COMMENT[*]

on

"Mind the Gap: AI Adoption in Europe and the US"

by Bick, Blandin, Deming, Fuchs-Schündeln, and Jessen

Raffaella Sadun

Harvard Business School, CEPR and NBER

May 19, 2026

Bick, Blandin, Deming, Fuchs-Schündeln, and Jessen provide some of the clearest and most internationally comparable evidence to date on generative AI adoption across the developed world. The paper combines two rounds of harmonized worker surveys across seven countries with firm-level data from the EU-ICT-Firm survey and the US Census BTOS. I emphasize three findings emerging from the paper. First, the paper documents a substantial US-Europe gap: 43 percent of US workers report using generative AI for their job, compared with 32 percent of European workers on average. Second, an Oaxaca-Blinder decomposition attributes roughly 55 percent of this gap to compositional differences in worker demographics, occupation, industry, and firm size. Third, and in my view most interestingly, an augmented decomposition that adds firm encouragement to the same controls accounts for nearly all of the US-Europe adoption gap.

This last finding is the reason I read the paper as a contribution to the measurement of generative AI diffusion, as well as to the study of organizational adjustment to a major technological change. The discussion below asks three questions: how confident we should be about the magnitude of the gap; what mechanisms could make firm encouragement so important, especially relative to earlier waves of Information and Communication Technologies (ICT); and what the evidence implies for research and policy.

[*] The final version of this comment will be published in the Spring 2026 issue of the *Brookings Papers on Economic Activity (BPEA)*.

I. Is the Gap Real?

I find the evidence persuasive that a US lead in generative AI adoption exists, although the precise magnitude remains open. Bick and others (2026) document a worker-level gap relative to the US of 7 to 17 percentage points across six European countries. The same broad pattern appears in Yotzov and others (2026), who survey senior executives through central bank panels, and in Hartley and others (2025), who use an independent worker panel covering the US, Canada, and several European countries. The fact that these estimates line up despite different respondents and survey instruments makes it hard to dismiss the gap as a survey artifact.

At the same time, the magnitude of the gap is highly uncertain. Bick and others provide a clear example of the measurement challenges: the share of EU firms using generative AI for "any business purpose" is roughly five times larger than the share using generative AI for the "production of goods and services." This matters because, before its December 2025 revision (U.S. Census Bureau 2025), the US Census BTOS used the narrower "producing goods or services" wording, while the EU-ICT-Firm survey asks about a list of specific generative AI technologies for any business function. The concern is that we may be comparing different *margins* of adoption, not only different *levels* of adoption. This should not be treated simply as measurement error. It reflects a deeper ambiguity about what constitutes economically meaningful use. Workers, ICT managers, and C-suite executives observe different parts of the process, and firm-weighted and employment-weighted sampling answer different questions. The same ambiguity plagued ICT measurement two decades ago (Syverson 2011).

The paper also presents evidence that the gap between the US and Europe may be widening. Between the 2025 and 2026 survey waves, US worker adoption increased by 3.6 percentage points, while adoption in the three lowest-adoption European countries (Germany, France, Italy) increased by only 0.1 to 1.1 percentage points. This divergence, and its possible implications for productivity, make measurement precision all the more urgent.

The measurement problem is not only how much AI is being used, but what kind of use is being measured and who is in a position to report it. This implies the need to define more precisely what constitutes generative AI adoption from a productivity standpoint, distinguishing individual or peripheral uses (drafting emails, searching for information) from deep integration into core

workflows. Furthermore, it is essential to understand who in the firm is best placed to report on adoption. Workers, ICT managers, and senior executives each observe different facets of the same phenomenon, and each of these actors introduces different perspectives (or, possibly, biases).

II. What Is Driving the Gap?

The correlates of the generative AI adoption gap are very similar to those documented for ICT in the 1990s and 2000s in terms of firm-size effects and management correlates (Bloom, Sadun, and Van Reenen 2012; Bresnahan, Brynjolfsson, and Hitt 2002). Within Europe, the cross-country ordering also looks familiar: the range of adoption rates documented in this paper echoes the ICT geography of two decades ago, with Nordic and Anglo-Saxon countries leading and Southern Europe lagging (Schivardi and Schmitz 2020). This points to a known mechanism from the ICT era: differences across countries and firms in the organizational capacity to absorb new technologies.

Should we expect the adoption of generative AI to depend on organizational factors in the same way it did in the past? I would expect generative AI adoption to be at least as dependent on organizational choices as ICT adoption was, and probably more so, for two reasons.

The first is *uncertainty* about relevant use cases. ICT was also uncertain in its early years: the productivity paradox of the late 1980s reflected precisely that firms did not know where computing would be useful. However, many ICT applications, once specified, produced outputs that could be checked against relatively clear rules: a payroll system either produced the right paychecks or did not. Generative AI introduces different kinds of uncertainty, rooted in its probabilistic foundations. First, the potential benefit of generative AI use is often discovered rather than pre-specified: a firm is unlikely to know in advance what a model will do well in its specific workflow and must discover this through using it. Second, the same input does not necessarily yield the same output. Even when outputs can be made more stable through deployment choices, small changes in prompt, context, or task framing can shift answers in ways that are difficult to evaluate without domain expertise. Lastly, evaluating whether an output is correct often requires the same expertise the technology is meant to augment. Autor and Thompson (2025) emphasize that generative AI is distinctive because it can perform work that previously required years of specialized training,

reshaping the value of expertise. The corollary of this capability, however, is that a junior worker may be in the position to use generative AI to produce a draft that looks like expert work without being able to judge whether it actually is. As a case in point, Dell'Acqua and others (2026) show that even within a single consulting firm, AI dramatically improves performance on some tasks while degrading it on similar-looking adjacent tasks (the "jagged technological frontier"), and workers cannot reliably tell the difference. The combination of the need for discovery, stochastic output, and costly verification creates a premium on rigorous experimentation: the capacity to try, potentially fail, learn, and iterate around concrete and firm-specific workflows. Firms with stronger approaches to experimentation should face higher expected returns to generative AI adoption, and cross-firm differences in these intangible capabilities are likely to appear as cross-firm differences in usage.

The second difference is disruption of expertise and authority. As the example of the junior worker suggests, generative AI does not simply automate routine tasks, as in the canonical framework of Autor, Levy, and Murnane (2003). It can also change who is able to perform work that previously required years of specialized training (Brynjolfsson, Li, and Raymond 2025), and it can allow experts to span knowledge silos (Dell'Acqua and others 2025). As the technology shifts production toward less experienced workers, the value of expertise also changes (Autor and Thompson 2025). Ide and Talamàs (2025) formalize this point in a model of knowledge hierarchies in which humans sort into routine "workers" and exception-handling "solvers." Their analysis implies that generative AI's effects (in particular, the effects of autonomous AI agents) need not be monotone across the skill distribution. Top experts may benefit because generative AI amplifies their reach and leaves them the hardest problems. Individuals whose tasks fall within the model's capability are more exposed to substitution. The sharpest adverse impact may fall on middle-tier experts and mid-career knowledge workers, whose tasks overlap closely with what generative AI can do and whose judgment does not benefit from the scaling that lifts the top.

This heterogeneity compounds the verification problem described above. In many settings, the natural verifiers of generative AI output are not top experts, whose time is scarce, but workers whose knowledge lies just above the model's capability. They can often assess whether the answer is right, but they may also be among the workers most exposed to substitution. Adoption therefore requires firms to elicit cooperation from people who may have private reasons to slow it down.

This potential source of resistance was less central in many ICT settings, where the workers displaced by automation were not necessarily the same workers needed to make the technology work safely.

These features make generative AI adoption more than a technology-deployment problem. Firms have to decide who may use the tool, who checks the output, and how authority changes when the tool works. The quasi-fixed nature of labor identified by Oi (1962) means that this adjustment is unlikely to be immediate or frictionless. The organizational barriers to generative AI adoption are therefore likely to be larger, and in some respects different, from those that slowed ICT diffusion.

Interpreting the Firm Encouragement Result

I read the paper's most interesting descriptive finding through this lens of organizational change: firm encouragement, including tool provision and active promotion of the technology, helps account for nearly all the average US-Europe gap in the augmented decomposition. Encouragement is not the underlying explanation, but the object to be explained. Why do European firms encourage generative AI use less than US firms?

Lower encouragement may reflect lower private, and possibly lower social, returns to AI adoption within the specific context of European institutional environments. Stronger employment protection makes reorganization costly; works councils require negotiation over workflow changes; different skill distributions alter the cost-benefit calculus of technological adoption (Draghi 2024). Under this interpretation, lower encouragement is a rational response to binding constraints rather than a failure, and policy aimed simply at equalizing adoption rates would be poorly targeted.

A second, not mutually exclusive, interpretation is that the gap reflects organizational frictions and coordination failures that create a wedge between private and social returns to experimentation and adaptation. Firms may lack information about which generative AI applications are productive, and early adopters may generate knowledge spillovers they cannot capture. Additionally, complementarities between technology, organization, and human capital (Milgrom and Roberts 1990) may create coordination problems: partial adoption yields low returns, discouraging further investment. Organizational inertia, including sunk costs in existing workflows and resistance to change from workers whose expertise is threatened, may prevent

adoption even when returns are high (David 1990). My reading is that the evidence points at least partly toward this second interpretation, although the current data cannot distinguish cleanly between institutional constraints and organizational frictions.

Understanding the relative weight of these interpretations is essential for policy because the same adoption gap can have different welfare implications. One way to make progress is to compare firms operating under the same national institutions but differing in organizational imprinting. The paper's data make this difficult to do directly, but firm demographics and ownership structures offer a promising path forward. If the gap in adoption between the US and Europe narrows substantially among younger firms (with presumably lower legacy constraints), or for specific ownership types (e.g., family-owned firms versus others), this may suggest a larger role for organizational factors within institutional constraints. Multinational ownership offers another test: Bloom, Sadun, and Van Reenen (2012) showed that US multinationals operating in Europe used IT more intensively than European firms in the same markets and had higher productivity returns from these investments, consistent with the constraint being organizational rather than purely institutional. If the same pattern holds for generative AI, especially within narrowly defined industries and use cases, it would be suggestive evidence against a purely efficient-adaptation interpretation.

Another avenue for future research is to collect direct firm-level measures of generative AI-specific organizational practices. A natural candidate would be to measure whether and how generative AI experimentation cycles are structured inside the firm. Management surveys such as the World Management Survey and the Census MOPS could be adapted to capture experimentation practices, decentralization of AI-related decisions, and the extent of workflow redesign. Capturing such measures in future survey waves would sharpen the interpretation of how the encouragement result plays out inside firms.

III. Broader Implications

The macroeconomic implication is that generative AI adoption may be slow even when the underlying technology is powerful. If adoption depends on organizational redesign rather than technology investment alone, then standard diffusion models may fail to predict the trajectory of generative AI adoption and its productivity impact. Brynjolfsson, Rock, and Syverson (2021)

describe a "productivity J-curve" in which general-purpose technologies initially reduce measured productivity as firms invest in intangible organizational complements. For generative AI, the intangible investment phase may be dominated by experimentation and reorganization rather than capital expenditure. The trough in measured productivity may therefore be steeper and longer than models calibrated on the ICT J-curve would suggest. McElheran and others (2025), for example, provide evidence of J-curve-shaped returns to industrial AI in US manufacturing, consistent with the view that organizational adjustment is a critical complement to AI capital. One would expect these dynamics to be even more salient for generative AI.

The nature of the frictions matters for the shape of diffusion. For example, if barriers are fixed costs and disruption costs (Caroli and Van Reenen 2001), such as resistance from incumbent workers or the need for simultaneous changes across multiple organizational dimensions, the dynamics differ fundamentally from smooth convex adjustment. Fixed costs and complementarities can generate threshold effects: firms remain in low-adoption states until conditions tip, then transition rapidly. This implies longer delays, and larger cross-country divergence than standard models predict. The widening gap documented in the paper's comparison of 2025 and 2026 adoption rates is consistent with this possibility.

The paper's new evidence on productivity and employment consequences is important. The positive association between generative AI adoption and industry-level productivity growth in both Europe and the US, combined with the null on employment, is consistent with early-stage dynamics in which generative AI generates efficiency gains without yet triggering large-scale labor reallocation. These are correlations, not causal estimates, and the authors are appropriately cautious. But they provide useful context for the modeling agenda.

The next empirical step is to get inside the firm. We need to know why encouragement varies, and which forms of encouragement matter: access alone, training, process redesign, or permission to experiment. Field experiments that vary the form of support within firms (for example, tool provision alone versus tool provision plus training) would help identify which frictions bind most. Cross-country firm-level data should also link organizational practices, generative AI adoption, firm characteristics, and productivity over time. The US Census BTOS provides a useful benchmark that other countries should aim to match, although measurement should go beyond adoption rates and time savings to capture the use of generative AI for new tasks and activities.

Finally, macro models should incorporate organizational frictions, including experimentation costs, coordination failures, and resistance to change, if they are to generate realistic predictions of diffusion speed and cross-country divergence.

The paper also has implications for generative AI policy. Much of the current policy debate is organized around two margins: risk and regulatory burden. The EU AI Act is centered on a risk-based regulatory framework for AI developers and deployers and is complemented by initiatives intended to support trustworthy AI, investment, and innovation (European Commission 2025). Recent US policy, including the March 2026 National Policy Framework for AI, emphasizes innovation, deployment, regulatory sandboxes, and the reduction of regulatory fragmentation (White House 2026). These frameworks address important margins, but the evidence in Bick and others points to a margin that sits inside the firm. The constraint may be whether managers know where to try generative AI, whether workers have incentives to participate, and whether the organization can change the workflow when the experiment succeeds.

To be clear, the potential role of organizational frictions in generative AI adoption does not necessarily warrant policy intervention. If lower encouragement in Europe reflects lower private and social returns to generative AI adoption under European institutional arrangements, policies aimed simply at raising adoption rates would be poorly targeted. If, instead, lower encouragement reflects information spillovers, coordination failures among complementary investments, fixed implementation costs, financing constraints, or bargaining frictions that prevent mutually beneficial reorganization, then the case for intervention is stronger. The policy objective should not be to close the adoption gap per se, but to reduce the inefficient component of the gap.

Even in this case, however, it would be critical to leverage policy instruments that are well suited to deal with the specific friction. For example, sectoral testbeds and demonstration projects, analogous to agricultural extension services in spirit (Griliches 1957), can lower the cost of discovery when firms do not yet know which use cases are productive. Cross-firm learning networks are relevant when early adopters generate knowledge that others can use but that the original firm cannot fully appropriate; evidence from peer effects, management experiments, and historical technology transfers suggests that such channels can matter (Conley and Udry 2010; Bloom and others 2013; Giorcelli 2019). Implementation vouchers or matching grants may be

justified for smaller firms when the binding constraint is the fixed cost of training, governance, and workflow redesign. And where generative AI changes the allocation of expertise or authority inside the firm, co-designed deployment processes may be needed; if worker cooperation is necessary for verification and implementation, they may be a productive complement to adoption (Atkin and others 2017). More broadly, if the main frictions are concentrated in incumbents, policy should also consider entry barriers and the reallocation of resources toward firms better able to use the technology productively (Bartelsman, Haltiwanger, and Scarpetta 2013).

IV. Conclusion

Bick, Blandin, Deming, Fuchs-Schündeln, and Jessen highlight substantial cross-country patterns of generative AI adoption across the developed world. The US-Europe gap is real, probably substantial in magnitude, and closely associated with firm-level decisions about whether and how to encourage generative AI use. The central unresolved question is how to interpret that firm-level margin. The right policy depends on whether lower European encouragement reflects rational responses to binding institutional constraints or wedges between private and social returns driven by information frictions, coordination failures, or organizational inertia. In the latter case, policies that reduce adjustment costs or improve the institutional conditions for experimentation may be warranted. Progress will require both better data and better theory: data that link precise measures of generative AI use with rich firm characteristics, organizational practices, and productivity outcomes over time; and theory that incorporates experimentation, workflow redesign, and organizational resistance into models of technology diffusion.

The US-Europe productivity divergence of the past thirty years was shaped in significant part by differential technology adoption and by the organizational capacity to exploit new tools. The patterns emerging around generative AI suggest that similar forces may again be at work.

References

Atkin, David, Azam Chaudhry, Shamyla Chaudry, Amit K. Khandelwal, and Eric Verhoogen. 2017. "Organizational Barriers to Technology Adoption: Evidence from Soccer-Ball Producers in Pakistan." *Quarterly Journal of Economics* 132(3): 1101–1164.

Autor, David H., Frank Levy, and Richard J. Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." Quarterly Journal of Economics 118(4): 1279-1333.

Autor, David, and Neil Thompson. 2025. "Expertise." Journal of the European Economic Association 23(4): 1203–1271.

Bartelsman, Eric, John Haltiwanger, and Stefano Scarpetta. 2013. "Cross-Country Differences in Productivity: The Role of Allocation and Selection." American Economic Review 103(1): 305–334.

Bick, Alexander, Adam Blandin, David J. Deming, Nicola Fuchs-Schündeln, and Jonas Jessen. 2026. "Mind the Gap: AI Adoption in Europe and the US." BPEA Conference Draft, Spring.

Bloom, Nicholas, Raffaella Sadun, and John Van Reenen. 2012. "Americans Do IT Better: US Multinationals and the Productivity Miracle." American Economic Review 102(1): 167-201.

Bloom, Nicholas, Benn Eifert, Aprajit Mahajan, David McKenzie, and John Roberts. 2013. "Does Management Matter? Evidence from India." Quarterly Journal of Economics 128(1): 1-51.

Bresnahan, Timothy F., Erik Brynjolfsson, and Lorin M. Hitt. 2002. "Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence." Quarterly Journal of Economics 117(1): 339-376.

Brynjolfsson, Erik, Danielle Li, and Lindsey R. Raymond. 2025. "Generative AI at Work." Quarterly Journal of Economics 140(2): 889–942.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. 2021. "The Productivity J-Curve: How Intangibles Complement General Purpose Technologies." American Economic Journal: Macroeconomics 13(1): 333-372.

Caroli, Eve, and John Van Reenen. 2001. "Skill-Biased Organizational Change? Evidence from a Panel of British and French Establishments." Quarterly Journal of Economics 116(4): 1449-1492.

Conley, Timothy G., and Christopher R. Udry. 2010. "Learning About a New Technology: Pineapple in Ghana." American Economic Review 100(1): 35–69.

David, Paul A. 1990. "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox." American Economic Review 80(2): 355–361.

Dell'Acqua, Fabrizio, Edward McFowland III, Ethan R. Mollick, Hila Lifshitz-Assaf, Katherine C. Kellogg, Saran Rajendran, Lisa Krayer, François Candelon, and Karim R. Lakhani. 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): 403–423.

Dell'Acqua, Fabrizio, Charles Ayoubi, Hila Lifshitz, Raffaella Sadun, Ethan Mollick, Lilach Mollick, Yi Han, Jeff Goldman, Hari Nair, Stew Taub, and Karim R. Lakhani. 2025. "The Cybernetic Teammate: A Field Experiment on Generative AI Reshaping Teamwork and Expertise." Harvard Business School Working Paper, No. 25-043.

Draghi, Mario. 2024. "The Future of European Competitiveness." Report to the European Commission, September.

European Commission. 2025. "AI Continent Action Plan." COM(2025) 165 final, Brussels.

Giorcelli, Michela. 2019. "The Long-Term Effects of Management and Technology Transfers." American Economic Review 109(1): 121-152.

Griliches, Zvi. 1957. "Hybrid Corn: An Exploration in the Economics of Technological Change." Econometrica 25(4): 501-522.

Hartley, Jonathan, Filip Jolevski, Vitor Melo, and Brendan Moore. 2025. "The Labor Market Effects of Generative Artificial Intelligence." Working Paper, SSRN 5136877.

Ide, Enrique, and Eduard Talamàs. 2025. "Artificial Intelligence in the Knowledge Economy." Journal of Political Economy 133(12): 3762–3800.

McElheran, Kristina, Mu-Jeung Yang, Zachary Kroff, and Erik Brynjolfsson. 2025. "The Rise of Industrial AI in America: Microfoundations of the Productivity J-Curve(s)." CES Working Paper 25-27, U.S. Census Bureau.

Milgrom, Paul, and John Roberts. 1990. "The Economics of Modern Manufacturing: Technology, Strategy, and Organization." American Economic Review 80(3): 511-528.

Oi, Walter Y. 1962. "Labor as a Quasi-Fixed Factor." Journal of Political Economy 70(6): 538-555.

Schivardi, Fabiano, and Tom Schmitz. 2020. "The IT Revolution and Southern Europe's Two Lost Decades." Journal of the European Economic Association 18(5): 2441–2486.

Syverson, Chad. 2011. "What Determines Productivity?" Journal of Economic Literature 49(2): 326-365.

U.S. Census Bureau. 2025. "BTOS AI Core Question Updates." December 3.

White House. 2026. "National Policy Framework for Artificial Intelligence: Legislative Recommendations." March 20.

Yotzov, Ivan, Jose Maria Barrero, Nicholas Bloom, Philip Bunn, Steven J. Davis, Kevin M. Foster, Aaron Jalca, Brent H. Meyer, Paul Mizen, Michael A. Navarrete, Pawel Smietanka, Gregory Thwaites, and Ben Zhe Wang. 2026. "Firm Data on AI." NBER Working Paper 34836.