A new challenge for cohesion policies: Is AI a new source of regional divergence?

Francesco Pigliaru

Università di Cagliari and CRENoS

European Union cohesion policy has long been guided by the premise that development constraints are territorially specific and that effective interventions must be adapted to local contexts. This place-based orientation has encouraged strategies that combine financial transfers with locally tailored investment priorities, multilevel governance, and a strong role for subnational actors in diagnosing constraints and shaping project portfolios. Yet a widespread concern has emerged that, despite sustained effort and substantial resources, the overall record on durable convergence across European territories has been less transformative than hoped, especially in regions facing persistent development traps.

The post-pandemic policy cycle has intensified the governance debate. Recovery-oriented initiatives have placed stronger emphasis on measurable deliverables, tighter timelines, and performance-oriented monitoring. Even if it is too early to judge long-run effects, these changes have shifted attention from the mere scale of funding to the institutional conditions that make implementation feasible: administrative capacity, project selection and procurement competence, credible monitoring, and the ability to sustain reforms when local incentives and constraints are unfavourable. This governance shift has reopened a core question for cohesion policy: should Europe remain anchored to place-based principles, or should it move toward forms of more centralized coordination and oversight—at least for some policy components—when local constraints persist and undermine effectiveness?

This paper argues that the governance debate cannot be separated from a set of structural forces that are reshaping Europe's growth prospects and the geography of opportunity: the rapid diffusion of generative artificial intelligence. Generative AI is a broad technology that can affect productivity across many activities, from administration and business services to design, coding, customer interaction, data analysis, and parts of research and development. Yet its impacts are unlikely to be spatially neutral. European territories differ markedly in human capital, institutional quality, managerial capability, digital infrastructure, and civic capital. These differences shape both the capacity to adopt AI effectively and the capacity to translate AI adoption into durable gains in productivity and growth. The central question of this paper is therefore: under what conditions does generative AI foster territorial convergence, and under what conditions does it generate new divergence across heterogeneous territories?

To address this question, we adopt a task-based analytical framework associated with recent work by Charles I. Jones and Christopher Tonetti (2026), conceptually related to the "weak links” logic of O-ring production as introduced by Michael Kremer (1993). The framework represents production as a collection of essential tasks that are strongly complementary. Complementarity is captured by a low elasticity of substitution across tasks: tasks cannot easily compensate for one another in the aggregate production process. In plain language, overall output is vulnerable to bottlenecks. Exceptional performance in some tasks cannot fully offset poor performance in other tasks if tasks must be completed together for production to succeed. This is the “weak links” intuition: aggregate performance depends disproportionately on the weakest tasks, because the production chain is only as strong as its most binding bottlenecks.

This framework yields two insights that are central for thinking about convergence and divergence. The first is that, as the set of machine-performed tasks expands, a larger share of production becomes exposed to the pace at which machine performance improves in those tasks. If machine capabilities improve faster than human capabilities in the relevant tasks, shifting tasks toward machines changes the growth composition of the economy, not only its level.

Second, the framework allows for the possibility that some tasks remain structurally hard to automate. Hard-to-automate tasks are defined by comparative advantage: even as AI improves, these tasks remain more cheaply performed by humans because they rely on tacit knowledge, embodied interaction, contextual judgment, trust-based relationships, or deep organizational embedding that is difficult to codify. When a nontrivial set of tasks remains in this category, weak links have a powerful implication: as AI boosts performance in the tasks that can be automated, the residual human tasks increasingly become the binding constraints. Put differently, the better AI becomes in what it can do, the more aggregate performance depends on what remains stubbornly human.

A simple Kremer-inspired toy model to illustrate divergence mechanisms

To illustrate how a task-based framework—where complementarity across tasks is central to production—can be used to assess how AI adoption may shape regional disparities, it is helpful to introduce a minimal “O-ring” toy model in the spirit of Kremer (1993). Consider an economy where production requires completing a fixed number of tasks, denoted by N. A subset of these tasks is performed using AI-enabled capital and automation, and the remaining tasks are performed by humans. Let n_A be the number of tasks run by AI, and let $N - n_A$ be the number of residual tasks run by human labour. Suppose that each AI-performed task is completed with “quality” q_A, while each human-performed task is completed with “quality” q_L. Output is proportional to effective capital services Kmultiplied by the product of task qualities, so that:

$$Y \;=\; K \cdot q_A^{n_A}\, q_L^{N-n_A} \;=\; K \cdot q_L^{N} \left(\frac{q_A}{q_L}\right)^{n_A}.$$

This formulation embodies the core weak-links intuition: tasks are complementary, and aggregate output multiplies the performance of each step. In such an environment, small differences in task quality can translate into large differences in aggregate output, especially when weaknesses occur in tasks that remain essential and cannot be bypassed.

(1) Same AI quality but different human-task quality generates large output gaps
Hold AI performance constant across regions (the same q_A everywhere) but allow residual human performance to differ: q_L is lower in lagging territories because of weaker human capital, poorer organizational practices, or weaker institutions. The toy model implies that the output difference induced by q_L is magnified by the number of human-performed tasks $N - n_A$. Even modest gaps in q_L can generate sizable output gaps because the effect compounds across multiple essential tasks. The key implication is that, when tasks are strongly complementary, territories with low productivity in residual human tasks can remain constrained even if they have access to the same frontier AI tools. Put differently, weak links make "residual capability" a first-order determinant of aggregate performance.

(2) Higher adoption—more AI-run tasks in advanced regions—is an additional source of regional gaps
Now consider differences in adoption: advanced regions automate a larger number of tasks, so they have higher n_A, while lagging regions automate fewer tasks and therefore have more residual human tasks. Using the second expression above, automation affects output through the factor $(q_A/q_L)^{n_A}$. If AI is genuinely more effective than humans in the tasks it takes over (in the sense that $q_A > q_L$ for those tasks), then $\frac{q_A}{q_L} > 1$and output rises with n_A. A region with a higher n_A therefore enjoys an extra multiplicative advantage, even if its human-task quality were identical to that of the lagging region. When combined with the first mechanism (lower q_L in lagging territories), this creates a "double disadvantage": lagging regions both retain more tasks on the human side and perform those residual tasks less effectively.

This provides a simple illustration of a broader claim: automation does not only shift who performs tasks; it changes the economic weight of the technology frontier. Regions that move more tasks onto the AI side become more exposed to the capabilities of AI-enabled production and less constrained by human-side bottlenecks—provided that the tasks moved are precisely those where AI has comparative advantage.

(3) Faster increases in AI adoption are an additional divergence force, even if AI quality is constant over time and identical across regions.
Finally, consider a setting where the per-task AI quality $q_A$is constant over time (for example $q_A = 0.9$ everywhere) and per-task human quality q_L is also stable within each territory. In this case, growth in output from this channel comes from changes in the number of tasks that are automated. The toy model implies that when $q_A > q_L$, a faster increase in n_A mechanically produces faster improvements in aggregate output, because more tasks move from a lower-

quality execution mode to a higher-quality execution mode. A region that automates tasks more rapidly will experience a stronger transitional growth boost than a region where adoption proceeds slowly, even though both regions face the same underlying AI technology. Divergence can therefore arise not only from different end states (different long-run n_A), but also from different speeds of transition (different trajectories of n_A over time).

This "speed" channel is especially relevant for cohesion policy because many territorial disparities emerge and persist through transitional dynamics: slow adoption can mean that lagging territories fall behind during the period when the technology is diffusing.

Taken together, the toy model delivers a clear message: under strong task complementarities, AI can generate divergence through three reinforcing mechanisms—lower productivity in residual human tasks, lower overall adoption (fewer AI-run tasks), and slower adoption (a flatter adoption trajectory). The toy model is deliberately stylized, but it makes transparent the logic behind the "weak links + adoption heterogeneity" narrative: the same technological opportunity can yield very different aggregate outcomes across territories when local capabilities determine both the productivity of residual tasks and the extent and speed of automation.

From the toy model to a general, micro-founded task framework

While the toy model is useful for intuition, it treats the division of tasks into AI-run and human-run tasks as a reduced-form choice and does not explicitly model the cost comparison that governs automation, nor the equilibrium interactions between wages, the cost of capital/AI services, and the allocation of inputs across tasks. These issues matter when moving from conceptual mechanisms to empirically interpretable predictions and policy design. The remainder of the paper therefore explores the same divergence channels using the more general and micro-founded task-based framework of Jones and Tonetti. In that framework, tasks are aggregated with complementarity, while within-task substitution between humans and AI-enabled capital is driven by relative cost-effectiveness. This allows us to formalize, in a disciplined way, how territorial heterogeneity in human capital, institutional quality, managerial capability, and adoption complements maps into different automation sets and different bottleneck structures—thereby delivering clearer comparative predictions about when AI fosters convergence and when it risks amplifying territorial divergence.

Policy implications for cohesion: place-based capability building and scalable coordination

The analysis suggests that the policy-design dilemma between place-based strategies and more centralized governance should be reframed as a problem of targeting bottlenecks and adoption complements. Place-based policy remains essential because bottlenecks and adoption frictions are territorially specific. Building human capital, upgrading managerial practices, and strengthening the organizational foundations of effective public administration require local diagnosis and sustained, context-sensitive capability building. Yet the same framework highlights why selective centralization can be efficient. Some complements to AI adoption are scalable and benefit from uniform standards: interoperable data architectures, cybersecurity protocols, procurement frameworks, evaluation methods, and platforms that reduce fixed costs for smaller and weaker administrations. Central coordination can also help overcome local political-economy traps that block reform or undermine implementation discipline. The relevant policy question is therefore not "centralization versus place-based" in the abstract, but rather which policy components should be centralized because they scale and reduce adoption costs, and which must remain place-based because they require contextual information and local ownership.

Overall, the paper reframes cohesion policy for an AI era in which the binding constraint may shift from access to technology to the ability to adopt it and to raise performance in the tasks that remain structurally human. In a weak-links world, those residual tasks increasingly set the ceiling for territorial economic performance. A cohesion strategy that takes this seriously must treat AI not only as a digital investment item but as a capability challenge: enabling lagging territories to automate where it is economically viable, and to raise productivity in the residual tasks that remain human. The contribution of the paper is to provide an analytical framework that makes these mechanisms explicit.