Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

AI Research Attention and Higher Education System Response: A Lead-Lag Panel Analysis of the EU-27

Gabriel Osei Forkuo

gabriel.forkuo@unitbv.ro

Transylvania University of Brașov

Emmanuel Osei-Dwomoh

Controller and Accountant General's Department

Research Article

Keywords: generative artificial intelligence, panel data, Granger causality, European Union, research investment, digital skills

Posted Date: September 8th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10951085/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Abstract

Background

Studies of generative artificial intelligence (GenAI) in higher education usually ask whether digital and policy conditions predict AI-related activity in the same year. Few studies ask whether AI-related activity actually precedes, or instead follows, broader system-level change. This study tests whether AI research and policy attention in one year predicts higher education system decision-making outcomes in the following year, and whether the reverse direction also holds, using an EU-27 country-year panel.

Method

We built one-year lags of three AI-attention measures, GenAI and higher education publications, a national AI policy index, and Google Trends search interest, together with three system outcomes: gross tertiary enrollment, research and development expenditure as a share of GDP, and the share of the population with above-basic digital skills. We estimated four two-way fixed-effects panel models with cluster-robust standard errors: three forward models testing whether lagged AI attention predicts each outcome, controlling for the outcome's own lag, and one reverse model testing whether lagged system outcomes predict future GenAI publication activity.

Results

All four models show strong persistence, with each outcome's own lag as the strongest predictor of its current value. Among the AI-attention measures, only lagged search interest significantly predicts next-year research and development expenditure (coefficient 0.0042, p = 0.003). Lagged GenAI publications show a marginal negative association with next-year tertiary enrollment (coefficient − 0.746, p = 0.092) that does not reach conventional significance. No AI-attention measure significantly predicts next-year digital skills. In the reverse direction, none of the three lagged system outcomes significantly predicts future GenAI publication activity.

Conclusion

Higher education systems show limited evidence of AI attention driving broader institutional change within a one-year horizon, with one specific exception: public search interest in GenAI appears to precede small increases in research investment. Persistence, not AI-driven disruption, describes most of the year-to-year variation in these system indicators.

1. INTRODUCTION

1.1 Background

The rapid emergence of generative artificial intelligence has prompted unprecedented academic and public interest across global higher education systems (Krishnan, 2025; MacArthur, 2025). Between 2020 and 2025, universities, researchers, and national ministries across the 27 European Union member states produced a surging volume of publications exploring generative AI applications (Frumin et al., 2026). Much of this initial scholarly engagement focused on classroom-level opportunities, assessment restructuring, and academic integrity challenges (Cotton et al., 2024; Martin et al., 2025; Zhang, 2025). Simultaneously, public search interest in generative AI expanded rapidly, reflecting broad societal curiosity and experimentation (Choi and Varian, 2012; Jun et al., 2018).

Despite this intense attention, national-level higher education responses have remained uneven and poorly understood across Europe (Shapiro and Lam, 2025; Yueh, 2025). Educational leaders face mounting pressure to modernize institutional infrastructure, update curricula, and allocate financial resources toward artificial intelligence capabilities (Krause et al., 2025; Tong et al., 2025). However, significant debate persists regarding whether surging research interest actively drives structural adaptation within national higher education systems or merely reflects short-term technological hype (Cabrera and Neville, 2025; Trinché, 2025).

To understand how higher education systems respond to emerging technologies, scholars frequently draw upon foundational innovation theories. Rogers (2003) emphasizes that technology diffusion depends on perceived relative advantage, compatibility, and institutional trialability. Similarly, the Technology-Organization-Environment framework posits that technological adoption is shaped by organizational readiness and external policy pressures (Tornatzky and Fleischer, 1990). In academic institutions, developing dynamic capabilities allows organizations to reallocate resources and adapt competencies to evolving technological environments (Teece et al., 1997). Despite these theoretical foundations, the existing literature on generative artificial intelligence in higher education exhibits three major empirical gaps.

First, the vast majority of existing studies focus on micro-level or institutional-level case evidence. Researchers have extensively documented student perceptions, faculty digital literacy, and classroom pedagogical experiments (Contractor and Reyes, 2025; Haroud and Saqri, 2025; Meletiadou, 2025b; Najafov, 2025). Other investigations examine adaptive learning tools, student executive function support, and gender-specific attitudes toward AI tools (Al-Samarraie et al., 2025; Marzuki, 2025; Mittler, 2025). While these micro-level studies provide valuable pedagogical insights, they cannot explain how entire national higher education systems respond over time (Reibe, 2025).

Second, the literature lacks multi-year, cross-national macroeconomic evidence. While digital diffusion and research intensity have been examined across European economies (Tutar et al., 2026), generative AI adoption is rarely linked to macroeconomic indicators. System-level outcomes such as gross tertiary

enrollment, national research and development expenditure, and advanced population digital skills remain unexamined in relation to AI research intensity (Duong and Vu, 2025; Meletiadou, 2025a).

Third, existing macro-level analyses rely almost entirely on contemporaneous, same-year correlations. Same-year correlations cannot establish temporal precedence or directional causality. They cannot determine whether research attention in one year leads to measurable higher education adjustments in subsequent years, or whether both indicators simply reflect unobserved technological trends (Pesaran, 2021). A rigorous empirical test must evaluate multi-year leads and lags while testing bidirectional relationships.

1.2 The present study and contributions

This study addresses these empirical gaps by conducting a lead-lag panel analysis across all 27 European Union member states from 2020 through 2025. We evaluate whether country-level generative AI research activity, national policy initiatives, and public search interest predict subsequent shifts in three core higher education system outcomes. These outcomes comprise gross tertiary enrollment, research and development expenditure as a share of gross domestic product, and the national share of individuals possessing above-basic digital skills. This study makes three primary contributions to the literature. First, we construct a reproducible, open-source country-year panel capturing AI research attention and higher education indicators across the EU-27. Second, we employ a bidirectional lead-lag econometric specification with autoregressive controls. This design implements a Granger-style test of temporal precedence, testing whether AI attention precedes system adaptation or whether established system capacity precedes AI research output. Third, we incorporate panel diagnostic procedures to address cross-sectional dependence and structural heterogeneity across European economies (Pesaran, 2021; Tutar et al., 2026).

This research carries direct practical significance for university administrators, national funding agencies, and education ministries. Policymakers are actively deciding whether to reallocate substantial capital toward specialized AI faculty hiring, student digital literacy programs, and campus computing infrastructure (Shapiro and Lam, 2025; Tong et al., 2025). If AI attention actively drives system-level expansion, early investment signals strategic alignment. If current indicators primarily follow historical trajectories, leaders should avoid premature resource reallocation driven by anecdotal assumptions.

We also establish clear scope boundaries. This study does not measure micro-level cognitive learning gains or individual classroom performance. Instead, it measures national-level system responses. A lack of national-level acceleration does not preclude meaningful localized innovation occurring within specific institutions or departments.

1.3 Aim, objectives and research questions

This study aimed to evaluate the directional, longitudinal relationship between artificial intelligence research attention and higher education system responses across the European Union.

To achieve this aim, the study pursued three specific objectives: First, to construct a standardized country-year panel of generative AI indicators and higher education system metrics across the EU-27 from 2020 through 2025. Second, to test whether lagged AI research attention, policy actions, and search interest predict subsequent-year higher education enrollment, research spending, and digital skills. Third, to examine whether established higher education system capacity conversely predicts future AI research output, testing for bidirectional feedback.

These objectives guided three central research questions:

Research Question 1: Does national generative AI research attention in year t-1 predict higher education enrollment, research spending, or workforce digital skills in year t?

Research Question 2: Which form of AI attention (academic publications, public search volume, or national policy frameworks) exhibits the strongest longitudinal association with higher education indicators?

Research Question 3: Does baseline national higher education capacity predict future generative AI research output, indicating a bidirectional relationship?

To address these questions, we test four core hypotheses (H):

H1: Lagged generative AI research publication activity positively predicts next-year gross tertiary enrollment across EU-27 economies.

H2: Lagged generative AI research activity, AI policy activity, and public search interest positively predict next-year research and development expenditure.

H3: Lagged generative AI research activity, AI policy activity, and public search interest positively predict next-year population shares with above-basic digital skills.

H4: Lagged higher education system capacity positively predicts next-year generative AI publication output, confirming a bidirectional relationship.

1.4 Organization of the paper

The remainder of this paper is structured as follows. Section 2 reviews relevant theoretical frameworks and prior empirical literature. Section 3 details the data sources, indicator definitions, panel construction, and econometric models. Section 4 presents the descriptive findings, lead-lag regressions, and robustness evaluations. Section 5 discusses the theoretical and practical implications, addresses study limitations and outlines directions for future research. Section 6 presents concluding remarks.

2. LITERATURE REVIEW AND CONCEPTUAL FRAMEWORK

2.1 From adoption counts to institutional change

Most published work on GenAI in higher education measures adoption or attitudes at a single point in time (Crompton and Burke, 2023). A large recent wave of this work documents how students, instructors, and institutions have reacted to GenAI tools since late 2022, covering topics from academic integrity and assessment design to teaching practice and institutional guidance (Bittle and El-Gayar, 2025; Carmi, 2025; Ghilay, 2025). A smaller body of work asks whether AI-related activity changes institutional behavior, such as curriculum design, research priorities, or investment decisions (Cukurova, 2024). Recent bibliometric work on the GenAI-in-education literature itself finds that the volume of published research grew rapidly between 2022 and 2024, mirroring the pattern we document in previous studies (Frumin et al., 2026). This second question, whether attention translates into institutional change, is harder to answer than the first, because it requires data collected consistently over several years, and because institutional change often lags behind the appearance of a new technology by a year or more, particularly for decisions such as budget allocation or enrollment planning, which are typically set before an academic year begins rather than adjusted within it. Case-level evidence from individual institutions illustrates this lag directly: Contractor and Reyes (2025) trace how one elite college's response to GenAI unfolded gradually across several semesters rather than in a single step, and Trinché (2025) proposes a structured decision matrix precisely because institutions were found to prioritize GenAI-related risks and responses unevenly and slowly. Our lead-lag design, using a one-year lag between AI attention and system outcomes, is built specifically to capture this kind of delayed institutional response, though we note in Section 6 that a single-year lag may still be too short to capture slower administrative and budgetary cycles.

Prior technology adoption research offers a useful analogy and a body of formal theory for thinking about this lag. Rogers' diffusion of innovations framework describes adoption as a process that unfolds through distinct stages, knowledge, persuasion, decision, implementation, and confirmation, that different adopters pass through at different speeds depending on the perceived relative advantage, compatibility, and complexity of the innovation (Rogers, 2003). Wejnert (2002) extends this framework by distinguishing characteristics of the innovation itself from characteristics of the adopting units and the environment they operate in, a distinction that maps naturally onto our own separation between AI-attention variables and system-level outcome variables. Studies applying diffusion of innovations theory specifically to higher education technology adoption, such as work on open educational resources, generally find that institution-level adoption decisions precede measurable shifts in outcome data by more than a single year, and that organizational and policy-level factors, not just individual attitudes, shape the pace of adoption (Menzli, Smirani, Boulahia, and Hadjouni, 2022). The Technology-Organization-Environment framework offers a complementary lens, treating adoption as jointly shaped by the technology itself, the internal characteristics of the organization, and the external environment in which it operates (Tornatzky and Fleischer, 1990), which is consistent with our decision to test both AI-attention variables and structural system characteristics rather than either alone. If GenAI follows a similarly staged institutional adoption timeline, a one-year lag would be expected to show weak or absent forward effects even if a genuine, slower relationship exists. We treat this as a live possibility rather than a settled fact, and we discuss it directly when interpreting our results in Sections 4 and 5,

since it changes how a null result at one year should be read: as evidence against a fast relationship specifically, not necessarily as evidence against any relationship at all.

2.2 Learning analytics, AI support, and educational decision-making

The learning analytics literature argues that data-driven signals, including signals about emerging technology use, can support educational decision-making at both the classroom and institutional level (Cukurova, 2024). Much of this literature, however, is normative. It describes how AI-related data could support decisions, rather than testing whether AI-related attention has, in practice, preceded measurable changes in system-level indicators such as enrollment, research investment, or skills development. Governance-focused studies of GenAI in higher education illustrate this gap directly: Tong, Zainol, Chong, and Renganathan (20) analyze institutional policies for GenAI use and find considerable variation in how formally these policies are documented, while Krishnan (2025) maps the broader landscape of institutional strategies without testing whether any specific strategy measurably changed system outcomes afterward. Country-level, multi-year data offer one way to test this empirically, even if they cannot capture the specific institutional decision-making processes that individual universities go through. A country-level test also has a specific advantage for policy purposes: national ministries and EU-level funding bodies operate at exactly this level of aggregation, so evidence about whether AI attention precedes system-level shifts speaks directly to the level at which many funding and strategy decisions are actually made, even if it cannot speak to what happens inside any single institution.

2.3 Granger-style testing in educational panel research

Granger causality testing, originally developed for macroeconomic time series, asks whether past values of one variable improve the prediction of another variable's current value, beyond what the second variable's own past already predicts (Granger, 1969). Applied to a country-year panel, this approach requires including the outcome's own lag as a control, so that any apparent effect of a lagged predictor reflects information beyond simple persistence. Dumitrescu and Hurlin (2012) extended this logic to heterogeneous panels, allowing the causal relationship to differ across countries rather than assuming a single common coefficient, and Lopez and Weber (2017) packaged this test into accessible statistical software, which has since made panel Granger testing more common in cross-country economic research. This method has been used in economics and public health panel research more than in educational technology research, where cross-sectional and same-year designs remain more common. A closely related recent study applies a similar panel design, though not a full Granger test, to a different but related question: Tutar, Nam, Bayar, Varişli, and Kantarcioğlu (2026) examine whether digital diffusion and research and development intensity are associated with adult learning participation across the EU-27 from 2015 to 2023, using two-way fixed effects with Driscoll-Kraay standard errors and a first-difference robustness check, the same general estimation logic we use here. Their central finding, that a positive level association between digital diffusion and workforce learning does not survive a first-difference test, closely parallels our own results and is one we return to directly in Section 5.

Applying this general logic here lets us ask a sharper question: not just whether AI attention and system capacity move together, but whether AI attention actually comes first. Full panel Granger implementations typically also include diagnostics that our shorter panel cannot support: cross-sectional dependence tests such as Pesaran's CD test (Pesaran, 2021), panel unit root tests such as the Im-Pesaran-Shin test (Im, Pesaran, and Shin, 2003) or Pesaran's cross-sectionally augmented version (Pesaran, 2007), and panel cointegration tests such as Westerlund's (2007) error-correction approach. These diagnostics generally require more time periods than our six-year panel provides, so we do not run them here, and we flag this as a specific limitation in Section 6 rather than presenting our simpler lagged fixed-effects design as a complete Granger causality test in the formal econometric sense.

It is important to be precise about what a Granger-style test can and cannot establish. Finding that a lagged predictor significantly improves prediction of a later outcome, beyond the outcome's own history, is evidence of temporal precedence and incremental predictive power. It is not, by itself, proof of a causal mechanism, since both series could still respond to a shared earlier cause that our model does not observe directly. We use the term Granger-style rather than causal throughout this paper for this reason, and we encourage readers to treat our significant findings as evidence of a robust predictive relationship worth further investigation, rather than as a fully identified causal effect.

2.4 Reverse causality and system capacity as a driver of AI attention

A plausible alternative story is that countries with stronger existing higher education systems, more research investment, larger enrollment, or more advanced digital skills, are simply better positioned to produce GenAI research regardless of any change in AI-specific attention. The absorptive capacity perspective in innovation economics offers a formal version of this argument: Cohen and Levinthal (1990) propose that an organization's or economy's prior investment in related knowledge and research capacity determines how effectively it can recognize, assimilate, and apply new external knowledge, including knowledge about an emerging technology such as GenAI. Under this account, existing research capacity would be a precondition for, rather than a consequence of, GenAI-related research growth. Dynamic capabilities theory offers a related but distinct mechanism, emphasizing an organization's capacity to reconfigure resources and routines in response to a changing environment (Teece, Pisano, and Shuen, 1997; Winter, 2003), which would predict that countries with more adaptable, better-resourced higher education systems respond to GenAI more quickly regardless of public attention levels. Under either account, the relationship between system capacity and AI research activity would run from capacity to activity rather than the reverse, and existing capacity in one year would predict AI publication growth in the following year. Testing this reverse direction directly, alongside the forward direction, avoids assuming in advance which way the relationship runs. This concern is not merely theoretical. Several studies have found that digital skills, tertiary attainment, and research capacity variables were not significant same-year predictors of GenAI publication activity once country and year effects were removed, but a same-year null result does not rule out a delayed effect, where existing capacity takes a year to translate into observable research output (Zahra and George, 2002; Jeon, Hong, Ohm, and Yang,

2015; Venturini, Rincon-Aznar, and Vecchi, 2013). The reverse model in this study is designed specifically to close that gap.

2.5 Theoretical grounding for the three system outcomes

Choosing enrollment, research investment, and digital skills as system-level decision-making proxies draws on three distinct theoretical traditions, each of which motivates a different one of our forward hypotheses. Human capital theory treats education and skills development as a productive investment decision, made by individuals, institutions, and governments, that responds to changing returns and changing perceived needs (Becker, 1964). Under this framing, digital skills investment should respond to signals about which specific technologies are worth building capacity around, which is the logic behind our test of whether AI attention predicts next-year digital skills growth. Innovation-driven growth theory treats research and development spending as a deliberate response to the availability of new, potentially valuable knowledge to invest in, with research and development expenditure and innovation activity developing jointly within an economy (Grossman and Helpman, 1991), which is the logic behind our test of whether AI attention predicts next-year research investment. Automation and technological-change research adds a further, more cautious perspective: Acemoglu and Restrepo (2020) show that new technologies can affect labor market and educational outcomes unevenly, sometimes with lags and sometimes with effects that run in an unexpected direction, which is consistent with our own finding of a marginal negative association between GenAI publication activity and next-year enrollment rather than a straightforward positive one.

Institutional theory provides a further reason to expect that country-level responses to a common technological shock, such as the ChatGPT launch, might diverge rather than converge, even among relatively similar European economies. North (1990) argues that institutions, understood as the formal and informal rules governing economic and social behavior, shape how quickly and in what form a society responds to new opportunities or shocks. Hall and Soskice (2001) extend this argument specifically to comparative capitalism, showing that similarly developed economies can respond very differently to the same external pressure depending on their institutional configuration, a pattern broadly consistent with the divergent institutional trajectories often observed between EU West and EU East member states. Cross-country governance indicators of the kind compiled by Kaufmann, Kraay, and Mastruzzi (2011) are commonly used to capture this institutional variation in panel research, though our own panel does not include a direct institutional-quality measure, a gap we note in Section 6. The European Commission's own digital and skills policy framework, the 2030 Digital Compass, explicitly assumes that expanding digital infrastructure and strengthening workforce skills move together (European Commission, 2021), an assumption our results allow us to test directly rather than take for granted.

2.6 The GenAI-in-higher-education literature as an evolving evidence base

Because our AI-attention measures are built from the volume and content of GenAI-related publications, it is useful to characterize what that literature actually contains. Studies of institutional response and guidance make up one substantial strand, including Shapiro and Lam (2025) on unequal GenAI access across institutions, Meletiadou (2025) on lecturers' AI literacy, and Meletiadou (2025) on sustainable learning development frameworks. A second strand documents adoption patterns and the pace of institutional uptake, including Kaur (2026) on educators' concerns using structured decision-modeling techniques, Schoon, Smit, Uzuegbunam, and Bosch (2026) on adoption and perceptions among academics in South Africa, MacArthur (2025) on the reception of GenAI in writing instruction, Kalaitzidis (2025) on how GenAI has been positioned as a response to pre-existing problems in higher education, Ilieva, Yankova, Ruseva, and Kabaivanov (2025) on frameworks for GenAI-driven assessment, Apata (2025) on frameworks moving from acceptance toward deeper adoption, and Mittler (2025) on a specific case study of GenAI use in practice.

A third strand focuses on digital skills and digital literacy specifically, the same conceptual territory as one of our three system outcomes, including Marzuki (2025) on GenAI as an adaptive learning tool, Al-Zubaidi (2025) on institutional guidelines and generational gaps in GenAI-related digital literacy, Haroud and Saqri (2025) on teacher and student perspectives on digital literacy and GenAI, Volosevici, Eftimie, and Ghita (2025) on GenAI integration processes, Duong and Vu (2025) on GenAI incorporation and entrepreneurial self-efficacy, Estanyol, Fernández-Ardèvol, López-Borrull, Fernández-de-Castro, and Costa Gálvez (2025) on perception and use among communication students, and Bittle and El-Gayar (2025) on GenAI and academic integrity as a research agenda. A fourth strand addresses curriculum and enrollment-adjacent decisions, the second of our three system outcomes, including Sallai, Cardoso Silva, and Barreto (2025) on the risk of GenAI bypassing the learning process, Zhang (2025) on challenges and opportunities for course-level learning, Sozon, Parnther, Wong, and Chowdhury (2025) on navigating benefits and challenges of GenAI integration, Al-Samarraie, Sarsam, Alzahrani, Chatterjee, and Swinnerton (2025) on gender differences in GenAI perceptions relevant to equitable curriculum design, Schofield and Zhou (2025) on integrating GenAI into undergraduate and postgraduate curricula directly, Baker, Evans, and Farrelly (2025) on potential curriculum models for GenAI integration, and Krause, Panchal, and Ubhe (2025) on how GenAI is reshaping learning processes across higher education more broadly.

A fifth strand addresses equity and access and related digital divide themes, including Martin, Tubaltseva, Harrison, and Rubin (2025) on participatory approaches to GenAI-integrated assessment design, Isak and Isak (2025) on GenAI adoption in the Somali higher education context, Liu, Li, and Zhao (2025) on GenAI's potential to bridge or widen inequalities in higher education, and Yueh (2025) on philosophical and sociological perspectives on GenAI's uneven reach. A sixth and final strand reports direct perception and survey evidence from students and instructors, including Reibe (2025) on GenAI within a specific pedagogical tradition, Emre (2025) on integrating GenAI tools into practice, Baskar (2025) on general GenAI applications in higher education, Najafov (2025) on educators' perceptions and pedagogical approaches, Cabrera and Neville (2025) on the gap between widespread use and limited trust in GenAI tools, Ursavaş, Yalçın, İslamoğlu, Bakır-Yalçın, and Cukurova (2025) on the role of social

norms in GenAI adoption among students, and Hummel (2025) on biographical orientations toward GenAI in learning and teaching. Read together, this literature indicates that GenAI-related research activity in higher education already spans institutional governance, digital literacy, curriculum design, equity, and individual perception, which supports our decision to treat GenAI publication volume as a reasonable, if imperfect, proxy for the breadth of a country's engagement with the technology, rather than a narrow technical research niche.

2.7 Why country-level system indicators, and what they can and cannot show

A reasonable objection to this study's design is that enrollment, research expenditure, and digital skills are broad system indicators shaped by many forces unrelated to GenAI, so a country-level test may lack the sensitivity to detect a genuine but modest AI-driven effect. We accept this objection as a real limitation rather than a minor caveat. At the same time, country-level indicators have a specific value that institution-level data cannot fully replace: they are the level at which national ministries set targets, at which the European Commission benchmarks progress through instruments such as the Digital Decade and 2030 Digital Compass frameworks (European Commission, 2021), and at which cross-country comparisons of the kind policymakers actually use are made. A test conducted at this level, even an imperfect one, speaks directly to the claims most often made in AI-and-education policy discussion, which are themselves usually pitched at the country or system level rather than the classroom level. We see country-level and institution-level evidence as complementary rather than substitutable, and we position this study explicitly as the former.

2.8 Conceptual framework and hypotheses

Figure 1 summarizes the lead-lag framework used in this study. AI research and policy attention in year t-1, comprising GenAI publications, the AI policy index, and search interest, is modeled as a potential predictor of three higher education system outcomes in year t: tertiary enrollment, research and development expenditure, and digital skills. Each outcome's own lag enters the model as an autoregressive control, so that only additional predictive power beyond persistence counts as evidence of a forward relationship. A separate reverse model tests whether the three system outcomes in year t-1 predict GenAI publication activity in year t, again controlling for the publication series' own lag. This two-directional design lets the data show whether AI attention precedes system change, whether system capacity precedes AI attention, whether both directions hold, or whether neither holds beyond simple persistence, which was our finding, reported in Section 4, as the dominant pattern in three of the four models tested.

The four hypotheses introduced in Section 1 map directly onto this framework: H1 to H3 test the forward path, grounded respectively in human capital theory, innovation-driven growth theory, and diffusion of innovations theory, and H4 tests the reverse path, grounded in absorptive capacity and dynamic capabilities theory.

3. METHODOLOGY

3.1 Panel and variable construction

We built a country-year panel covering all 27 EU member states from 2020 to 2025 (N = 162 country-year observations), drawing on five underlying open data sources: Eurostat, the World Bank, OpenAlex, Google Trends, and the OECD AI Policy Observatory. Full details of these sources appear in the accompanying data dictionary. For this study, we constructed one-year lags, within each country, of three AI-attention variables (log-transformed GenAI and higher education publications, the AI policy index, and the Google Trends composite search interest index) and of three system outcome variables (gross tertiary enrollment, research and development expenditure as a share of GDP, and the share of the population with above-basic digital skills). We forward- and backward-filled missing values within each country before constructing lags, and we treat this as an open limitation given the incomplete annual coverage of the underlying Eurostat and World Bank series.

3.2 Outcome selection and its limits

We selected these three outcomes because they are the closest available proxies, within our open-data panel, to country-level higher education decision-making, even though they do not capture institution-level learning analytics adoption or curriculum-level learning design decisions, which were part of the original scope for this study but were not available without paid or restricted-access data sources such as the European Tertiary Education Register's full microdata or individual learning management system logs. Gross tertiary enrollment reflects national decisions and conditions around system access and capacity. Research and development expenditure as a share of GDP reflects national and institutional research investment decisions. The share of the population with above-basic digital skills reflects the cumulative result of workforce and curriculum decisions aimed at building digital capability. We present these as reasonable, available proxies for system-level decision-making rather than as a direct measurement of AI-specific learning analytics adoption, and we return to this distinction in Section 6.

3.3 Econometric strategy

We estimated four two-way fixed-effects panel models, using a demeaning and cluster-robust standard error approach. For each of the three forward models, the dependent variable is one system outcome in year t, and the regressors are that outcome's own value in year t-1 together with the three AI-attention variables in year t-1. Country and year fixed effects are removed through iterative double demeaning before estimation, and standard errors are clustered by country. For the reverse model, the dependent variable is log-transformed GenAI publications in year t, and the regressors are that series' own value in year t-1 together with the three system outcomes in year t-1, using the same fixed-effects and clustering approach.

This design directly tests Granger-style temporal precedence within a fixed-effects panel framework. A significant coefficient on a lagged AI-attention variable in a forward model indicates that AI attention predicts the outcome beyond what the outcome's own past trajectory already predicts. A significant

coefficient on a lagged system outcome in the reverse model indicates that system capacity predicts future AI attention beyond what AI attention's own past trajectory already predicts. Because our panel spans only six years, we limit lags to a single year, which preserves the largest possible number of observations, though it also limits our ability to detect relationships that unfold over a longer institutional planning cycle, a point we return to in Section 6.

All models were estimated using a custom Python implementation of two-way demeaning and Liang-Zeger cluster-robust covariance estimation, applied here to the lagged specification. Model outputs are reported in a dedicated results workbook alongside the lagged analysis panel.

3.4 Sample construction for lagged models

Building a one-year lag within each country necessarily drops each country's first available year from the lagged sample, since no prior-year value exists to lag from. Combined with missing-data patterns in the underlying Eurostat and World Bank series, this reduces the effective sample for the forward models to between 125 and 130 country-year observations across 25 to 26 countries, compared with 150 to 162 observations in the full contemporaneous panel. We report the exact sample size and number of countries for each model alongside its results in Section 4, rather than using a single pooled sample size across all four models, since listwise deletion for missing lagged regressors differs slightly by outcome.

3.5 Robustness considerations

Two design choices deserve explicit comment. First, we used a single one-year lag rather than testing multiple lag lengths, because our six-year panel does not have enough time periods to estimate longer lags with reasonable statistical power once fixed effects are removed; a two-year lag, for example, would leave only four usable time periods per country after accounting for both the lag itself and the double demeaning step. Second, we estimated each forward outcome in a separate model rather than a single combined system, so that missing data in one outcome variable does not reduce the usable sample for the other two outcomes. A joint vector autoregression treating all six series together, of the kind used in cross-country patent and innovation research (Guloglu and Tekin, 2012), would be a natural extension once more years of data accumulate, and we return to this point in Section 6. We also do not report Pesaran (2021) cross-sectional dependence diagnostics or Westerlund (2007) cointegration tests here, both standard components of a fuller panel causality workflow (Tutar, Nam, Bayar, Varişli, and Kantarcioğlu, 2026), because our six-year time dimension is too short for these tests to have reasonable power; we flag this explicitly as a methodological limitation in Section 6 rather than omit it silently.

3.6 Design rationale for variable roles

This study repositions three variables, tertiary attainment, research capacity, and digital skills, as outcomes rather than as predictors, reorganizing six regressors used elsewhere in the same-year GenAI-adoption literature into a lead-lag structure. This choice is deliberate rather than incidental. Much of the existing literature asks which contemporaneous digital, educational, and policy characteristics predict GenAI publication activity in the same year, treating tertiary attainment, research capacity, and digital

skills purely as predictors. This study asks the complementary question of whether GenAI-related activity, together with policy attention and search interest, itself predicts subsequent movement in a subset of those same higher education system characteristics, and whether the reverse relationship holds as well. The design is therefore best understood as examining the same underlying system from a different temporal angle, one-year lead-lag precedence rather than same-year association, rather than as an arbitrary reshuffling of variables.

4. RESULTS

4.1 Descriptive trends in the three outcome variables

Table 1 presents the descriptive statistics for outcome and AI-attention variables used in the lagged models. Similarly, Fig. 2 shows the EU-27 mean higher education system indicators from 2020 to 2025. Gross tertiary enrollment across the panel averages 78.8 percent of the relevant age cohort (SD = 22.8, range 20.0 to 166.4), reflecting substantial cross-country variation in how enrollment is measured and reported, including countries where enrollment exceeds the standard age cohort due to older or part-time students. Research and development expenditure averages 1.75 percent of GDP (SD = 0.89, range 0.46 to 3.60). The share of the population with above-basic digital skills averages 29.8 percent (SD = 11.4, range 7.7 to 55.6). Table 1 reports descriptive statistics for the outcome and AI-attention variables used in the lagged models. Figure 2 plots the EU-27 mean of each series across the six-year study period. Tertiary enrollment shows a mild upward drift over the period. Research and development expenditure shows a slower, more gradual increase. Digital skills shows the most uneven year-to-year pattern of the three, consistent with the partial and non-annual coverage of the underlying Eurostat digital skills survey noted in Section 3.1.

Table 1
Descriptive statistics for outcome and AI-attention variables used in the lagged models.

Variable	N	Mean	SD	Min	Max
Tertiary enrollment, gross (%)	122	78.81	22.76	19.96	166.39
R&D expenditure (% GDP)	108	1.75	0.89	0.46	3.60
Digital skills, above basic (%)	78	29.80	11.42	7.73	55.61
GenAI + HE publications (ln, 1 + x)	162	1.35	1.46	0.00	4.53
AI policy index (0–2)	162	1.56	0.57	0.00	2.00
Trends composite (search interest)	156	2.76	4.83	0.00	19.46

4.2 Forward models: does lagged AI attention predict system outcomes

Table 2 reports the three forward models. Figure 3 shows the lead-lag coefficients (standardized z-statistics) for AI attention (t-1) predicting higher education outcomes (t). In each model, the outcome's own one-year lag is the strongest and most consistent predictor of its current value, with coefficients of 0.716 (tertiary enrollment, $p < 0.001$), 0.447 (R&D expenditure, $p < 0.001$), and 0.619 (digital skills, $p < 0.001$), confirming that these system indicators are highly persistent from one year to the next, which is expected given that none of them reflects a genuinely fast-moving process.

Table 2
Forward two-way fixed-effects models: AI attention (t-1) predicting higher education system outcomes (t).

Outcome model	**Regressor**	**Coef.**	**Cluster SE**	**p**
Tertiary enrollment (t) N = 130, 26 countries, R^2=0.571	Enrollment (t-1)	0.7158	0.0737	< 0.001
	GenAI pubs (t-1)	-0.7459	0.4430	0.092
	AI policy index (t-1)	-0.2798	0.4554	0.539
	Trends composite (t-1)	0.0212	0.1273	0.868
R&D expenditure (t) N = 130, 26 countries, R^2=0.336	R&D expenditure (t-1)	0.4471	0.0382	< 0.001
	GenAI pubs (t-1)	0.0056	0.0040	0.156
	AI policy index (t-1)	0.0309	0.0350	0.377
	Trends composite (t-1)	0.0042	0.0014	0.003
Digital skills, above basic (t) N = 125, 25 countries, R^2=0.278	Digital skills (t-1)	0.6188	0.0619	< 0.001
	GenAI pubs (t-1)	0.6404	0.4211	0.128
	AI policy index (t-1)	0.0689	0.3595	0.848
	Trends composite (t-1)	0.1940	0.1969	0.325

Among the three AI-attention predictors, only one reaches conventional significance across the three forward models. Lagged search interest significantly predicts next-year research and development expenditure (coefficient 0.0042, $p = 0.003$), even though search interest itself is a relatively blunt measure of public curiosity rather than a direct signal of research funding decisions. Lagged GenAI publications show a negative coefficient on next-year tertiary enrollment that approaches but does not cross conventional significance (coefficient − 0.746, $p = 0.092$). Lagged AI policy index and lagged search interest are not significant predictors of tertiary enrollment. None of the three AI-attention variables significantly predicts next-year digital skills, though lagged GenAI publications show the

largest, if still non-significant, coefficient in that model (0.640, p = 0.128). Figure 3 presents these coefficients as standardized z-statistics across all three forward models for direct visual comparison.

This pattern gives limited support to H1 through H3. H2 receives partial support through the search interest and R&D expenditure link. H1 and H3 are not supported at conventional significance levels, though the enrollment result is close enough to significance to warrant caution rather than a firm negative conclusion.

4.3 Magnitude and practical significance

Statistical significance alone does not establish practical importance, so it is worth converting the one significant forward coefficient into a more concrete magnitude. The coefficient of 0.0042 on lagged search interest, predicting research and development expenditure as a share of GDP, implies that a 10-point increase in the Google Trends composite index, roughly the size of the jump the EU-27 mean experienced between 2024 and 2025, is associated with a next-year increase in R&D expenditure of about 0.042 percentage points of GDP, holding the outcome's own lag and the other regressors constant. Given that R&D expenditure across the panel averages 1.75 percent of GDP with a standard deviation of 0.89, an increase of this size is real but modest, representing less than one-twentieth of one standard deviation in the outcome. We report this magnitude explicitly so that readers can judge for themselves whether this represents a practically meaningful effect or a statistically detectable but substantively small one; we lean toward the latter reading, and we say so plainly rather than allowing statistical significance to imply more practical weight than the effect size supports.

4.4 Reverse model: does lagged system capacity predict AI attention

Table 3 reports the reverse model, in which lagged GenAI publications and lagged system outcomes are used to predict current GenAI publications. The autoregressive term, lagged GenAI publications, is significant and positive (coefficient 0.441, $p < 0.001$), consistent with the persistence seen in the forward models. None of the three lagged system outcomes, tertiary enrollment, R&D expenditure, or digital skills, significantly predicts next-year GenAI publication activity. This does not support H4. Higher education system capacity, as measured here, does not appear to be a leading indicator of future GenAI research activity once the publication series' own past trajectory is controlled.

Table 3
Reverse two-way fixed-effects model: higher education system outcomes (t-1) predicting GenAI publications (t). N = 130; 26 countries; $R^2 = 0.187$.

Regressor	Coef.	Cluster SE	p
GenAI pubs (t-1)	0.4408	0.0983	< 0.001
Tertiary enrollment (t-1)	-0.0008	0.0140	0.952
R&D expenditure (t-1)	0.4491	0.4526	0.321
Digital skills, above basic (t-1)	0.0102	0.0208	0.624

4.5 Summary of the four models

Taken together, the four models point to a specific and fairly narrow pattern rather than a broad bidirectional relationship. Three of the four models show essentially no significant cross-linkage between AI attention and system capacity, in either direction, beyond each series' own persistence. The exception is the positive link from lagged search interest to next-year research and development expenditure, which held up at a stringent significance threshold ($p = 0.003$) despite the modest sample available for lagged models (N = 130, 26 countries). We discuss the possible interpretations and limits of this single significant cross-linkage in Section 5.

It is also worth summarizing what did not turn out to matter, since a study built around testing several plausible relationships should report negative results as clearly as positive ones. Lagged AI policy activity was not a significant predictor of any of the three forward outcomes, nor was it a significant outcome in the reverse model. Lagged GenAI publications were not a significant predictor of digital skills or research investment, and their negative, near-significant coefficient on enrollment should be read as an open question rather than a confirmed effect. None of the three system-capacity variables significantly predicted future GenAI publications in the reverse model. Readers looking for a simple, uniform story in which AI attention broadly drives, or is broadly driven by, higher education system capacity will not find one in this panel. What the data show instead is a set of largely independent, highly persistent series, connected by one specific and modest cross-linkage rather than a general pattern of mutual reinforcement.

A useful way to summarize the four models together is to note what they rule out as much as what they support. They rule out a strong, broad forward relationship in which AI research and policy attention drives rapid change across enrollment, research spending, and digital skills simultaneously; had such a relationship existed, we would have expected to see significant lagged AI-attention coefficients in at least two, and plausibly all three, of the forward models, rather than in only one. They also rule out a strong reverse relationship in which existing system capacity is what allows AI research activity to grow; had that been the dominant mechanism, we would have expected at least one of the three lagged system outcomes to significantly predict future publications, which none did. What remains, after ruling out both of these broader stories, is a narrower and more specific finding: public attention to GenAI and

next-year research investment move together in a way that survives a conservative lag-and-autoregressive-control test, while the rest of the system, enrollment and digital skills in particular, has so far shown no measurable one-year response to AI-specific attention of any kind.

5. DISCUSSION

5.1. Key findings

This study set out to test whether AI research and policy attention precedes higher education system change, using a lead-lag panel design applied to the EU-27 country-year panel described in Section 3. Two findings stand out.

First, persistence, not AI-driven change, describes most of the year-to-year variation in the three system outcomes we examined. Each outcome's own lag was, by a wide margin, the strongest predictor of its current value in every model. This is not a surprising result on its own, since enrollment, research spending, and digital skills are all slow-moving system characteristics that would be expected to carry over strongly from one year to the next. It is nonetheless an important corrective to any framing that treats GenAI-era attention as a force capable of quickly reshaping these system indicators. Within the one-year horizon we can test, the data do not support such a framing for enrollment or digital skills, and they support it only narrowly for research investment.

Second, the one significant cross-linkage we found, from lagged search interest to next-year research and development expenditure, deserves a careful, non-overreaching interpretation. Search interest is a measure of public curiosity and search behavior, not a direct measure of institutional or governmental funding decisions. A plausible reading is that public attention to GenAI, concentrated heavily after the November 2022 ChatGPT launch, coincided with or slightly preceded a period in which several EU countries also increased research spending, for reasons that may or may not be causally connected to GenAI specifically. Another plausible reading is that both series responded to a shared underlying trend, general momentum in digital and AI-related investment during this period, without search interest itself doing any of the causal work. This second reading gains some support from Tutar, Nam, Bayar, Varişli, and Kantarcioğlu (2026), who find a positive level association between digital diffusion, research and development intensity, and adult learning participation across the EU-27 that does not survive a first-difference specification, and who explicitly caution against reading their level associations as causal or dynamic effects. Our design, using a single lag and cluster-robust standard errors, cannot fully rule out either interpretation, and we present this result as suggestive rather than conclusive.

The absence of a significant relationship between lagged GenAI publications and next-year digital skills is also informative, given how often public discussion assumes that growing AI research activity should translate fairly directly into workforce digital skills development. Our results do not support that assumption within the timeframe available to us. Digital skills, as measured through Eurostat's basic and above-basic skills surveys, likely respond to a broader and slower set of drivers, including national

education policy, labor market demand, and general digital infrastructure investment, that are not well captured by a single year's lag in AI-specific research or policy attention. This reading is consistent with human capital theory's emphasis on skills investment as a longer-horizon decision shaped by expected returns across an entire career rather than by a single year's research headlines (Becker, 1964), and with survey evidence from within the GenAI-in-higher-education literature itself showing that institutional digital literacy initiatives tend to lag behind tool adoption by students and instructors rather than lead it (Haroud and Saqri, 2025; Al-Zubaidi, 2025).

The marginal, non-significant negative association between lagged GenAI publications and next-year enrollment also deserves brief comment, since a negative sign runs against the simple expectation that more AI research activity would coincide with growing student interest and higher enrollment. One possibility, consistent with the uneven and sometimes counterintuitive distributional effects of new technologies documented in the broader automation literature (Acemoglu and Restrepo, 2020), is that countries experiencing a rapid rise in GenAI-related research attention are not necessarily the same countries experiencing rising tertiary enrollment in the same period, so the observed association may reflect compositional differences across countries rather than any within-country dynamic. Because this coefficient does not reach conventional significance, we do not treat it as an established pattern, but we also do not discard it, since a marginal result close to conventional thresholds in a modest panel deserves monitoring in future work rather than automatic dismissal.

The reverse model result, that past system capacity does not significantly predict future GenAI publication growth, is a genuinely informative null result set against absorptive capacity theory, which would have predicted that countries with more pre-existing research capacity should show faster subsequent growth in GenAI-specific output (Cohen and Levinthal, 1990), and against dynamic capabilities accounts that treat resource reconfiguration capacity as a precondition for responding to a new technology (Teece, Pisano, and Shuen, 1997). Neither mechanism finds support here. Together, the two studies suggest a consistent picture: GenAI research activity in European higher education has been driven primarily by a rapid, largely exogenous shift in public attention following the ChatGPT launch, rather than by the prior accumulation of digital skills, research capacity, or policy groundwork in any given country. Countries did not need years of prior capacity-building to begin producing GenAI research at scale. They needed researchers and the public paying attention at the same time, and that attention arrived on a similar timeline across most of the EU-27 following a single global technological event, a pattern broadly consistent with diffusion of innovations accounts in which a sufficiently disruptive innovation can compress the knowledge and persuasion stages of adoption across many adopting units simultaneously (Rogers, 2003; Wejnert, 2002).

For institutions and policymakers, this carries a specific, moderately cautious message. Building general digital infrastructure and skills remains valuable for many reasons, but our results do not support treating it as the primary lever for building GenAI-specific research capacity in the short run. Similarly, a rise in AI-related research attention should not, by itself, be treated as a leading indicator that broader enrollment, funding, or skills outcomes are about to shift within the following year. The one exception,

the tentative link between search interest and research investment, suggests that tracking public attention trends could offer institutions and funders early, if imperfect, signals about where research demand may be heading, a point we develop further as a direction for future research in Section 6.

The search interest finding is also worth situating within this study's own forward and reverse results. This study shows that the same search interest signal carries some forward-looking information about research investment one year later, while carrying no significant forward-looking information about enrollment or digital skills. This is a genuinely useful, if narrow, finding for anyone trying to decide which of the many available indicators to monitor going forward: a country's basic digital skills score or its formal AI strategy documents carry comparatively little predictive signal in this panel, while public search behavior carries a disproportionate share of what predictive signal exists here.

At the same time, we want to avoid overstating the reach of a single administrative-style indicator. Google Trends search interest is shaped by many factors beyond genuine research-relevant curiosity, including media coverage cycles, the wording of specific search queries, and platform-level changes in how Google reports trend data over time. A single significant coefficient, however carefully estimated, is not sufficient grounds for building major funding or strategy decisions around search interest monitoring alone. We present it as one useful, low-cost signal among several that institutions and funders might reasonably track, not as a validated leading indicator ready for direct policy use.

It is also worth noting what these findings do not show. This study asked whether the search interest variable, together with publications and policy activity, also predicts what happens next, and largely found that it does not, with one exception. Read on their own, these results suggest that public attention operates mainly as a same-year companion to research activity rather than as an early warning signal for broader system change. This is a more modest role for public attention than some AI-adoption narratives assume, but it is also a more precisely specified one, and it is the role that our data actually support.

It is worth pausing on why a null result for enrollment and digital skills is itself informative, rather than simply an absence of findings. Enrollment decisions in most EU higher education systems are shaped by demographic trends, tuition and funding policy, and broader labor market signals that operate over multi-year horizons and are largely independent of any single emerging technology. Our finding that lagged AI attention does not significantly predict next-year enrollment is consistent with enrollment being driven by these slower, more structural forces rather than by a fast-moving technology narrative. Similarly, national digital skills levels reflect years of accumulated investment in schooling, adult education, and labor-market training, so it would be somewhat surprising, on reflection, if a single year of AI-specific research or policy attention shifted that measure noticeably. Framed this way, our null results are not merely a failure to detect an effect. They are consistent with a plausible, and in our view more realistic, account of how these system indicators actually move.

We also note that our results do not speak to whether GenAI is changing teaching and assessment practice inside individual classrooms, which is a separate and, in our view, probably faster-moving

process than the country-level system indicators we examine here. A university department might change how it assesses student writing within a single semester in response to GenAI, well before that change shows up in national enrollment or digital skills statistics, if it ever does directly. Our country-level null results should therefore not be read as evidence that GenAI is having no effect on higher education practice. They should be read specifically as evidence that AI-related attention has not yet produced a measurable, system-wide shift in the three specific indicators we examined, within the one-year window our data can test.

This distinction between classroom-level and system-level change also has a bearing on how the wider educational technology literature should interpret null results of this kind. A long history of technology-in-education research has shown that individual teachers and departments often adopt new tools well ahead of any visible shift in national statistics, sometimes by several years (Laurillard, 2012). GenAI research and interaction may currently be in this early, classroom-level phase across much of the EU-27, a phase that our system-level indicators are simply not built to detect. If that is correct, a repeat of this study in three or four years, once GenAI use has had more time to work through curriculum committees, accreditation cycles, and national statistical reporting, might find forward relationships that are not yet visible in the 2020 to 2025 window we study here. We see this as a reason for continued monitoring rather than a reason to dismiss the country-level approach altogether, since system-level indicators remain the level at which most policy decisions and funding allocations are actually made and evaluated.

A further point concerns how our findings should inform the balance between infrastructure investment and demand-side attention in AI-related education policy. Some national strategies have emphasized building digital infrastructure and skills capacity as a precondition for AI adoption, on the assumption that capacity constraints are the main barrier to research and institutional engagement. This study's finding that lagged system capacity does not predict future GenAI activity suggests that capacity constraints, at least as measured by our variables, have not been the binding constraint on GenAI research growth in the EU-27 during this period. This does not mean capacity investment is unimportant for other reasons, but it does suggest that framing capacity-building as the primary lever for accelerating GenAI research engagement may not match what the data show about how that engagement has actually grown so far.

5.2. Limitations and future research

This study is subject to data limitations inherent to the underlying open data sources, including substantial missing data in several Eurostat and World Bank series and the treatment of missing GenAI publication counts as zero rather than as genuinely unobserved. Constructing one-year lags compounds these limitations somewhat, since a country with a gap in one year loses both that year's observation and the following year's lagged predictor, which further reduces the effective sample size for the lagged models (N = 125 to 130, compared with N = 150 to 162 in the full contemporaneous panel).

Second, and specific to this study's design, a one-year lag may be too short to capture some forms of institutional response. Enrollment planning, research budget cycles, and national digital skills strategies often operate on multi-year planning horizons, so a genuine effect of AI attention on these outcomes might appear only after two or three years, which our six-year panel cannot test with the statistical power needed for longer lags. Third, as discussed in Section 3.2, our three outcome variables are country-level proxies for higher education decision-making, not direct measures of institutional learning analytics adoption or curriculum-level learning design choices, which were part of the original scope for this study but were not available through open data sources at the time of writing. Fourth, our reverse model tests only whether system capacity predicts future GenAI publications, not whether it predicts future AI policy activity or search interest specifically; a fuller vector autoregression treating all six variables jointly, rather than the two separate directional models used here, would provide a more complete picture of the dynamic relationships between these series, at the cost of requiring a larger panel than our current six-year window supports. Fifth, our cluster-robust standard errors assume that dependence between countries is fully captured by within-country clustering, an assumption that may understate uncertainty if some cross-country dependence, such as shared EU funding programs, is not captured by this approach.

Future research should extend this design in three directions. First, once additional years of data become available, extending the panel and testing two- and three-year lags would clarify whether the null forward results reported here for enrollment and digital skills reflect a genuinely weak relationship or simply a lag structure that is too short. Second, replacing our country-level proxies with genuine institution-level learning analytics adoption data, gathered directly from a sample of universities rather than from open national statistics, would provide a more direct test of the original research question this study was designed around. Third, a full panel vector autoregression across all six series, AI attention and system outcomes together, would let researchers test the entire system of relationships jointly rather than through the two separate directional models used in this study.

A fourth direction concerns the search interest and research investment link specifically. Because this is the only significant cross-linkage found across four models, it deserves targeted follow-up rather than being treated as a settled finding. A useful next step would be to test whether this relationship holds at a finer geographic level, such as regions or individual universities within a country, and whether it holds specifically for research funding categories connected to computing and artificial intelligence, rather than for total research and development expenditure across all fields, which is the more aggregated measure available in our current panel. If the relationship is genuine and specific to AI-relevant research funding, disaggregating the outcome variable in this way should produce an even clearer and larger effect; if the relationship instead reflects a broader, unrelated funding trend that happened to coincide with rising search interest, disaggregation should weaken or eliminate it.

Two further limitations deserve mention. Sixth, our AI policy index combines two binary indicators, whether a country had an active national AI strategy and whether that strategy referenced higher education, into a simple 0 to 2 score. This index reached its maximum value for nearly all EU-27 countries by 2022, leaving little year-to-year variation for a lagged model to exploit in later years, which

likely understates whatever genuine policy effect exists and should be read as a data limitation rather than firm evidence that policy activity has no forward effect on system outcomes. Seventh, this study relies on OpenAlex as the sole bibliometric source for GenAI publication counts. A different search strategy or a different bibliometric database, such as Scopus or Web of Science, could return a different publication series, and while we have no specific reason to expect this would change our substantive conclusions, we cannot rule it out without direct comparison, which was outside the scope of what free, open data sources could support for this study.

Finally, we want to be direct about a broader methodological point that applies to both studies in this pair. Open, free data sources make a reproducible, six-year, 27-country panel possible without requiring institutional access or paid subscriptions, which is a genuine advantage for transparency and replication. That same openness means the panel is built from whatever variables happen to be tracked consistently by Eurostat, the World Bank, OpenAlex, Google Trends, and the OECD AI Policy Observatory, rather than from variables specifically designed to measure GenAI-related institutional decision-making. Future work with access to richer, more specific data sources, whether through institutional partnerships, paid bibliometric databases, or direct survey instruments, will likely be able to test sharper versions of the questions raised here than an open-data panel alone can support.

One additional limitation concerns the interpretation of our single significant finding. Because we tested four models, each with three or four coefficients, a total of thirteen coefficients across the study, a strict correction for multiple comparisons would raise the bar needed to call any single result significant. Our reported p-value of 0.003 for the search interest and R&D expenditure link would survive a standard Bonferroni correction across these thirteen tests, but readers should still treat this as one plausible finding among many tested hypotheses rather than as a preregistered, singularly targeted result. We did not preregister this analysis, and we report this openly so that readers can weigh the finding accordingly.

6. CONCLUSIONS

This study tested whether AI research and policy attention in European higher education precedes system-level change, using a lead-lag panel design applied to an EU-27 country-year panel. Across four models, persistence dominated: each outcome's own past value was the strongest predictor of its current value in every specification. Among the AI-attention measures tested, only lagged public search interest showed a significant forward relationship, predicting next-year research and development expenditure, while GenAI publications and AI policy activity did not significantly predict next-year enrollment, research spending, or digital skills. The reverse test found no evidence that prior system capacity predicts future GenAI publication growth. Together, these results suggest that GenAI research growth in European higher education has followed a largely public-attention-driven path, rather than a path built on pre-existing digital or research capacity, and that this attention has not yet translated into broad, measurable shifts in higher education system indicators within a one-year horizon. Four answers follow from our research questions. On H1, lagged GenAI publications did not significantly predict next-year tertiary enrollment, though the coefficient was negative and approached significance, a result we

treat as inconclusive rather than as evidence of a genuine negative effect. On H2, lagged search interest, though not lagged publications or policy activity, significantly predicted next-year research and development expenditure, giving partial support to this hypothesis. On H3, none of the three AI-attention measures significantly predicted next-year digital skills, offering no support for this hypothesis within our one-year lag structure. On H4, lagged system outcomes did not significantly predict future GenAI publication activity, offering no support for a reverse causal path from system capacity to AI research growth. Researchers and policymakers should treat claims that AI attention is already reshaping enrollment, research investment, or skills outcomes with caution, while continuing to monitor these relationships as more years of data become available. As the EU-27 panel extends beyond 2025, and as more years accumulate on both sides of the November 2022 ChatGPT launch, the lead-lag tests reported here can be repeated with longer lags and a larger effective sample, which should sharpen rather than overturn the cautious picture presented in this study. Until then, the most defensible reading of the evidence is a narrow one: public attention to GenAI has grown quickly and has tracked research output closely in the same year, but it has not yet shown a clear, broad footprint in the slower-moving system indicators that higher education policy ultimately cares about.

Declarations

Author Contribution

G.O.F.: Conceptualization, Methodology, Formal analysis, Data curation, Validation, Investigation, Visualization, Writing – original draft, review & editing, Resources, Supervision, Project administration. E.O-D.: Conceptualization, Methodology, Formal analysis, Data curation, Validation, Investigation, Visualization, Writing – original draft, review & editing, Resources, Project administration.

Acknowledgement

The authors would like to thank the Department of Forest Engineering, Forest Management Planning, and Terrestrial Measurements, Faculty of Silviculture and Forest Engineering, Transilvania University of Brasov, for providing some of the resources needed for this study.

Data Availability

The datasets and the Python 3.12 script used for data analysis are available in the figshare repository at: https://figshare.com/s/66ef3b432d267dff2e79.

References

1. Acemoglu, D., & Restrepo, P. (2020). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, *128*(6), 2188–2244. https://doi.org/10.1086/705716

2. Al-Samarraie, H., Sarsam, S. M., Alzahrani, A. I., Chatterjee, A., & Swinnerton, B. J. (2025). Gender perceptions of generative AI in higher education. *Journal of Applied Research in Higher Education*. Advance online publication. https://doi.org/10.1108/jarhe-02-2024-0109
3. Al-Zubaidi, K. (2025). *The role of generative AI in higher education: Institutional guidelines, generational gaps, and the Grok 4 challenge* [Preprint]. OSF Preprints. https://doi.org/10.31235/osf.io/rq8xg_v1
4. Apata, O. E. (2025). *From acceptance to empowerment: A new framework for ethical generative AI adoption and student persistence in higher education* [Preprint]. SSRN. https://doi.org/10.2139/ssrn.5724262
5. Baker, N., Evans, W., & Farrelly, T. (2025). Potential curriculum models for intentionally incorporating generative AI in higher education. In *Advances in computational intelligence and robotics* (pp. 41–64). IGI Global. https://doi.org/10.4018/979-8-3373-7579-3.ch003
6. Baskar, S. (2025). Generative AI. In *Advances in computational intelligence and robotics* (pp. 111–134). IGI Global. https://doi.org/10.4018/979-8-3693-5623-4.ch006
7. Becker, G. S. (1964). *Human capital: A theoretical and empirical analysis, with special reference to education*. University of Chicago Press.
8. Bittle, K., & El-Gayar, O. (2025). Generative AI and academic integrity in higher education: A systematic review and research agenda. *Information*, *16*(4). Article 296. https://doi.org/10.3390/info16040296
9. Cabrera, C., & Neville, R. (2025). *Widely used, barely trusted? Understanding student perceptions on the use of generative AI tools in higher education* [Preprint]. OSF Preprints. https://doi.org/10.31235/osf.io/hxc4u_v1
10. Carmi, G. (2025). Learning with generative AI: An empirical study of students in higher education. *Education Sciences*, *15*(12). Article 1696. https://doi.org/10.3390/educsci15121696
11. Choi, H., & Varian, H. (2012). Predicting the present with Google Trends. *Economic Record*, *88*(s1), 2–9. https://doi.org/10.1111/j.1475-4932.2012.00809.x
12. Cohen, W. M., & Levinthal, D. A. (1990). Absorptive capacity: A new perspective on learning and innovation. *Administrative Science Quarterly*, *35*(1), 128–152. https://doi.org/10.2307/2393553
13. Contractor, Z., & Reyes, G. (2025). *Generative AI in higher education: Evidence from an elite college* [Preprint]. SSRN. https://doi.org/10.2139/ssrn.5392712
14. Cotton, D. R. E., Cotton, P. A., & Shipway, J. R. (2024). Chatting and cheating: Ensuring academic integrity in the era of ChatGPT. *Innovations in Education and Teaching International*, *61*(2), 228–239. https://doi.org/10.1080/14703297.2023.2190148
15. Crompton, H., & Burke, D. (2023). Artificial intelligence in higher education: The state of the field. *International Journal of Educational Technology in Higher Education*, *20*., Article 22. https://doi.org/10.1186/s41239-023-00392-8
16. Cukurova, M. (2024). *The interplay of learning, analytics, and artificial intelligence in education: A vision for hybrid intelligence* (arXiv:2403.16081) [Preprint]. arXiv. https://arxiv.org/abs/2403.16081

17. Dumitrescu, E. I., & Hurlin, C. (2012). Testing for Granger non-causality in heterogeneous panels. *Economic Modelling*, *29*(4), 1450–1460. https://doi.org/10.1016/j.econmod.2012.02.014
18. Duong, C. D., & Vu, T. N. (2025). Entrepreneurial education and higher education students' e-entrepreneurial intention: A moderated mediation model of generative AI incorporation and e-entrepreneurial self-efficacy. *Higher Education, Skills and Work-Based Learning*. Advance online publication. https://doi.org/10.1108/heswbl-12-2024-0390
19. Emre, İ. E. (2025). Integrating generative AI tools in higher education. *American Journal of STEM Education*. Advance online publication. https://doi.org/10.32674/2gxk8g16
20. Estanyol, E., Fernández-Ardèvol, M., López-Borrull, A., Fernández-de-Castro, P., & Costa Gálvez, L. (2025). Perception and use of generative AI among corporate communication students in higher education: Adjusting expectations. *Cuadernos info*, *62*, 1–21. https://doi.org/10.7764/cdi.62.89632
21. European Commission (2021). *2030 Digital Compass: The European way for the Digital Decade* (COM(2021) 118 final). https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021DC0118
22. Frumin, I., Vorochkov, A., Kiryushina, M., Platonova, D., & Terentiev, E. (2026). Mapping the generative AI research in higher education: 2022–2024 insights. *Higher Education Quarterly*. Advance online publication. https://doi.org/10.1111/hequ.70075
23. Ghilay, Y. (2025). *Generative AI in higher education: Teaching, assessment, and research in the AI era* [Preprint]. SSRN. https://doi.org/10.2139/ssrn.5345033
24. Granger, C. W. J. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, *37*(3), 424–438. https://doi.org/10.2307/1912791
25. Grossman, G. M., & Helpman, E. (1991). *Innovation and growth in the global economy*. MIT Press.
26. Guloglu, B., & Tekin, R. B. (2012). A panel causality analysis of the relationship among research and development, innovation, and economic growth in high-income OECD countries. *Eurasian Economic Review*, *2*(1), 32–47. https://doi.org/10.14208/BF03353831
27. Hall, P. A., & Soskice, D. (Eds.). (2001). *Varieties of capitalism: The institutional foundations of comparative advantage*. Oxford University Press.
28. Haroud, S., & Saqri, N. (2025). Generative AI in higher education: Teachers' and students' perspectives on support, replacement, and digital literacy. *Education Sciences*, *15*(4). Article 396. https://doi.org/10.3390/educsci15040396
29. Hummel, S. (2025). Higher education under generative AI: Biographical orientations of democratic learning and teaching. *Education Sciences*, *15*(12). Article 1572. https://doi.org/10.3390/educsci15121572
30. Ilieva, G., Yankova, T., Ruseva, M., & Kabaivanov, S. (2025). A framework for generative AI-driven assessment in higher education. *[Preprint] Preprints*. https://doi.org/10.20944/preprints202505.0800.v1
31. Im, K. S., Pesaran, M. H., & Shin, Y. (2003). Testing for unit roots in heterogeneous panels. *Journal of Econometrics*, *115*(1), 53–74. https://doi.org/10.1016/S0304-4076(03)00092-7

32. Isak, M. A., & Isak, A. A. (2025). *Assessing generative AI tools in Somali higher education and its impact on student and faculty performance at Somali National University* [Preprint]. Research Square. https://doi.org/10.21203/rs.3.rs-7292768/v1
33. Jeon, J., Hong, S., Ohm, J., & Yang, T. (2015). Causal relationships among technology acquisition, absorptive capacity, and innovation performance: Evidence from the pharmaceutical industry. *PLOS ONE*, *10*(7). https://doi.org/10.1371/journal.pone.0131642. Article e0131642.
34. Jun, S. P., Yoo, H. S., & Choi, S. (2018). Ten years (2008–2017) of Google Trends research: A systematic literature review. *Technological Forecasting and Social Change*, *130*, 320–337. https://doi.org/10.1016/j.techfore.2017.09.023
35. Kalaitzidis, T. J. (2025). *How generative AI fixes what higher education broke* [Preprint]. OSF Preprints. https://doi.org/10.31235/osf.io/yh6rs_v2
36. Kaufmann, D., Kraay, A., & Mastruzzi, M. (2011). The worldwide governance indicators: Methodology and analytical issues. *Hague Journal on the Rule of Law*, *3*(2), 220–246. https://doi.org/10.1017/S1876404511200046
37. Kaur, I. (2026). Generative AI in higher education: A deep dive into educators' concerns using the DEMATEL technique. *Asian Education and Development Studies*. Advance online publication. https://doi.org/10.1108/aeds-04-2025-0176
38. Krause, S., Panchal, B. H., & Ubhe, N. (2025). Evolution of learning: Assessing the transformative impact of generative AI on higher education. *Frontiers of Digital Education*. Advance online publication. https://doi.org/10.1007/s44366-025-0058-7
39. Krishnan, S. (2025). Scoping the landscape: Opportunities, challenges, and strategies for generative AI in higher education. *Indian Journal of Computer Science*, *10*(4). Article 175747. https://doi.org/10.17010/ijcs/2025/v10/i4/175747
40. Laurillard, D. (2012). *Teaching as a design science: Building pedagogical patterns for learning and technology*. Routledge.
41. Liu, Y., Li, Y., & Zhao, T. (2025). Bridging inequalities in higher education with the use of Gen AI. In *Advances in computational intelligence and robotics* (pp. 51–72). IGI Global. https://doi.org/10.4018/979-8-3693-5623-4.ch003
42. Lopez, L., & Weber, S. (2017). Testing for Granger causality in panel data. *The Stata Journal*, *17*(4), 972–984. https://doi.org/10.1177/1536867X1801700412
43. MacArthur, M. (2025). Generative AI and writing in higher education: Editor's introduction. *Critical AI*, *2*(1). https://doi.org/10.1215/2834703x-12095937
44. Martin, A. F., Tubaltseva, S., Harrison, A., & Rubin, G. J. (2025). *Participatory co-design and evaluation of a novel approach to generative AI-integrated coursework assessment in higher education* [Preprint]. Preprints. https://doi.org/10.20944/preprints202505.0016.v1
45. Marzuki, A. G. (2025). *Integrating generative AI as an adaptive learning tool in English language instruction in higher education* [Preprint]. Preprints. https://doi.org/10.20944/preprints202508.0426.v1

46. Meletiadou, E. (2025a). Generative artificial intelligence for sustainable learning development. In *Advances in computational intelligence and robotics* (pp. 23–50). IGI Global. https://doi.org/10.4018/979-8-3693-5623-4.ch002
47. Meletiadou, E. (2025b). The effective use of generative AI in higher education exploring lecturers' artificial intelligence literacy. In *Advances in computational intelligence and robotics* (pp. 1–22). IGI Global. https://doi.org/10.4018/979-8-3693-5623-4.ch001
48. Menzli, L. J., Smirani, L. K., Boulahia, J. A., & Hadjouni, M. (2022). Investigation of open educational resources adoption in higher education using Rogers' diffusion of innovation theory. *Heliyon*, *8*(7). https://doi.org/10.1016/j.heliyon.2022.e09885. Article e09885.
49. Mittler, S. (2025). Harnessing generative AI to overcome executive dysfunction in higher education: A case study. In *11th International Conference on Higher Education Advances (HEAd'25)*. Universitat Politècnica de València. https://doi.org/10.4995/head25.2025.20077
50. Najafov, R. (2025). *Educators' perceptions and pedagogical approaches in the era of generative AI integration: A qualitative study in higher education* [Preprint]. *Research Square*. https://doi.org/10.21203/rs.3.rs-8330375/v1
51. North, D. C. (1990). *Institutions, institutional change and economic performance*. Cambridge University Press.
52. Pesaran, M. H. (2007). A simple panel unit root test in the presence of cross-section dependence. *Journal of Applied Econometrics*, *22*(2), 265–312. https://doi.org/10.1002/jae.951
53. Pesaran, M. H. (2021). General diagnostic tests for cross-sectional dependence in panels. *Empirical Economics*, *60*(1), 13–50. https://doi.org/10.1007/s00181-020-01875-7
54. Reibe, N. (2025). Generative AI and the Ignatian pedagogical paradigm. *Jesuit Higher Education: A Journal*, *14*(1). https://doi.org/10.15365/2164-7666.1520
55. Rogers, E. M. (2003). *Diffusion of innovations* (5th ed.). Free.
56. Sallai, D., Silva, C., J., & Barreto, M. (2025). *Approach generative AI tools proactively or risk bypassing the learning process in higher education* [Preprint]. SSRN. https://doi.org/10.2139/ssrn.5259749
57. Schofield, L., & Zhou, X. (2025). Integrating generative AI into undergraduate and postgraduate curricula: An interactive workshop. *Journal of Learning Development in Higher Education*, *37*. https://doi.org/10.47408/jldhe.vi37.1712
58. Schoon, A., Smit, A., Uzuegbunam, C., & Bosch, T. (2026). Adoption and perceptions of generative AI among South African academics. *Teaching in Higher Education*. Advance online publication. https://doi.org/10.1080/13562517.2025.2571911
59. Shapiro, S. J., & Lam, V. D. (2025). Bridging the AI divide. In *Advances in computational intelligence and robotics* (pp. 115–138). IGI Global. https://doi.org/10.4018/979-8-3373-1195-1.ch006
60. Sozon, M., Parnther, C., Wong, W. L., & Chowdhury, M. A. (2025). Generative AI in higher education: Navigating benefits and challenges in the technological era. *Journal of Applied Research in Higher Education*. Advance online publication. https://doi.org/10.1108/jarhe-02-2025-0103

61. Teece, D. J., Pisano, G., & Shuen, A. (1997). Dynamic capabilities and strategic management. *Strategic Management Journal*, *18*(7), 509–533. https://doi.org/10.1002/(SICI)1097-0266(199708)18:7%3C509::AID-SMJ882%3E3.0.CO;2-Z
62. Tong, A., Zainol, Z., Chong, T. S., & Renganathan, K. (2025). AI governance on young consumers in higher education: A content analysis of policies for generative AI. *Young Consumers*. Advance online publication. https://doi.org/10.1108/yc-10-2024-2303
63. Tornatzky, L. G., & Fleischer, M. (1990). *The processes of technological innovation*. Lexington Books.
64. Trinché, T. L. A. (2025). *Decision matrix for prioritizing generative AI risks in higher education* [Preprint]. Research Square. https://doi.org/10.21203/rs.3.rs-7494968/v1
65. Tutar, H., Nam, S., Bayar, M., Varişli, N., & Kantarcioglu, N. (2026). Digital diffusion, R&D intensity, and adult learning participation: Panel evidence from EU-27 economies. *Administrative Sciences*, *16*(7). Article 311. https://doi.org/10.3390/admsci16070311
66. Ursavaş, Ö. F., Yalçın, Y., İslamoglu, H., Bakır-Yalçın, E., & Cukurova, M. (2025). Rethinking the importance of social norms in generative AI adoption: Investigating the acceptance and use of generative AI among higher education students. *International Journal of Educational Technology in Higher Education*. Advance online publication. https://doi.org/10.1186/s41239-025-00535-z
67. Venturini, F., Rincon-Aznar, A., & Vecchi, M. (2013). *ICT as a general purpose technology: Spillovers, absorptive capacity and productivity performance* (NIESR Discussion Paper No. 416). National Institute of Economic and Social Research.
68. Volosevici, D., Eftimie, S., & Ghita, I. A. (2025). Integration of generative AI into higher education. In *Advances in computational intelligence and robotics* (pp. 142–165). IGI Global. https://doi.org/10.4018/979-8-3373-7579-3.ch007
69. Wejnert, B. (2002). Integrating models of diffusion of innovations: A conceptual framework. *Annual Review of Sociology*, *28*, 297–326. https://doi.org/10.1146/annurev.soc.28.110601.141051
70. Westerlund, J. (2007). Testing for error correction in panel data. *Oxford Bulletin of Economics and Statistics*, *69*(6), 709–748. https://doi.org/10.1111/j.1468-0084.2007.00477.x
71. Winter, S. G. (2003). Understanding dynamic capabilities. *Strategic Management Journal*, *24*(10), 991–995. https://doi.org/10.1002/smj.318
72. Yueh, R. (2025). Generative AI in higher education: Philosophical, sociological, and pedagogical perspectives. *International Journal of Arts Humanities and Social Science*, *6*(4). https://doi.org/10.56734/ijahss.v6n4a6
73. Zahra, S. A., & George, G. (2002). Absorptive capacity: A review, reconceptualization, and extension. *Academy of Management Review*, *27*(2), 185–203. https://doi.org/10.5465/AMR.2002.6587995
74. Zhang, J. (2025). Generative AI in higher education: Challenges and opportunities for course learning. *Advances in Social Sciences Research Journal*, *12*(1). https://doi.org/10.14738/assrj.1201.18121

Figures

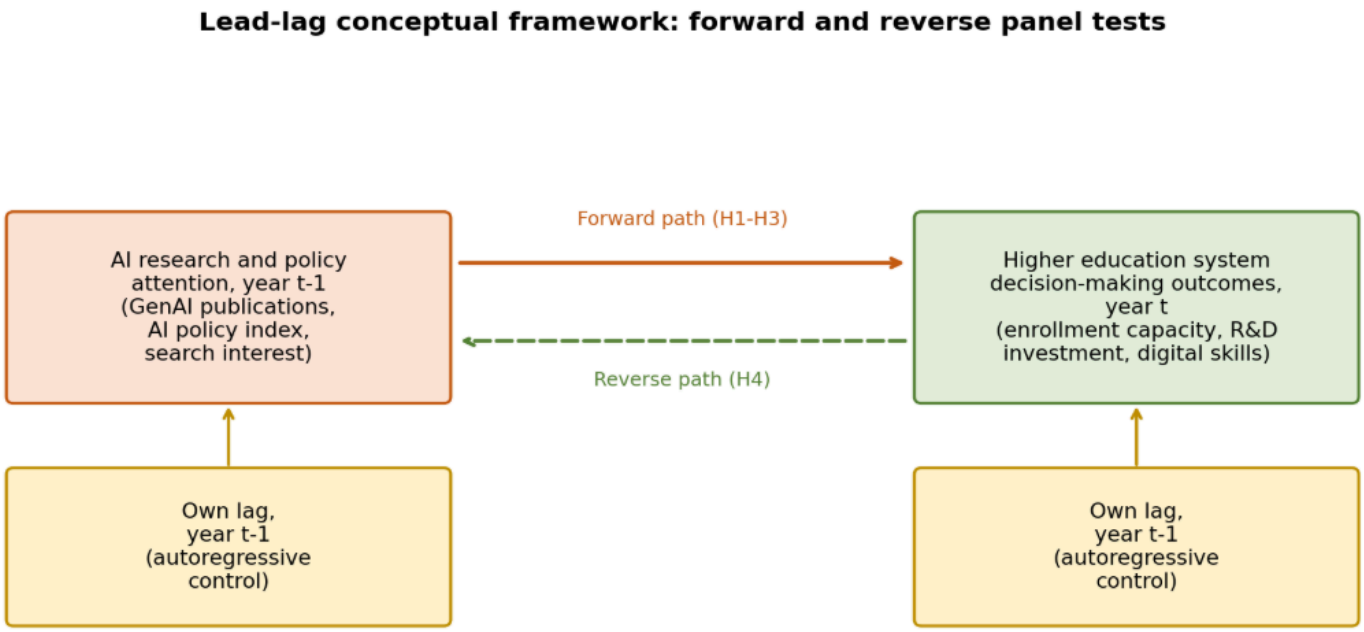


Figure 1

Lead-lag conceptual framework: forward and reverse panel tests.

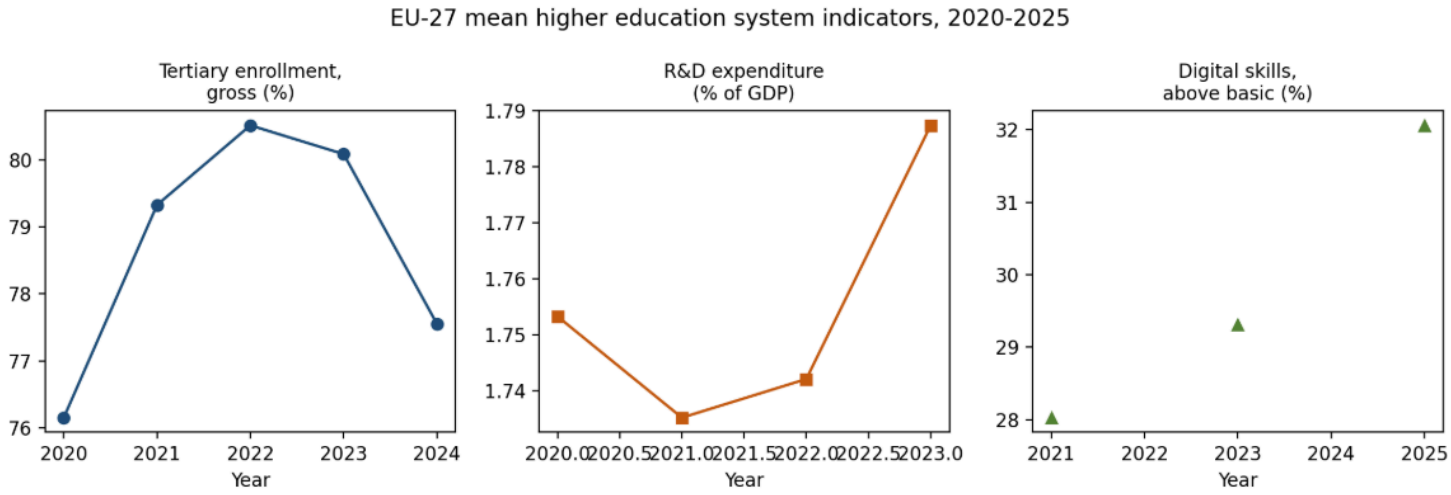


Figure 2

EU-27 mean higher education system indicators, 2020-2025.

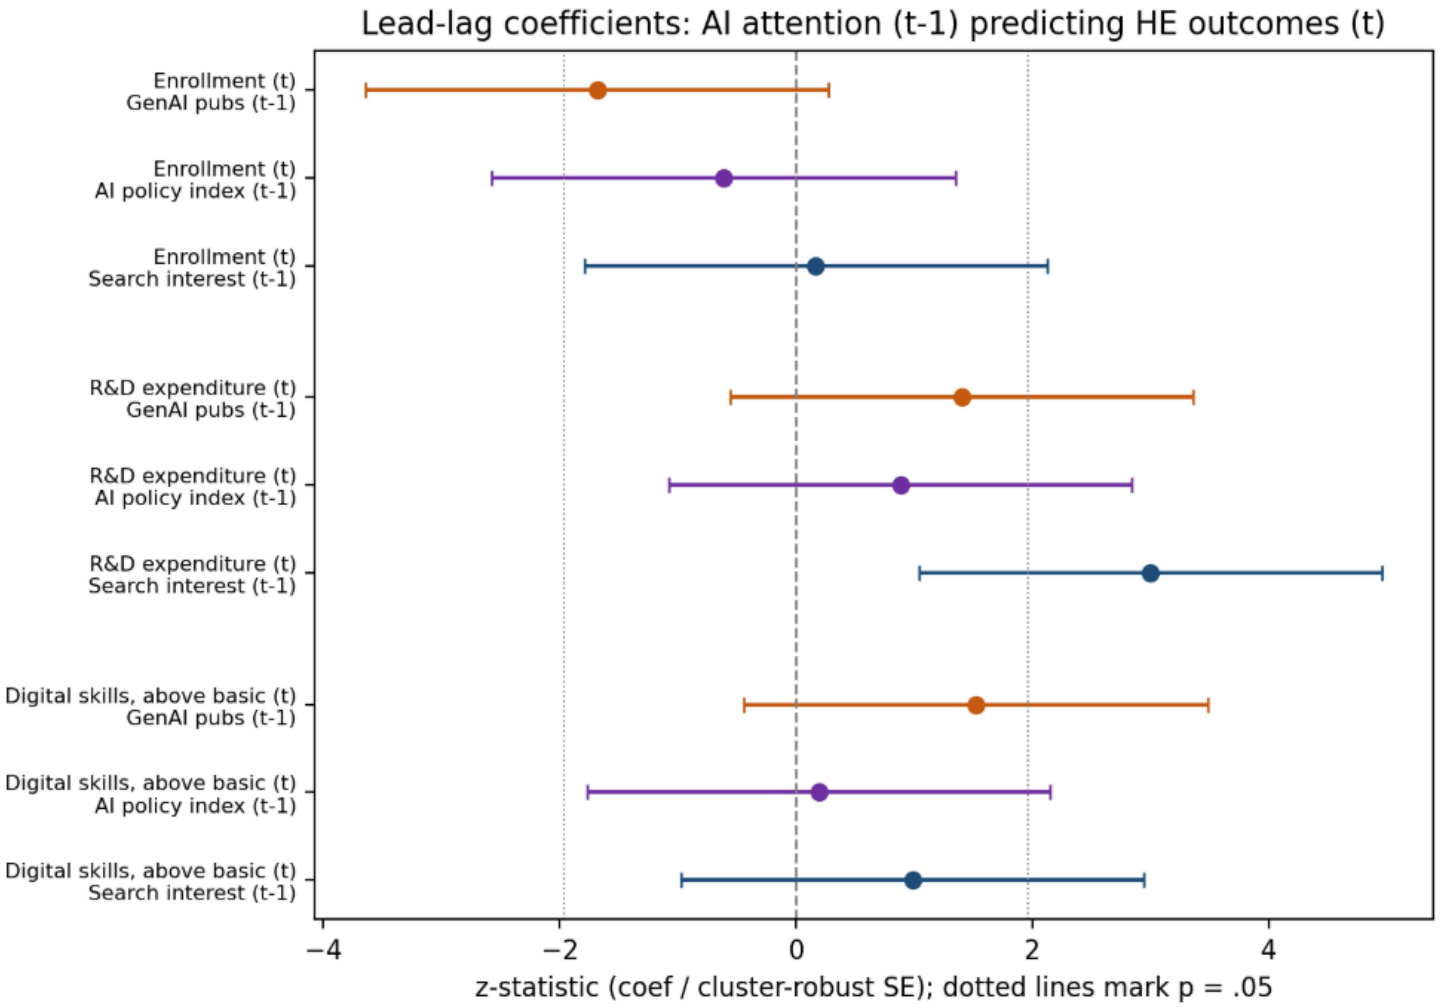


Figure 3

Lead-lag coefficients (standardized z-statistics): AI attention (t-1) predicting higher education outcomes (t).