move from experimentation and search toward convergence on a more persistent set of tasks and workflow components (Bresnahan and Trajtenberg, 1995; Helpman, 1998; David, 1990; Bresnahan and Greenstein, 1996; Brynjolfsson et al., 2021). Finally, weeks with internal sharing activity exhibit higher subsequent retention, consistent with the idea that circulation of reusable artifacts helps sustain use within organizations (Foster and Rosenzweig, 1995; Conley and Udry, 2010; Rogers et al., 2014). Overall, the descriptive patterns suggest that the key challenge in organizational adoption is not simply access to the technology, but the slower process of discovering use cases, building complements, and embedding the technology into recurring work routines.

# 4 Identifying accelerants of diffusion

Section 3 documented descriptive patterns in how ChatGPT diffuses inside firms. We now turn to two empirical settings that help identify forces accelerating that diffusion. The first is organizational enablement: trainings, workshops, and related interventions that plausibly reduce implementation frictions and help firms translate established know-how into broader organizational use. The second is the release of new model generations, which shifts the technology frontier available to all firms.

These two settings require different identification strategies. In the case of enablement, we observe firm-level interventions that can be analyzed using matched pairs of treated and control firms around the timing of the intervention. In the case of model releases, the relevant shock occurs in calendar time and affects the frontier available to all firms rather than a treated subset. Following the logic of Temporal Causal Inference (Turjeman and Feinberg,

2024), we therefore shift the temporal reference frame from calendar time to firm tenure: the relevant comparison is between firms at similar points in their post-adoption trajectories that face different model environments because they adopt in different calendar periods. Using this logic, we are able to estimate the effect of major model transitions on within-firm diffusion.

## 4.1 Enablement as an Organizational Shock

Formal enablement provides a natural setting for identifying whether structured organizational support accelerates diffusion. Trainings, workshops, and related onboarding interventions are intended to reduce knowledge barriers, surface promising use cases, and help employees incorporate the technology into recurring work. If organizational frictions are an important barrier to diffusion, then firms that receive enablement should subsequently exhibit faster growth in adoption, broader participation, and more intensive use.

To estimate these effects, we implement a matched-pair difference-in-differences design at the firm level. We begin with treated firms that receive an enablement intervention and their matched control firms from the enablement matching procedure. For each treated-control pair, we define an event window spanning eight weeks before and eight weeks after the treated firm's enablement date, and compute firm-level averages of each outcome separately over the pre-period and post-period. The estimated treatment effect is the standard difference-in-differences estimand: the change in the treated firm's mean outcome from the pre- to post-period minus the corresponding change for its matched control. We estimate this specification for four outcomes—cumulative adopters, weekly active users, weekly message entropy, and weekly total messages—as well as $\log(1 + y)$ transforms of each outcome. Confidence intervals are obtained by bootstrap resampling of matched treated-control pairs, preserving the matched-pair structure in each draw. The matching procedure delivers strong balance on adoption timing, industry, size buckets, seat counts, and employee size in the matched sample.

We find that enablement increases diffusion along multiple margins. Following training, treated firms exhibit higher cumulative adopters, more weekly active users, higher message

entropy, and greater message volume relative to matched controls. The estimated effects are positive in both levels and log specifications, and event-study profiles are relatively flat prior to treatment and shift after week 0 in the expected direction, with effects building over subsequent weeks.

Taken together, these results suggest that structured enablement accelerates diffusion by broadening participation, increasing usage intensity, and making adoption less concentrated within the firm. More broadly, they imply that organizational frictions are an important barrier to diffusion, and that targeted support can help firms translate early experimentation into broader and more sustained organizational use.

## 4.2 Model Releases as Technology-Supply Shocks

Model releases provide a second setting for identifying forces that accelerate diffusion, but the empirical problem is different from the enablement case. Major model releases occur in calendar time and affect the technology frontier available to all firms, so there is no natural contemporaneous untreated group. A simple before-versus-after comparison would therefore confound the effects of improved models with secular changes in adoption, learning, and demand.

To address this problem, we adapt the logic of Temporal Causal Inference (Turjeman and Feinberg, 2024; Huang et al., 2023; Geng et al., 2025). The core idea is to shift the temporal reference frame from calendar time to firm tenure. When a model transition affects all firms in the same calendar period, there is no obvious contemporaneous control group. Instead, we compare firms at similar points in their post-adoption trajectories that face different model environments because they adopted in different calendar periods. In this sense, earlier-adopting firms provide counterfactual diffusion trajectories for later-adopting firms observed at comparable tenure positions but under an earlier technology frontier. Following Turjeman and Feinberg (2024), we first use this temporal alignment logic to construct comparisons across firms at similar tenure positions, and then implement temporal causal forest methods

(Wager and Athey, 2018) to estimate how diffusion outcomes would have evolved absent a given model transition.

We focus on the shift to GPT-5 in August 2025, which provides the main technology-supply shock in our analysis. Our estimates from temporal causal forest analyses suggest that GPT-5 materially accelerates within-firm diffusion: cumulative adopters, weekly active users, and total messages rise relative to counterfactual trajectories, while message entropy also increases, implying broader participation within firms.[2]

These results indicate that the faster diffusion of later cohorts is not solely a matter of accumulated organizational learning. Improvements in the underlying model frontier also have immediate and measurable effects on within-firm adoption, participation, and usage intensity. In this sense, diffusion depends not only on firms learning how to use the technology, but also on continued improvements in what the technology can do.

## 4.3 Summary

Taken together, the evidence in this section identifies two distinct forces that accelerate diffusion inside firms. First, structured enablement increases adoption, broadens participation, raises usage intensity, and makes adoption less concentrated within the firm. Second, improvements in the underlying model frontier—captured here by the shift to GPT-5—also materially accelerate within-firm diffusion, increasing adoption, participation, and usage relative to counterfactual trajectories. These results imply that the faster diffusion observed in later cohorts reflects both organizational capability-building and continued technological progress. More broadly, they suggest that diffusion depends not only on firms learning how to use the technology, but also on ongoing improvements in what the technology can do.

[2]To make outcomes comparable across firms of very different sizes, we normalize weekly active users, cumulative adopters, and weekly total messages by their mean values over the three weeks prior to the GPT-5 model release. We do not apply this normalization to message entropy, since it is already scale-free.

# 5 Conclusions

This paper studies how a general-purpose AI technology diffuses inside organizations using longitudinal usage data from ChatGPT Enterprise. We document a staged diffusion process: over tenure, cumulative adopters, active users, message volume, and within-firm breadth all rise; firms expand the range of tasks for which they use ChatGPT; and usage becomes increasingly embedded in production workflows through broader use of tools, custom GPTs, and connectors. We also find suggestive evidence that internal sharing of reusable ChatGPT artifacts is associated with more persistent use. The paper's main contribution is to make these intermediate stages of diffusion visible at the micro level in a large cross-section of firms, and to complement that descriptive evidence with causal estimates of two concrete accelerants. Structured enablement increases adoption, participation, and usage intensity while reducing concentration, and improvements in the model frontier—captured here by the shift to GPT-5—likewise accelerate within-firm diffusion.

More broadly, the findings suggest that the economic significance of generative AI will depend less on initial access than on the speed with which firms can turn broad technical capability into durable organizational routines. Diffusion inside firms is uneven, cumulative, and complement-intensive: it requires broader participation, task discovery, and workflow embedding, and it is accelerated by both organizational support and advances in the technology frontier. That perspective helps explain why AI diffusion can appear rapid at the frontier yet gradual inside organizations.

The paper has two important limitations. First, our data do not identify the causal effect of ChatGPT adoption on downstream firm outcomes such as productivity, profits, wages, or employment. We can observe how usage spreads and intensifies inside firms, but not whether those changes translate into measurable performance gains. Second, because the sample is composed of ChatGPT Enterprise customers, the results should be interpreted as evidence on diffusion dynamics in this population rather than universal estimates for all firms or all AI tools. In addition, some mechanisms—especially internal sharing and peer propagation—are

supported by descriptive rather than causal evidence.

A natural agenda for future work is to connect these diffusion stages more directly to firm outcomes and to measure organizational mechanisms in greater detail. Richer data on workflows, managerial practices, team structure, and internal networks could clarify how experimentation becomes durable integration, why some firms diffuse more quickly than others, and how organizational change interacts with continued advances in frontier models. More generally, understanding the economic impact of generative AI will require linking the microeconomics of diffusion inside firms to the broader question of when new technical capabilities become routine and productive parts of organizational life.

# 6 Acknowledgments

The authors thank Nick Otis and Cassandra Duchan-Solis. Chatterji, Rakholia, and Weeratunga are full-time employees of OpenAI. Holtz and Tambe are paid part-time visiting researchers at OpenAI.

# References

Acemoglu, D. and Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2):3–30.

Aghion, P. and Howitt, P. W. (2008). *The economics of growth*. MIT press.

Attewell, P. (1992). Technology diffusion and organizational learning: The case of business computing. *Organization Science*, 3(1):1–19.

Battisti, G., Hollenstein, H., Stoneman, P., and Wörter, M. (2007). Inter and intra firm diffusion of ict in the united kingdom (uk) and switzerland (ch): An internationally comparative study based on firm-level data. *Economics of Innovation and New Technology*, 16(7-8):669–687.

Battisti, G. and Stoneman, P. (2003). Inter- and intra-firm effects in the diffusion of new process technology. *Research Policy*, 32(9):1641–1655.

Bresnahan, T. (2024). What innovation paths for AI to become a GPT? *Journal of Economics & Management Strategy*, 33(2):305–316.

Bresnahan, T. and Greenstein, S. (1996). Technical progress and co-invention in computing and in the uses of computers. *Brookings Papers on Economic Activity. Microeconomics*, pages 1–83.

Bresnahan, T. F., Brynjolfsson, E., and Hitt, L. M. (2002). Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence. *Quarterly Journal of Economics*, 117(1):339–376.

Bresnahan, T. F. and Trajtenberg, M. (1995). General purpose technologies "engines of growth"? *Journal of Econometrics*, 65(1):83–108.

Brynjolfsson, E., Li, D., and Raymond, L. R. (2023). Generative AI at work. Technical Report w31161, National Bureau of Economic Research.

Brynjolfsson, E., Rock, D., and Syverson, C. (2021). The productivity j-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics*, 13(1):333–372.

Chatterji, A., Cunningham, T., Deming, D. J., Hitzig, Z., Ong, C., Shan, C. Y., and Wadman, K. (2025). How people use ChatGPT. Technical Report w34255, National Bureau of Economic Research.

Conley, T. G. and Udry, C. R. (2010). Learning about a new technology: Pineapple in ghana. *American economic review*, 100(1):35–69.

Cooper, R. B. and Zmud, R. W. (1990). Information technology implementation research: A technological diffusion approach. *Management Science*, 36(2):123–139.

David, P. A. (1990). The dynamo and the computer: An historical perspective on the modern productivity paradox. *The American Economic Review*, 80(2):355–361.

Eloundou, T., Manning, S., Mishkin, P., and Rock, D. (2024). GPTs are GPTs: Labor market impact potential of LLMs. *Science*, 384(6702):1306–1308.

Fichman, R. G. and Kemerer, C. F. (1997). The assimilation of software process innovations: An organizational learning perspective. *Management Science*, 43(10):1345–1363.

Fichman, R. G. and Kemerer, C. F. (1999). The illusory diffusion of innovation: An examination of assimilation gaps. *Information Systems Research*, 10(3):255–275.

Foster, A. D. and Rosenzweig, M. R. (1995). Learning by doing and learning from others: Human capital and technical change in agriculture. *Journal of political Economy*, 103(6):1176–1209.

Fuentelsaz, L., Gómez, J., and Polo, Y. (2003). Intrafirm diffusion of new technologies: An empirical application. *Research Policy*, 32(4):533–551.

Geng, Y., Mayya, R., and Sun, C. (2025). Token tradability as a platform governance mechanism: Evidence from a policy change. *Available at SSRN 4766534*.

Goldfarb, A., Taska, B., and Teodoridis, F. (2023). Could machine learning be a general purpose technology? a comparison of emerging technologies using data from online job postings. *Research Policy*, 52(1):104653.

Griliches, Z. (1957). Hybrid corn: An exploration in the economics of technological change. *Econometrica*, 25(4):501–522.

Hall, B. H. and Khan, B. (2003). Adoption of new technology. Technical Report w9730, National Bureau of Economic Research.

Helpman, E., editor (1998). *General Purpose Technologies and Economic Growth*. MIT Press.

Huang, H., Fu, R., and Ghose, A. (2023). Generative ai and content creators: Evidence from digital art platforms. *Available at SSRN 4670714*.

Jasperson, J. S., Carter, P. E., and Zmud, R. W. (2005). A comprehensive conceptualization of post-adoptive behaviors associated with information technology enabled work systems. *MIS Quarterly*, 29(3):525–557.

Jovanovic, B. and Rousseau, P. L. (2005). General purpose technologies. In *Handbook of economic growth*, volume 1, pages 1181–1224. Elsevier.

Mansfield, E. (1961). Technical change and the rate of imitation. *Econometrica*, 29(4):741–766.

Mansfield, E. (1963). Intrafirm rates of diffusion of an innovation. *The Review of Economics and Statistics*, 45(4):348–359.

McElheran, K., Li, J. F., Brynjolfsson, E., Kroff, Z., Dinlersoz, E., Foster, L., and Zolas, N. (2024). Ai adoption in america: Who, what, and where. *Journal of Economics & Management Strategy*, 33(2):375–415.

Nelson, R. R. and Phelps, E. S. (1966). Investment in humans, technological diffusion, and economic growth. *The American economic review*, 56(1/2):69–75.

Rogers, E. M., Singhal, A., and Quinlan, M. M. (2014). Diffusion of innovations. In *An integrated approach to communication theory and research*, pages 432–448. Routledge.

Turjeman, D. and Feinberg, F. M. (2024). When the data are out: measuring behavioral changes following a data breach. *Marketing Science*, 43(2):440–461.

Wager, S. and Athey, S. (2018). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242.

Yotzov, I., Barrero, J. M., Bloom, N., Bunn, P., Davis, S. J., Foster, K. M., Jalca, A., Meyer, B. H., Mizen, P., Navarrete, M. A., et al. (2026). Firm data on ai. Technical report, National Bureau of Economic Research.

# TABLE 1: SUMMARY OF ENTERPRISE MEASURES

| Measure | Description |
|---|---|
| **Organizational diffusion and intensity** | |
| Cumulative adopters | Number of employees in firm $f$ who have ever used ChatGPT up to week $t$. |
| Weekly active users (WAU) | Number of unique active users in firm $f$, week $t$. |
| Total messages | Message volume $M_{ft} = \sum_i m_{ift}$ in firm-week $(f, t)$. |
| User-share entropy $\text{Entropy}_{ft}^{\text{user}}$ | Shannon entropy $\text{Entropy}_{ft}^{\text{user}} = -\sum_i s_{ift} \log s_{ift}$, where $s_{ift} = m_{ift}/M_{ft}$; higher values indicate broader participation across users. For firm-weeks with $M_{ft} = 0$, entropy is set to 0 by convention. |
| **Use-case breadth and maturation** (available Q4 2025 onward) | |
| Task-share composition | Distribution of message shares across task categories in firm $f$, week $t$. |
| Distinct task categories | Number of unique classifier-assigned task categories used in firm $f$, week $t$. |
| Task-share concentration (Gini) | Gini coefficient over task-category shares, $G_{ft} = \sum_c \sum_{c'} \lvert s_{cft} - s_{c'ft} \rvert / (2C)$, where $s_{cft}$ is the share of messages in category $c$ and $C$ is the number of task categories; higher values indicate greater concentration in a small number of task categories. |
| **Workflow embedding** | |
| Distinct tools used | Number of unique tools used by firm $f$, week $t$. |
| Tool-share entropy $\text{Entropy}_{ft}^{\text{tool}}$ | Shannon entropy $\text{Entropy}_{ft}^{\text{tool}} = -\sum_j s_{jft}^{\text{tool}} \log s_{jft}^{\text{tool}}$, where $s_{jft}^{\text{tool}}$ is the message share of tool $j$ in firm $f$, week $t$; higher values indicate usage spread across a broader set of tools. For firm-weeks with no tool usage, entropy is set to 0 by convention. |
| Distinct custom GPTs used | Number of unique custom GPTs used by firm $f$, week $t$. |
| Custom-GPT-share entropy $\text{Entropy}_{ft}^{\text{gpt}}$ | Shannon entropy $\text{Entropy}_{ft}^{\text{gpt}} = -\sum_j s_{jft}^{\text{gpt}} \log s_{jft}^{\text{gpt}}$, where $s_{jft}^{\text{gpt}}$ is the message share of custom GPT $j$ in firm $f$, week $t$; higher values indicate usage spread across a broader set of custom GPTs. For firm-weeks with no custom GPT usage, entropy is set to 0 by convention. |
| Distinct connectors used | Number of unique connectors used by firm $f$, week $t$. |
| Connector-share entropy $\text{Entropy}_{ft}^{\text{conn}}$ | Shannon entropy $\text{Entropy}_{ft}^{\text{conn}} = -\sum_j s_{jft}^{\text{conn}} \log s_{jft}^{\text{conn}}$, where $s_{jft}^{\text{conn}}$ is the message share of connector $j$ in firm $f$, week $t$; higher values indicate usage spread across a broader set of connectors. For firm-weeks with no connector usage, entropy is set to 0 by convention. |
| **Within-firm sharing and persistence** | |
| Sharing exposure indicator | Indicator equal to 1 if, in workspace-week $(f, t)$, at least one non-owner recipient is observed on any internal sharing surface (shared GPTs, collaborative Projects, or internal share links), and 0 otherwise. |
| One-week user retention | Retention rate from week $t$ to $t+1$, $R_{ft \to t+1} = \#\{\text{users active in both } t \text{ and } t+1\} / \#\{\text{users active in } t\}$. |