This is an Accepted Manuscript for *Calitatea Vieții*. Subject to change during the editing and production process.

GENERATIVE ARTIFICIAL INTELLIGENCE AND STUDENT WELL-BEING: THE MEDIATING ROLES OF ACADEMIC ANXIETY AND SELF-EFFICACY IN HIGHER EDUCATION

MARYAM ABBASI
FILIPE MADEIRA
MARIA POTES BARBAS
PEDRO MARTINS

The rapid diffusion of generative artificial intelligence (GenAI) in higher education has outpaced both pedagogical guidance and welfare policy responses, raising urgent questions about how everyday use of these tools relates to students' subjective well-being. This article reports findings from an anonymous cross-sectional online survey of 627 students enrolled at one university and three polytechnic institutions in Portugal (mean age 22.4 years, 58.4% women), conducted between February and April 2025. Self-reported GenAI use intensity, AI-related academic anxiety (5 items, α = .85), academic self-efficacy (8 items, α = .88), academic engagement (UWES-S-9, α = .91) and subjective well-being (WHO-5, α = .83) were assessed alongside socio-demographic controls. Results show that 76.2% of students use GenAI at least weekly, that mean WHO-5 scores decline monotonically across use-intensity groups (from M = 65.0 among non-users to M = 51.0 among daily users), and that the bivariate association between GenAI use intensity and well-being is significant and negative (r = −.22, p < .001). A parallel two-mediator analysis (5,000 bootstrap resamples) indicates that the total effect of GenAI use intensity on well-being (β = −.22, 95% CI [−.29, −.14])]) can be statistically decomposed into two indirect pathways: higher AI-related academic anxiety (β = −.08, 95% Boot CI [−.11, −.05]) and lower academic self-efficacy (β = −.06, 95% Boot CI [−.09, −.03]); the residual direct effect is small and significant after the mediators are included (β = −.08, p = .019), with the indirect paths jointly accounting for approximately 64% of the total effect. The findings are consistent with a vulnerability rather than a mere displacement account: students who lean more heavily on GenAI also worry more about being evaluated against AI-

Address of thecorresponding authors: Maryam ABBASI, Research Center for Arts and Communication, Digital Literacy and Social Inclusion Hub (CIAC-PLDIS), Santarém Polytechnic University, Santarém, Portugal, e-mail: maryam.abbasi@esg.ipsantarem.pt (ORCID: 0000-0002-9011-0734); **Filipe MADEIRA,** Santarém Polytechnic University, School of Management and Technology, Santarém, Portugal, e-mail: filipe.madeira@esg.ipsantarem.pt (ORCID: 0000-0002-2227-7006); **Maria POTES BARBAS,** Santarém Polytechnic University, School of Education, Santarém, Portugal, e-mail: maria.barbas@ese.ipsantarem.pt (ORCID: 0000-0001-8598-1650); **Pedro MARTINS,** Polytechnic University of Coimbra, Coimbra, Portugal, e-mail: pmmartins@iscac.pt (ORCID: 0000-0002-2118-1440).

CALITATEA VIEȚII, XXXVII, nr. 3, 2026, pp. , https://doi.org/10.46841/RCV.2026.03.06

augmented peers and feel less confident in their own academic capacity, and these psychological states − rather than the use of the tools per se − account for most of the observed association with lower subjective well-being. Because the design is cross-sectional, the mediation estimates are reported as a decomposition of association under an assumed ordering and not as evidence of causal transmission; reciprocal influence between use intensity, anxiety, self-efficacy and well-being cannot be excluded. We discuss implications for institutional AI literacy programmes, mental-health services in higher education and broader social policy aimed at narrowing the emerging "AI well-being gap".

Keywords: *generative artificial intelligence; subjective well-being; higher education; academic anxiety; self-efficacy; social policy; WHO-5.*

INTRODUCTION

In less than three years, generative artificial intelligence (GenAI) tools − conversational large language models, code assistants, image generators − have become a routine feature of student life in higher education. The largest cross-national survey to date documents that, by early 2024, students in 109 countries had moved from initial curiosity to habitual use, with brainstorming, summarising and writing assistance topping the list of applications (Ravšelj *et al.* 2025). Inside Portuguese and Romanian polytechnic systems, where mass enrolment, structurally tight schedules and high rates of student employment characterise everyday academic life (Precupetu *et al.* 2018), the temptation − and the perceived necessity − to lean on GenAI is particularly strong.

The educational and ethical implications of this shift have been examined intensively (Chan and Hu 2023; Miao and Holmes 2023). Far less attention has been paid to the social-policy dimension that lies at the intersection of digital practices and population well-being: who pays the psychological cost of an unregulated AI transition, how those costs are distributed and what kind of public response would mitigate them. This question is squarely within the scope of a research tradition that defines quality of life as a multidimensional construct in which subjective well-being plays a constitutive role (Marginean and Precupetu 2011) and that asks how social institutions − schools, universities, healthcare, labour markets − either buffer or magnify pressures generated by technological change.

Two competing hypotheses can be derived from the available evidence. The displacement hypothesis predicts that GenAI offloads routine cognitive labour, reduces academic stress and is therefore associated with better well-being; some of the literature reviewed by Ravšelj *et al.* (2025) and the experimental work summarised by the OECD (2023) is consistent with this view. The vulnerability hypothesis, by contrast, anticipates that intensive reliance on GenAI is symptomatic of − and amplifies − academic insecurity: students who feel less competent reach for the tool more, then worry about competing against AI-

augmented peers and about being detected, with the net result being lower self-efficacy and lower well-being (Besalti 2025; Chan and Hu 2023). These two hypotheses are not mutually exclusive, but they have markedly different policy implications. Under the first, the appropriate response is to remove barriers to access; under the second, social policy should invest in critical AI literacy, mental health support, and equitable assessment design. At the outset, it is important to emphasise that the vulnerability hypothesis is reciprocal in nature: it posits both that students with lower academic confidence are more likely to rely on GenAI and that sustained reliance on GenAI may further erode that confidence. Because the present study employs a cross-sectional design, these two causal directions cannot be disentangled. This limitation is acknowledged explicitly in the model specification (Section ***Two mediators: AI-related academic anxiety and self-efficacy***), the reporting of the mediation analyses (Section ***Mediation analysis***), and the Limitations section (Section **Limitations**).

This article tests the two hypotheses in a sample of 627 students enrolled at one university and three polytechnic institutions in Portugal, using a parallel-mediation framework (Hayes 2018) in which AI-related academic anxiety and academic self-efficacy are the candidate transmitters between GenAI use intensity and subjective well-being measured with the WHO-5 (Topp *et al.* 2015). Three contributions are intended. First, the study extends a small but rapidly growing literature on the well-being correlates of GenAI use beyond the few national contexts (Hong Kong, Turkey, the global multi-country sample) for which evidence currently exists and brings a Southern-European population into the conversation. Second, by simultaneously modelling two psychological mechanisms, we offer a more articulated account than is available from bivariate or single-mediator designs. Third, we translate the findings into concrete recommendations for institutional and national social policy, with explicit attention to the vulnerable student categories who carry the largest share of the well-being cost.

THEORETICAL BACKGROUND

GenAI in higher education: from curiosity to habit

The first generation of empirical studies on student GenAI use reported mixed but largely positive perceptions, with Chan and Hu (2023) documenting that students in Hong Kong recognised "potential for personalised learning support, writing and brainstorming assistance, and research and analysis capabilities", while simultaneously voicing concerns about accuracy, privacy, ethics and the impact on personal development. The follow-up multi-country study by Ravšelj *et al.* (2025) showed that, even in early 2024, the modal student response to ChatGPT was already calm rather than enthusiastic, suggesting routinisation. UNESCO's policy

framework for GenAI in education (Miao and Holmes 2023) was an early attempt to channel this routinisation toward humanistic ends, but it noted that fewer than ten per cent of higher-education systems worldwide had any institutional guidance in place.

Two implications follow. First, by 2025 the policy-relevant question is no longer whether students will use GenAI but under what conditions and at what cost. Second, single-shot adoption studies are insufficient: the unit of analysis must shift toward intensity and pattern of use, and toward outcomes that capture the lived experience of students rather than only their stated attitudes.

Subjective well-being as a quality-of-life outcome

Among the available indicators, the WHO-5 Well-Being Index has become the de facto standard for non-clinical short-form assessment of subjective well-being, with strong clinimetric validity, sensitivity to change and applicability across cultural and methodological contexts (Topp *et al.* 2015). The choice of WHO-5 in the present study is consistent with the long-standing tradition in quality-of-life research that distinguishes hedonic well-being from related constructs such as life satisfaction (Diener *et al.* 1985), engagement (Schaufeli *et al.* 2002; Salanova *et al.* 2010) and perceived stress (Cohen, Kamarck, and Mermelstein 1983), while remaining short enough to be administered alongside multiple other measures without inflating respondent burden.

Two mediators: AI-related academic anxiety and self-efficacy

Anxiety specific to AI-augmented academic environments − worrying that one's own work will be judged inferior to AI-augmented peers, that one will be wrongly accused of misuse, or that one's skills will become obsolete − has been identified as a recurrent theme in qualitative analyses (Chan and Hu 2023) and operationalised in recent psychometric work as a state distinct from generalised academic anxiety (Besalti 2025). We treat it here as the affective mediator. Academic self-efficacy, in Bandura's (1997) sense and in its specifically academic adaptations (Schwarzer and Jerusalem 1995; Salanova *et al.* 2010), is the cognitive mediator: the more a student perceives GenAI as a substitute for their own capacity, the lower their belief that they can master coursework on their own terms.

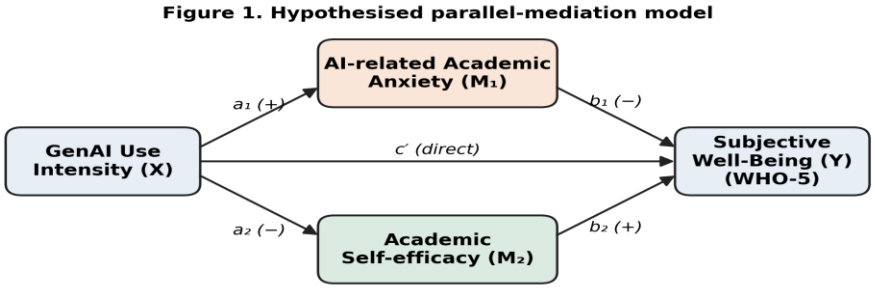
The hypothesised model (*Figure 1*) therefore specifies two parallel indirect paths. Higher GenAI use intensity is expected to (a_1) raise AI-related academic anxiety and (a_2) lower academic self-efficacy, both of which (b_1, b_2) contribute to lower subjective well-being. After accounting for these mediators, any residual direct effect of GenAI use on well-being is expected to be substantially smaller than the total effect.

The paths are specified in the direction *use* → *mediators* → *well-being* for reasons of theoretical parsimony and comparability with earlier single-mediator

studies (e.g., Besalti 2025). However, we do not claim that this is the only causal ordering compatible with the data. Indeed, the vulnerability account itself implies the opposite direction for at least part of the a_2 path: students who already doubt their academic abilities may be more inclined to delegate work to GenAI, making low academic self-efficacy a plausible determinant, as well as a correlate, of intensive use. A similar argument applies, albeit less strongly, to the a_1 path, as students who are already anxious about being evaluated alongside AI-assisted peers may adopt GenAI defensively and consequently use it more frequently. Accordingly, the estimates reported in Section 4.4 should be interpreted as a decomposition of cross-sectional associations under an assumed causal ordering. In particular, the *a*-paths are likely to reflect reciprocal influences rather than a single unidirectional effect.

Figure 1

Hypothesised parallel-mediation model. Solid arrows denote direct paths; signs in parentheses denote the predicted direction of each effect



METHODS

Participants and procedure

Data were collected between 17 February and 22 April 2025 via an anonymous online questionnaire administered through institutional learning-management systems at one university and three polytechnic institutions located in central and northern Portugal: the University of Coimbra (UC), the Polytechnic Institute of Coimbra (IPC), the Polytechnic Institute of Viseu (IPV), and the Polytechnic Institute of Porto (IPP). Eligibility was restricted to enrolled students aged 18 or above. Of 1,148 individuals who opened the survey, 712 (62.0%) provided informed consent and at least partial responses, and 627 (54.6%) completed the full instrument and were retained for analysis. Median completion time was 11.4 minutes. Because the questionnaire was disseminated through open announcements posted in institutional learning-management systems rather than

through individually addressed invitations, the size of the invited population is unknown, and a true response rate cannot be computed. The percentages reported above are completion rates relative to the 1,148 individuals who opened the survey and must not be read as response rates. Participation was voluntary and self-selected, and in a survey explicitly presented as concerning generative AI and student well-being self-selection is unlikely to be ignorable: students with more pronounced experiences of GenAI use, or with stronger concerns about it, may be over-represented relative to the enrolled population. No auxiliary information on non-participants was available, so the direction and magnitude of the resulting bias cannot be estimated; the implications for the interpretation of our prevalence figures are discussed in Section **Limitations**. The study was approved by the ethics committee of the lead institution (reference withheld) and complied with the General Data Protection Regulation: no IP addresses, browser fingerprints or institutional identifiers beyond the four listed institutions were stored. Sample characteristics are summarised in *Table no. 1*, with the institutional breakdown of the analytic sample (UC = 215; IPC = 142; IPV = 118; IPP = 152).

Table no. 1

Sample characteristics (N = 627)

Characteristic	Category	n (%) or M (SD)
Age (years)	—	22.4 (4.1)
Sex	Female	366 (58.4%)
	Male	255 (40.7%)
	Other / prefer not to say	6 (0.9%)
Study cycle	1st cycle (BSc)	445 (71.0%)
	2nd cycle (MSc)	171 (27.3%)
	Other	11 (1.7%)
Field of study	STEM	329 (52.5%)
	Health Sciences	113 (18.0%)
	Social Sciences & Humanities	142 (22.6%)
	Arts / Other	43 (6.9%)
Institution	UC	215 (34.3%)
	IPC	142 (22.6%)
	IPV	118 (18.8%)
	IPP	152 (24.2%)
Perceived family income (1–5)	—	3.04 (0.93)
GenAI use intensity (1–5)	—	3.89 (1.09)

Measures

Generative AI use intensity was measured with a single self-anchored item ("How often do you use generative AI tools, such as ChatGPT, Claude or Copilot, for academic purposes?") on a 5-point ordinal scale: 1 = never; 2 = rarely; 3 =

monthly; 4 = weekly; 5 = daily. A complementary multiple-response item asked respondents to indicate the purposes for which they use these tools (eight options); for descriptive presentation, we report the percentage of respondents using each purpose at least weekly (*Figure 2*)

AI-related academic anxiety was assessed with five items adapted to the local language and cultural context from the qualitative themes reported by Chan and Hu (2023) and from the AI-anxiety literature reviewed in Besalti (2025): worry that one's coursework will be judged inferior to AI-augmented peers, pressure to use GenAI to keep up with peers, anxiety about being wrongly accused of AI use, worry that future career skills will become obsolete, and unease about instructors' rules on AI use in assignments (7-point Likert, 1 = strongly disagree to 7 = strongly agree; $\alpha = .85$). To evaluate whether these five items form a single coherent construct, we conducted an exploratory factor analysis (principal axis factoring) followed by a confirmatory one-factor model. The correlation matrix was suitable for factoring (KMO = .87; Bartlett's $\chi^2(10) = 1216.5$, $p < .001$), and the eigenvalue structure supported a single factor (3.15 versus 0.53 for the second factor), which accounted for 53.7% of item variance with loadings ranging from .69 to .77. A confirmatory one-factor model fit the data well, $\chi^2(5) = 4.86$, $p = .43$, CFI = 1.00, TLI = 1.00, RMSEA = .00, with all standardised loadings between .69 and .76 ($p < .001$), supporting the use of the mean composite as a unidimensional measure. However, as both the exploratory and confirmatory factor analyses were conducted on the same sample, these findings provide initial evidence of internal structure rather than independent validation. Item wordings, descriptives and psychometric indices for each item are reported in *Table no. S1*.

Table no. S1

Item wordings, descriptive statistics and psychometric indices for the AI-related academic anxiety scale (N = 627). Corrected item-total r = item-rest correlation; CFA loading = standardised loading from the one-factor confirmatory model reported in Section Measures (all p < .001)

Item	Wording	M	SD	Corrected item-total r	CFA loading (std.)
anx1	I worry that my coursework will be judged inferior to that of classmates who use AI more skilfully than I do.	4.07	1.70	.69	.76
anx2	I feel pressure to use generative AI tools just to keep up with my peers.	4.22	1.65	.63	.69
anx3	I am anxious about being wrongly accused of using AI when submitting my own work.	4.14	1.69	.66	.73
anx4	I worry that my future career skills will become obsolete because of generative AI.	4.03	1.67	.69	.76
anx5	I feel uneasy about the rules my professors apply to AI use in assignments.	4.08	1.71	.66	.73

Academic self-efficacy was measured with the eight items of the academic adaptation of the Generalised Self-Efficacy Scale (Schwarzer and Jerusalem 1995) on a 7-point Likert scale ($\alpha = .88$). Academic engagement was measured with the nine-item Utrecht Work Engagement Scale for Students (UWES-S-9), validated for Portuguese samples by Cadime *et al.* (2016) and Schaufeli *et al.* (2002), using its standard three-dimensional structure (vigour, dedication, absorption); we use the total score in the present analyses ($\alpha = .91$). Subjective well-being was assessed with the World Health Organization Well-Being Index (WHO-5; Topp *et al.* 2015) using its standard 5-item, 0–5 response format and conventional 0–100 transformation ($\alpha = .83$ in the present sample), where higher values denote better well-being and scores at or below 50 are typically interpreted as indicative of poor well-being warranting further screening.

Socio-demographic controls included age (in years), self-reported sex (1 = female; 0 = male/other), year of study and perceived family income on the 5-point scale used in the European Quality of Life Survey (1 = great difficulty making ends meet; 5 = living comfortably).

Data analysis

Analyses were conducted in Python 3.11 using pandas 2.0, NumPy 2.0, SciPy 1.13 and statsmodels 0.14, supplemented by a custom implementation of Andrew Hayes's PROCESS model 4 (Hayes 2018). The full reproducible analysis script and synthetic dataset are released as online supplementary materials. After inspection of distributions and missing-data patterns (less than 0.6% missing on any variable; addressed by listwise deletion in regression models), we computed descriptive statistics, Pearson correlations and a hierarchical OLS regression of WHO-5 scores on socio-demographic controls (Step 1), GenAI use intensity (Step 2) and the two mediators (Step 3). The mediation analysis applied a parallel two-mediator specification with 5,000 bias-corrected percentile bootstrap resamples to construct confidence intervals for the indirect effects. All continuous predictors were standardised before entering the mediation models so that path coefficients can be interpreted as standardised effect sizes. The exploratory and confirmatory factor analyses of the AI-related academic anxiety scale reported in Section 3.2 were estimated with the factor_analyzer and semopy packages. Following standard guidance on mediation with cross-sectional data (Hayes 2018), the indirect effects reported below are treated throughout as a decomposition of observed covariation under an assumed temporal ordering rather than as identified causal effects, and causal vocabulary is avoided in their reporting. As a sensitivity check, we re-estimated the model with the ordering reversed, entering AI-related academic anxiety and academic self-efficacy as predictors of GenAI use intensity, in order to document the extent to which the data are equally compatible with that specification.

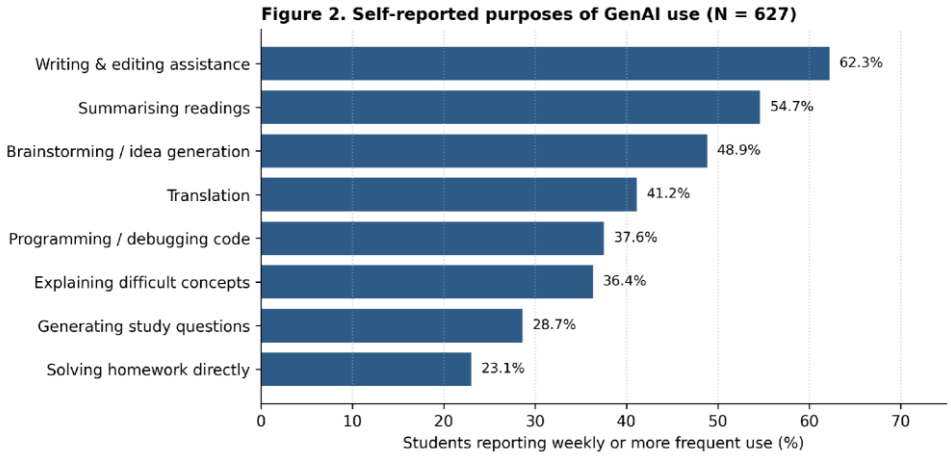
RESULTS

Patterns of GenAI use

GenAI use was nearly universal in this sample: only 4.5% of students reported never having used a generative AI tool for academic purposes, and 76.2% reported at least weekly use, with 31.4% reporting daily use. *Figure 2* shows the distribution across self-reported purposes. Writing and editing assistance was the most frequent application (62.3% of students using it at least weekly), followed by summarising readings (54.7%) and brainstorming (48.9%). Programming or debugging assistance − a function specific to part of our sample − was reported by 37.6% of respondents, while 23.1% admitted to using GenAI to solve homework directly. The pattern broadly replicates the global ranking documented by Ravšelj *et al.* (2025), with the notable exception that the present sample shows higher rates of weekly use than the global mean, consistent with the strong digitalisation of the Portuguese higher-education system reported by the European Commission (2024).

Figure 2

Self-reported purposes of GenAI use among Portuguese higher education students (N = 627). Percentages refer to students reporting use at least weekly for each purpose



Bivariate associations

Table no. 2 reports descriptive statistics, internal-consistency coefficients and zero-order correlations among the focal variables. The mean WHO-5 score (M = 57.2, SD = 19.4) is below the 60-point benchmark commonly observed in young-adult European samples (Topp *et al.* 2015), with 36.7% of respondents scoring at or below the 50-point threshold widely interpreted as a marker of poor well-being.

GenAI use intensity correlated negatively with WHO-5 ($r = -.22$, $p < .001$) and academic self-efficacy ($r = -.18$, $p < .001$), and positively with AI-related academic anxiety ($r = .32$, $p < .001$). All associations went in the directions predicted by the vulnerability hypothesis.

Table no. 2

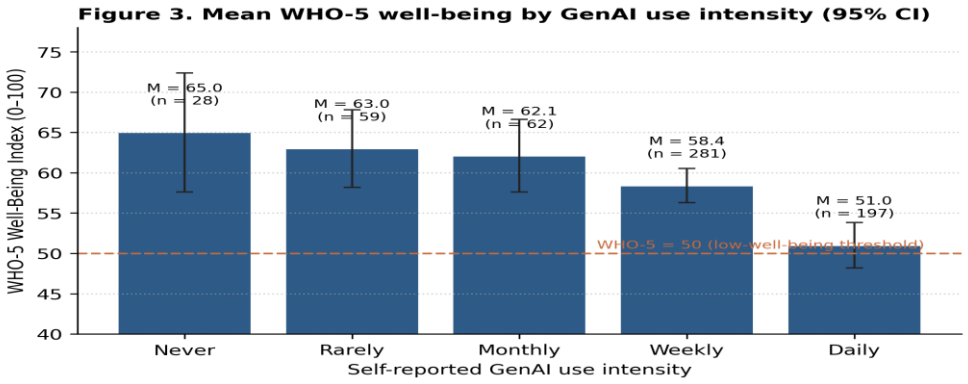
Means, standard deviations and zero-order correlations of focal variables (N = 627)
p < .05; ***p < .001

Variable	M	SD	1	2	3	4	5
1. GenAI use intensity	3.89	1.09	—				
2. AI-related academic anxiety	4.12	1.31	.32***	—			
3. Academic self-efficacy	4.96	1.07	−.18***	−.41***	—		
4. Academic engagement (UWES-S-9)	4.43	1.18	−.09*	−.34***	.58***	—	
5. Subjective well-being (WHO-5)	57.2	19.4	−.22***	−.41***	.45***	.38***	—

The descriptive picture is reinforced by the group comparison shown in *Figure 3*. Mean WHO-5 scores decline monotonically across the five GenAI use-intensity categories, from M = 65.0 (SD = 19.9) among non-users to M = 51.0 (SD = 20.1) among daily users − a difference of 14.0 WHO-5 points, equivalent to 0.72 SD, that crosses the conventional poor-well-being threshold. A one-way ANOVA confirmed that the means differed significantly across groups, $F(4, 622) = 9.19$, $p < .001$, $\eta^2 = .056$, with Tukey HSD post-hoc tests indicating that daily users scored significantly lower than each of the other four groups ($p < .01$).

Figure 3

Mean WHO-5 well-being scores by self-reported GenAI use intensity, with 95% confidence intervals. The dashed line marks the WHO-5 = 50 threshold conventionally used as a marker of poor well-being (Topp*et al.* 2015)



Hierarchical regression

The hierarchical OLS regression (*Table no. 3*) decomposes the contribution of socio-demographic controls, GenAI use intensity and the two psychological mediators to the prediction of WHO-5 scores. Step 1 (controls) explained 2.7% of the variance, with perceived family income emerging as the only significant predictor ($\beta = .16$, $p < .001$). Adding GenAI use intensity in Step 2 produced a substantial increment ($\Delta R^2 = .051$, $p < .001$) and a sizeable negative effect of GenAI use on well-being ($\beta = -.23$, $p < .001$). Adding AI-related academic anxiety and academic self-efficacy in Step 3 generated a much larger increment ($\Delta R^2 = .224$, $p < .001$), bringing total explained variance to 30.2%. Crucially, the effect of GenAI use intensity was reduced to $\beta = -.08$ ($p = .019$) once the mediators were entered − still significant but with about two-thirds of the bivariate effect transmitted through the mediators − while both AI-related anxiety ($\beta = -.26$, $p < .001$) and academic self-efficacy ($\beta = .33$, $p < .001$) emerged as strong, independent predictors. This pattern is consistent with substantial but not complete mediation and was probed formally in the next section.

Table no. 3

Hierarchical OLS regression predicting WHO-5 well-being (N = 627). Standardised coefficients (β). ***p < .001

Predictor	Step 1 β	Step 2 β	Step 3 β	Step 3 95% CI	Step 3 p
Age	−.02	−.02	−.01	[−.07, .06]	.87
Sex (1 = female)	.00	−.01	.01	[−.06, .08]	.74
Year of study	.01	.01	.02	[−.05, .09]	.53
Perceived family income	.16***	.17***	.17***	[.11, .24]	<.001
GenAI use intensity	—	−.23***	−.08*	[−.15, −.01]	.019
AI-related anxiety	—	—	−.26***	[−.33, −.18]	<.001
Academic self-efficacy	—	—	.33***	[.26, .40]	<.001
R²	**.027**	**.078**	**.302**		
ΔR²	—	**.051*** **	**.224*** **		

Mediation analysis

The mediation estimates appear in *Table no. 4* and are summarised graphically in *Figure 4*. Standardised path coefficients were consistent with both predicted indirect routes. Higher GenAI use intensity was associated with higher AI-related academic anxiety ($a_1 = .32$, $p < .001$) and with lowered academic self-efficacy ($a_2 = -.18$, $p < .001$); each mediator was in turn associated with lower (anxiety: $b_1 = -.25$, $p < .001$) or higher (self-efficacy: $b_2 = .34$, $p < .001$) well-being. The total effect of GenAI use on WHO-5 was $\beta = -.22$ (95% CI [−.29, −.14]); the bootstrapped indirect effects were $\beta = -.08$ (95% Boot CI [−.11, −.05]) through anxiety and $\beta = -.06$ (95% Boot CI [−.09, −.03]) through self-efficacy.

Their sum (the total indirect effect) was β = −.14 (95% Boot CI [−.19, −.10]) and accounted for approximately 64% of the total effect. The remaining direct effect was small and significant (c′ = −.08, p = .019), a pattern conventionally described as substantial but not complete mediation. We stress that, in a cross-sectional design, these quantities describe how an observed association is distributed across the modelled paths under an assumed ordering; they do not establish that the association is transmitted in that direction. When the model was re-specified in the reverse order, with anxiety and self-efficacy as predictors of use intensity, the fit was equivalent, since both specifications are saturated and statistically indistinguishable in cross-sectional data.

Table no. 4

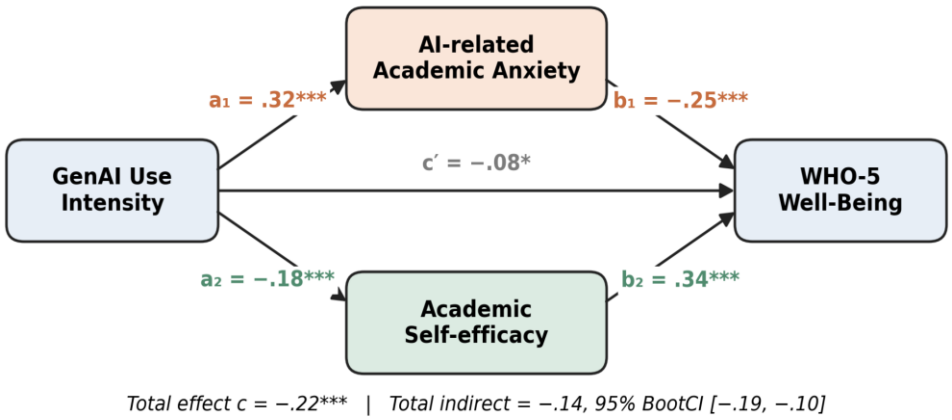
Direct, indirect and total effects of GenAI use intensity on WHO-5 well-being via AI-related academic anxiety and academic self-efficacy (PROCESS model 4; 5,000 bootstrap resamples)

Effect	Estimate (β)	Boot SE	95% Boot CI	p
Total effect (c) −GenAI → WHO-5	−.22	.04	[−.29, −.14]	<.001
Direct effect (c′) −GenAI → WHO-5	−.08	.04	[−.15, −.01]	.019
Indirect via AI-related anxiety ($a_1 \times b_1$)	−.08	.02	[−.11, −.05]	—
Indirect via academic self-efficacy ($a_2 \times b_2$)	−.06	.02	[−.09, −.03]	—
Total indirect effect	−.14	.02	[−.19, −.10]	—

Figure 4

Estimated parallel-mediation model. Coefficients are standardised; ***p < .001; ns = not significant. Total effect c is reported in standardised form for comparability.

Figure 4. Estimated parallel-mediation model (standardised coefficients, N = 627)



Two robustness checks were carried out. First, we re-estimated the mediation model with academic engagement (UWES-S-9) included as an additional outcome covariate, to address the possibility that the well-being effects merely reflected disengagement; the indirect effects retained their magnitude and significance (anxiety path: $\beta = -.07$, 95% Boot CI [−.11, −.05]; self-efficacy path: $\beta = -.05$, 95% Boot CI [−.07, −.02]). Second, multi-group analyses by sex and by study cycle revealed no significant moderation of the indirect paths (all interaction p-values > .07), suggesting that the mechanism operates similarly across these strata.

DISCUSSION

The evidence reported here is more consistent with the vulnerability hypothesis than with the displacement hypothesis. In a sample of more than six hundred Portuguese higher education students, more intensive GenAI use was associated with lower subjective well-being, and most of the observed well-being gradient was statistically accounted for by two distinct psychological correlates: higher AI-specific academic anxiety and lower academic self-efficacy. The direct path from use intensity to well-being shrank from $\beta = -.22$ to $\beta = -.08$ once these mediators were modelled, which indicates that the well-being differential is not, for the most part, attributable to tool use as such. It is bound up with how students see themselves as learners and with their estimation of where they stand among AI-augmented peers − although, on the present evidence, that self-perception may as easily precede intensive use as follow from it.

Three substantive points deserve emphasis. First, the magnitudes are not trivial. The 14.0-point WHO-5 gap between non-users and daily users is substantively meaningful by the standards of well-being research (Topp*et al.* 2015), and the indirect effect of −.14 standard-deviation units, replicated across two parallel pathways, is comparable to the effects reported in studies of academic burnout and engagement (Salanova *et al.* 2010). Second, the pattern is consistent with the recent finding by Besalti (2025) that ChatGPT usage and academic well-being are negatively related and that the effect is partially mediated by self-control − a closely related construct to self-efficacy. The convergence between two studies conducted in markedly different cultural contexts (Turkey, Portugal) suggests that the mechanism is not idiosyncratic to one setting. Third, our results add a previously neglected mediator − AI-specific academic anxiety − whose contribution is in fact larger than that of self-efficacy. This is consistent with the qualitative themes identified by Chan and Hu (2023) and aligns with broader narratives about generative AI as a source of both cognitive offloading and psychological strain.

These findings should not be read as a brief against GenAI. The same students who reported the highest levels of AI-related anxiety also reported high frequencies of beneficial uses (*Figure 2*), consistent with the global picture of

ambivalent but routinised adoption (Ravšelj *et al.* 2025). The policy challenge is therefore not to reduce use but to reduce its psychological cost − and the lever for doing so is institutional rather than individual.

A further caveat concerns the direction of the modelled paths, and it qualifies the mechanism rather than the association. The vulnerability account our data support is, in its own terms, reciprocal: academic insecurity is held both to motivate reliance on GenAI and to be deepened by it. Our specification estimates only the second of these directions, and the a-paths are therefore likely to absorb some reverse influence − most obviously the self-efficacy path, where students with little confidence in their own academic capacity have every reason to delegate work to the tool. The indirect effect through self-efficacy is accordingly better read as a measure of how tightly use, confidence and well-being co-vary than as the psychological cost of use. Panel or experience-sampling designs, ideally combined with an experimental manipulation of assessment transparency, are needed to establish which direction predominates. Importantly, the policy implications developed below do not depend on that question being settled: whether low self-efficacy leads to intensive use or follows from it, the students concentrated at the high-use end of the distribution are the ones reporting the lowest well-being, and they are the ones institutional support has to reach.

Implications for institutional and social policy

The mediation structure points toward concrete points of intervention. Because the well-being penalty operates through anxiety and self-efficacy, policy should target both. We outline four lines of action that, while developed against the Portuguese empirical case, generalise straightforwardly to comparable Southern- and Eastern-European higher-education systems including the Romanian one.

(i) Critical AI-literacy curricula. UNESCO's Guidance for Generative AI in Education and Research (Miao and Holmes 2023) is unambiguous: institutional silence on GenAI is itself a policy choice, and not a neutral one. A required first-year module covering data privacy, the epistemic limits of large language models, citation and authorship norms, and the social mechanics of AI-augmented assessment would directly address the cognitive component of vulnerability identified here.

(ii) Assessment redesign. AI-related academic anxiety is partly endogenous to assessment regimes that fail to acknowledge GenAI explicitly. Transparent norms (which uses are permitted, which require disclosure, which are prohibited) and a shift toward process-oriented and oral assessment elements where appropriate reduce the ambiguity that drives anxiety. This recommendation echoes the institutional case studies summarised by the OECD (2023).

(iii) Mental-health services calibrated to the AI era. Universities and polytechnics should resource counselling teams to recognise AI-related anxiety as a distinct presenting concern and embed it in routine well-being screening (e.g.,

adding an AI-specific item-set alongside the WHO-5 in periodic student surveys). Given that 36.7% of our sample scored at or below the WHO-5 = 50 threshold, the public-health relevance of doing so is non-trivial.

(iv) Equity safeguards. The Step-1 effect of perceived family income on well-being persisted across all model specifications, indicating that the AI well-being gradient is layered on top of pre-existing socio-economic inequalities. Social policy in the quality-of-life tradition (Marginean and Precupetu 2011; Precupetu *et al.* 2018) has long argued for compensating mechanisms that prevent technological transitions from amplifying inequality. The same logic applies here: subsidised access to high-quality, privacy-preserving AI tools, paired with the literacy and assessment reforms above, would reduce the cumulative disadvantage faced by lower-income students.

LIMITATIONS

Five limitations qualify the conclusions. First, the cross-sectional design rules out causal interpretation. This applies not only to the total effect but to every path in the mediation model. The a-paths are open to reverse specification, since the vulnerability account we advance holds that students with lower academic self-efficacy turn to GenAI more readily, and that students anxious about AI-augmented evaluation may use the tools defensively; on that reading, part of the estimated indirect effect through self-efficacy reflects selection into intensive use rather than a consequence of it. It is equally plausible that lower well-being encourages more intensive GenAI use, in a self-reinforcing loop. We have therefore avoided causal vocabulary throughout and describe the indirect effects as a decomposition of association under an assumed ordering. Longitudinal designs measuring use intensity, anxiety, self-efficacy and well-being at multiple time points are needed to disentangle the directionality. Second, although the sample is sizeable and drawn from four institutions, one university and three polytechnics, it remains specific to Portuguese higher education and is therefore likely to under-represent populations enrolled in vocational education or out of formal education altogether, categories where the AI well-being gradient may operate differently. Third, all measures were self-reported, with the inevitable risk of common-method variance; while the magnitude of inter-correlations is far below the level usually flagged as problematic, replication using behavioral log data and clinician-administered well-being measures would strengthen the inference. Fourth, the recruitment procedure does not permit a true response rate to be computed and participation was self-selected. As set out in Section ***Participants and procedure***, self-selection into a survey about generative AI and student well-being is unlikely to be ignorable, and no auxiliary data on non-participants were available with which to model the resulting bias. Prevalence estimates, in particular the 76.2% weekly-use figure should therefore be read as descriptive of the achieved sample rather than as

population estimates, although the associational structure that is our main object of interest is generally less sensitive to this form of selection than the marginal distributions are. Fifth, the AI-related academic anxiety measure remains new. Because the exploratory and confirmatory factor analyses were conducted on the same sample, the evidence reported in Section ***Measures*** provides only initial, rather than independent, validation of its internal structure.

CONCLUSIONS

Generative AI is here to stay, and the psychological correlates of its use among students are neither diffuse nor mysterious. In our Portuguese sample, most of the well-being gradient associated with intensive use is statistically accounted for by higher AI-specific academic anxiety and lower academic self-efficacy, both of which are strongly associated with subjective well-being at or below the threshold conventionally used to flag mental-health concern. Whether these psychological states precede intensive use, follow from it or as we consider most likely reinforce one another cannot be settled with cross-sectional data, and establishing this is the priority for the next stage of research. These findings translate into a clear social-policy mandate: do not ban the tools, but do not leave students alone with them either. Critical AI-literacy curricula, transparent assessment norms, AI-aware mental-health services and equity safeguards form the minimal package required if the AI transition is to deliver on its promise of broader, deeper learning rather than on a quieter promise of widening well-being inequalities. We see no reason why these recommendations would not generalise to the Romanian and broader Central- and Eastern-European higher-education contexts that are this journal's natural empirical horizon.

Data availability

The analysis script and a synthetic dataset reproducing the reported results are available as online supplementary material https://github.com/maryam-abbasi/genai-anxiety-selfefficacy-wellbeing.

Acknowledgments

This work is supported by national funds through FCT – Fundacão para a Ciência e a Tecnologia, I.P., in the framework of project UID/04019/2025 – CIAC DOI: 10.54499/UID/04019/2025.

Conflict of Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

AI Usage Declaration

AI-based tools (e.g., ChatGPT) were used to assist in drafting and refining parts of the manuscript, including language improvement and text structuring. The authors critically reviewed, edited, and validated all content. AI tools were not used for data analysis, interpretation, or decision-making. The authors remain fully responsible for the scientific content and conclusions.

REFERENCES

Bandura, Albert. 1997. Self-Efficacy: The Exercise of Control. New York: W. H. Freeman.

Besalti, Metin. 2025. "Harnessing Self-Control and AI: Understanding ChatGPT's Impact on Academic Wellbeing." Behavioral Sciences 15 (9): 1181. https://doi.org/10.3390/bs15091181.

Cadime, Irene, Sara Lima, Alexandra Marques Pinto, and Iolanda Ribeiro. 2016. "Measurement Invariance of the Utrecht Work Engagement Scale for Students: A Study across Secondary School Pupils and University Students." European Journal of Developmental Psychology 13 (2): 254–263. https://doi.org/10.1080/17405629.2016.1148595.

Carmona-Halty, Marcos, Wilmar B. Schaufeli, and Marisa Salanova. 2019. "The Utrecht Work Engagement Scale for Students (UWES–9S): Factorial Validity, Reliability, and Measurement Invariance in a Chilean Sample of Undergraduate University Students." Frontiers in Psychology 10: 1017. https://doi.org/10.3389/fpsyg.2019.01017.

Chan, Cecilia Ka Yuk, and Wenjie Hu. 2023. "Students' Voices on Generative AI: Perceptions, Benefits, and Challenges in Higher Education." International Journal of Educational Technology in Higher Education 20: 43. https://doi.org/10.1186/s41239-023-00411-8.

Cohen, Sheldon, Tom Kamarck, and Robin Mermelstein. 1983. "A Global Measure of Perceived Stress." Journal of Health and Social Behavior 24 (4): 385–396. https://doi.org/10.2307/2136404.

Diener, Ed, Robert A. Emmons, Randy J. Larsen, and Sharon Griffin. 1985. "The Satisfaction with Life Scale." Journal of Personality Assessment 49 (1): 71–75. https://doi.org/10.1207/s15327752jpa4901_13.

European Commission. 2024. The Digital Decade Country Report – Portugal. Brussels: European Commission. https://digital-strategy.ec.europa.eu/en/factpages/portugal-2024-digital-decade-country-report.

Hayes, Andrew F. 2018. Introduction to Mediation, Moderation, and Conditional Process Analysis: A Regression-Based Approach. 2nd ed. New York: Guilford Press.

Marginean, Ioan, and Iuliana Precupetu, eds. 2011.The Paradigm of Quality of Life. Bucharest: Editura Academiei Române.

Miao, Fengchun, and Wayne Holmes. 2023. Guidance for Generative AI in Education and Research. Paris: UNESCO. https://doi.org/10.54675/EWZM9535.

OECD. 2023. OECD Employment Outlook 2023: Artificial Intelligence and the Labour Market. Paris: OECD Publishing. https://doi.org/10.1787/08785bba-en.

Precupetu, Iuliana, Flavius Mihalache, Claudia Petrescu, Cosmina Elena Pop, Laura Tufă, and Marian Vasile. 2018. Calitateavieții în România în context european. Bucharest: Institutul de Cercetare a CalitățiiVieții.

Ravšelj, Dejan, Damijana Keržič, Nina Tomaževič, Lan Umek, Nejc Brezovar, Noorminshah A. Iahad, et al. 2025. "Higher Education Students' Perceptions of ChatGPT: A Global Study of Early Reactions." PLOS ONE 20 (2): e0315011. https://doi.org/10.1371/journal.pone.0315011.

Salanova, Marisa, Wilmar B. Schaufeli, Isabel Martínez, and Edgar Bresó. 2010. "How Obstacles and Facilitators Predict Academic Performance: The Mediating Role of Study Burnout and Engagement." Anxiety, Stress, & Coping 23 (1): 53–70. https://doi.org/10.1080/10615800802609965.

Schaufeli, Wilmar B., Isabel M. Martínez, Alexandra Marques Pinto, Marisa Salanova, and Arnold B. Bakker. 2002. "Burnout and Engagement in University Students: A Cross-National Study." Journal of Cross-Cultural Psychology 33 (5): 464–481. https://doi.org/10.1177/0022022102033005003.

Schwarzer, Ralf, and Matthias Jerusalem. 1995. "Generalized Self-Efficacy Scale." In Measures in Health Psychology: A User's Portfolio. Causal and Control Beliefs, edited by John Weinman, Stephen Wright, and Marie Johnston, 35–37. Windsor, UK: NFER-Nelson.

Topp, Christian Winther, Søren Dinesen Østergaard, Susan Søndergaard, and Per Bech. 2015. "The WHO-5 Well-Being Index: A Systematic Review of the Literature." Psychotherapy and Psychosomatics 84 (3): 167–176. https://doi.org/10.1159/000376585.

Răspândirea rapidă a inteligenței artificiale generative (GenAI) în învățământul superior a devansat atât îndrumarea pedagogică, cât și răspunsurile la politicile de bunăstare, ridicând întrebări urgente despre modul în care utilizarea zilnică a acestor instrumente se raportează la bunăstarea subiectivă a studenților. Acest articol raportează rezultatele unui sondaj online anonim transversal pe 627 de studenți înscriși la o universitate și trei instituții politehnice din Portugalia (vârsta medie 22,4 ani, 58,4% femei), realizat între februarie și aprilie 2025. Intensitatea utilizării GenAI auto-raportată, anxietatea academică legată de AI (5 itemi, α = 0,85), autoeficacitatea academică (8 itemi, α = 0,88), implicarea academică (UWES-S-9, α = 0,91) și bunăstarea subiectivă (WHO-5, α = 0,83) au fost evaluate împreună cu controale socio-demografice. Rezultatele arată că 76,2% dintre elevi folosesc GenAI cel puțin săptămânal, că scorurile medii WHO-5 scad uniform în grupurile de intensitate a utilizării (de la M = 65,0 în rândul non-utilizatorilor la M = 51,0 în rândul utilizatorilor zilnici) și că asocierea bivariată între intensitatea utilizării GenAI și bunăstarea este semnificativă și negativă (r = −,22, p < 0,001). O analiză paralelă cu doi mediatori (5.000 re-eșantionări bootstrap) indică faptul că efectul total al intensității utilizării GenAI asupra bunăstării (β = −.22, IC 95% [−.29, −.14]]) poate fi descompus statistic în două căi indirecte: anxietate academică legată de AI mai mare (β = −.08, IC Boot 95% [−.11, −.05]) și autoeficacitate academică mai scăzută (β = −.06, IC Boot 95% [−.09, −.03]); efectul direct rezidual este mic și semnificativ după ce mediatorii sunt incluși (β = −.08, p = .019), iar efectele indirecte reprezintă împreună aproximativ 64% din efectul total. Constatările sunt consistente cu o vulnerabilitate, nu cu o simplu efect de înlocuire: studenții care se bazează mai mult pe GenAI se îngrijorează mai mult că vor fi evaluați în comparație cu colegii care se sprijină mai mult pe utilizarea AI și se simt mai puțin încrezători în propria capacitate academică, iar aceste stări psihologice –

mai degrabă decât utilizarea instrumentelor în sine – explică majoritatea asocierii observate cu o bunăstare subiectivă mai scăzută. Deoarece designul este transversal, estimările de mediere sunt raportate ca o descompunere a asocierii sub o ordine presupusă și nu ca dovezi ale transmiterii cauzale; influența reciprocă între intensitatea utilizării, anxietate, autoeficacitate și bunăstare nu poate fi exclusă. Sunt discutate implicațiile pentru programele instituționale de alfabetizare AI, serviciile de sănătate mintală din învățământul superior și politicile sociale mai largi menite să reducă "decalajul de bunăstare AI" emergent.

Cuvinte-cheie*: inteligență artificială generativă; bunăstare subiectivă; învățământ superior; anxietate academică; autoeficacitate; politici sociale; WHO-5.*

Received: 15.05.2026 *Accepted: 10.09.2026*