Cardiff Economics Working Papers

Working Paper No. E2026/14

Avoidant AI Use and Career Exploration Among University Students: The Curvilinear Moderating Role of AI-Shock Awareness

Bo Zhang, Yizhu Wang, Peng Zhou

November 2026

ISSN 1749-6010

Cardiff Business School
Cardiff University
Colum Drive
Cardiff CF10 3EU
United Kingdom
t: +44 (0)29 2087 4000
f: +44 (0)29 2087 4419
business.cardiff.ac.uk

This working paper is produced for discussion purpose. These working papers are expected to be published in due course, in revised form.
Cardiff Economics Working Papers are available online from:
http://econpapers.repec.org/paper/cdfwpaper/ and
https://carbsecon.com/wp/index_abstracts.html
Enquiries: EconWP@cardiff.ac.uk

Article

Avoidant AI Use and Career Exploration Among University Students: The Curvilinear Moderating Role of AI-Shock Awareness

Bo Zhang [1], Yizhu Wang [1] and Peng Zhou [2,*]

[1] Management School, Capital Normal University, Haidian District, Beijing 100086, China; bozhang@cnu.edu.cn (B.Z.); wyz0404@hotmail.com (Y.W.)
[2] Cardiff Business School, Cardiff University, Cardiff CF10 3EU, UK
* Correspondence: zhoup1@cardiff.ac.uk

Abstract

Some university students choose not to use artificial intelligence (AI) tools, yet research on students and AI has focused largely on adoption and has treated avoidance as maladaptive. That view cannot explain when avoiding AI might support, rather than undermine, student development. Drawing on transactional coping theory, threat-appraisal reasoning, and self-regulation theory, this study examined whether avoidant AI use is associated with career exploration and whether AI-shock awareness moderates that association linearly or curvilinearly. We hypothesized a positive main effect, a negative linear moderation, and a U-shaped curvilinear moderation. Survey data from 289 students at multiple Chinese universities (recruited via snowball sampling; 289 valid cases from 314 returned questionnaires) were analyzed with three-step hierarchical regression and bias-corrected and accelerated bootstrap confidence intervals (5000 resamples). Avoidant AI use was positively associated with both environment exploration and self-exploration. The linear interaction was negative in sign, but unreliable for environment exploration and only marginally supported for self-exploration. The quadratic interaction was significant for both outcomes: the association held at low awareness for both outcomes and at high awareness for self-exploration, but not at moderate awareness. For environment exploration, the high-awareness rebound rested on a small subgroup and is treated as preliminary. These results are consistent with the view that the developmental consequences of avoiding AI depend on the self-regulatory cost of maintaining that stance, a mechanism we treat as inferred rather than directly measured, and they carry implications for career education and AI-literacy instruction in higher education.

Keywords: technology avoidance; career exploration; AI-shock awareness; self-regulation; curvilinear moderation; higher education

1. Introduction

Artificial intelligence tools have become widely available in higher education. A global survey of over 23,000 students across 109 countries found that most students had used generative AI for academic purposes, from brainstorming to drafting assignments (Ravselj et al., 2025). Reviews similarly document AI's rapid spread through assessment, advising, and learning support in universities (Crompton & Burke, 2023), yet not every student embraces these tools. Some choose not to use them. They decline chatbot assistance, write without AI support, and gather information through traditional channels

(Chan & Hu, 2023). This pattern is not a marginal curiosity. In classrooms where AI is increasingly embedded in learning platforms, advising systems, and career services (Crompton & Burke, 2023), students who opt out represent a substantial and poorly understood segment. The stakes reach beyond study habits. AI is reshaping not only how students learn, but how they prepare for work. Career information, job-matching platforms, and application tools are themselves becoming AI-mediated, so a student's stance toward AI now forms part of their career preparation environment. For higher education, then, understanding students who avoid AI is not a niche concern. It is a prerequisite for designing career education that reaches all students rather than only willing adopters. The question motivating this study is direct. What does avoiding AI mean for students' development, and specifically for their career exploration, the information-gathering and self-reflection through which students prepare for the transition from university to work (Stumpf et al., 1983)?

The conversation about students and AI has been dominated by an adoption-centered perspective. Research on technology acceptance has consistently linked successful adaptation to higher usage, greater digital competence, and willingness to engage with AI tools (Ravselj et al., 2025). The implicit logic is straightforward: the technology is valuable, so the goal is to get students to use it. This lens has shaped institutional responses, from AI-literacy curricula to career-platform integration, all organized around increasing uptake. However, adoption is only half the picture. When students are navigating a technology that could transform their future occupations and diminish their professional standing, they do not simply decide to use it or not. They regulate how they engage with it, testing, stepping back, deliberating, and sometimes deliberately avoiding it as part of a broader process we call AI engagement regulation. From this perspective, avoiding AI is not necessarily a failure to adapt. It can be a deliberate stance that reflects how students are making sense of AI's place in their careers.

The prevailing view in the literature would diagnose such avoidance as maladaptive. In the information systems domain, avoidance has been linked to underutilization and declining job performance (Bala & Venkatesh, 2016; Beaudry & Pinsonneault, 2005). In educational contexts, technology non-use is typically framed as a deficit to be remedied rather than a behavior to be understood (Satchell & Dourish, 2009). Using AI is not cost-free; it demands cognitive effort to formulate prompts, evaluate outputs, and integrate results, and it can provoke anxiety about one's own skills and future prospects (Greenhalgh & Rosenblatt, 1984; Tlili et al., 2023). For some students, avoidance may be the psychologically cheaper stance. Whether avoidance is in fact the cheaper stance, and under which appraisal conditions, is the open question this study addresses. In doing so, we distinguish two forms of career exploration, outward scanning of occupational opportunities (environment exploration) and inward reflection on interests and values (self-exploration), and we ask how their association with avoidance varies across levels of AI-shock awareness, the perception that AI threatens one's career future.

This pattern is difficult to explain with existing frameworks. Coping theory holds that avoidance is neither inherently adaptive nor inherently maladaptive; its consequences depend on context (Lazarus & Folkman, 1984). That principle, however, does not specify which features of the context make avoidance of a beneficial technology more or less costly. Career construction theory identifies exploration as a central adaptive behavior (Savickas, 2005), but has not considered technology avoidance among its antecedents. The technology non-use literature argues that non-use can be deliberate, meaningful, and productive (Satchell & Dourish, 2009), yet offers no account of when it is productive and when it merely drains the user. What is missing is a framework that specifies when avoiding a

technology is psychologically cheap and when it is expensive, and how that cost shapes the resources available for exploration.

We develop such a framework by combining three theoretical lenses. Transactional coping theory supplies the construct of avoidant AI use and the contingency principle that its consequences depend on appraisal context. Threat-appraisal reasoning generates a linear prediction. Because avoidance becomes increasingly emotion-focused as awareness rises, the avoidance-exploration link should weaken with awareness. Self-regulation theory (Carver & Scheier, 1981, 1998) supplies the deeper mechanism. Maintaining an avoidant stance consumes self-regulatory resources in proportion to the internal conflict it generates, and we argue that this conflict follows an inverted-U across awareness. On this account, habitual avoidance at low awareness and resolved, committed avoidance at high awareness should be relatively cheap, whereas contested avoidance at moderate awareness should be expensive. Because exploration draws on the same resource pool, the avoidance-exploration link should be U-shaped across awareness. Together, these lenses reframe what it means for students to manage their relationship with AI in an era of uncertainty about which skills and careers remain viable.

This study makes three contributions. First, it reframes technology avoidance from a uniformly maladaptive behavior to a context-dependent one whose developmental consequences are governed by self-regulatory demand. Second, it proposes that this demand follows an inverted-U across AI-shock awareness and derives from it a U-shaped, rather than linear, moderation of the avoidance-exploration link. Third, it extends career construction theory by identifying a technology-domain behavior, namely, how students position themselves toward AI, as an antecedent of career exploration. It further articulates the shared-resource logic through which adapting behavior in one domain finances or constrains adaptation in another. In doing so, this study connects the technology adaptation literature to career development research and offers a differentiated basis for career education and AI-literacy instruction in higher education.

2. Theoretical Framework and Hypotheses

2.1. Theoretical Foundations

Three bodies of theory inform our model. The first is the transactional theory of stress and coping (Lazarus & Folkman, 1984). It holds that people appraise events in two steps. First, they judge whether an event is irrelevant, benign, or threatening. Then they judge what resources they have to manage it. Coping responses come in two kinds. Problem-focused coping aims to change the situation. Emotion-focused coping aims to manage the distress the situation causes. The theory's central claim for us is simple. No coping response is always adaptive or always maladaptive. Avoidance is an emotion-focused strategy. It can preserve functioning when direct action is impossible or premature. It can also undermine functioning when it blocks necessary engagement. Decades of research confirm this. The same strategy predicts well-being in one context and distress in another. The field has therefore moved from listing strategies to asking when each one works. Adaptiveness depends on the fit between the response and the appraisal context, not on the response itself. This maps directly onto our question. Deliberately avoiding AI is an emotion-focused response to perceived technological disruption. The theory tells us to ask not whether such avoidance helps or harms in general, but under which appraisal conditions it does each.

The second body of theory carries this logic into the technology domain. Beaudry and Pinsonneault (2005) developed the coping model of user adaptation (CMUA). They showed that employees respond to a new technology in one of four ways. The response depends on whether they see the technology as an opportunity or a threat, and on how

much control they feel over the situation. Two of these responses are problem-focused. They aim to exploit the technology's benefits. The other two are emotion-focused. They aim to restore equilibrium. The one that matters most here is self-preservation, withdrawing from the technology to reduce disturbance. Bala and Venkatesh (2016) embedded these responses in a wider network of causes and outcomes. They defined avoidance as "the degree to which an employee tries not to use an IT when accomplishing his or her tasks" (Bala & Venkatesh, 2016, Section 2.2.4), an emotion-focused, self-preservation response to an appraised threat. Wang et al. (2025) translated this framework to AI. They showed that AI-shock awareness, the perception that AI threatens one's career future, is linked to avoidant AI use through threat appraisal. This lineage does two things for our model. It shows that avoidant AI use is a coping behavior grounded in theory, not a casual preference. It also points to a directional expectation about how avoidance should vary with awareness, which we develop in Section 2.4.

The third body of theory is self-regulation theory (Carver & Scheier, 1981, 1998). It holds that goal-directed behavior runs on a limited capacity for deliberate control. Maintaining any behavioral stance draws on this capacity. How much it draws depends on the internal conflict the stance creates. Habitual behaviors run with little effort. Behaviors backed by a resolved commitment become self-sustaining once the decision is made. Behaviors that stay contested drain the most. When someone is torn between acting and not acting, they enter what Gollwitzer (1990) called an action crisis. This unresolved conflict consumes cognitive and emotional capacity over and over. Research on ego depletion shows the consequence. Sustained self-control leaves fewer resources for later goal-directed activity (Baumeister et al., 1998). Conservation-of-resources theory makes the same point. People try to protect finite psychological resources and suffer when those resources are threatened with loss (Hobfoll, 1989). This resource logic supports both the main-effect prediction and the curvilinear moderation we develop below.

The three lenses play distinct roles. Coping theory defines what avoidant AI use is and insists that its consequences are conditional. Threat-appraisal reasoning specifies one direction that conditionality could take. It predicts a linear weakening of avoidance's consequences as awareness grows. Self-regulation theory specifies the mechanism. It identifies the variable cost of maintaining an avoidant stance, which generates both the main effect and a curvilinear alternative to the linear account. Together, the lenses point to a broader question. The question is not whether students should adopt AI. It is how they regulate their engagement with it, what we call AI engagement regulation, and how this regulation shapes their development. We now define the two focal constructs and then derive the hypotheses.

2.2. Avoidant AI Use and Career Exploration

Avoidant AI use is the degree to which a person tries not to use AI when completing learning or career tasks. Bala and Venkatesh (2016) defined technology avoidance as effortful withdrawal from an available and relevant technology. Wang et al. (2025) adapted the construct to AI. Our measure captures active withdrawal. It asks about trying to perform tasks without AI, looking for ways to avoid relying on it, and staying away from it as much as possible. One boundary deserves honesty. Avoidance can arise from different sources. In Bala and Venkatesh's (2016) model, avoidance follows threat appraisal, but it is also stronger when people perceive little control over the technology, including limited ability or resources. Non-use can therefore reflect motivated withdrawal or perceived incapacity, and our items, like the parent scale, measure the behavioral act of withdrawal rather than its motive. We therefore describe avoidant AI use as active withdrawal rather than deliberate non-use, and we return to this boundary in Section 5.3. Two further

boundaries matter. First, avoidant AI use is not a personality trait. A student who avoids AI may be highly proactive elsewhere, and avoidance in one domain can coexist with engagement in another. Second, it is domain-specific. Its consequences should be judged against outcomes that draw on shared psychological resources, not against abstract ideas about adjustment.

Career exploration is the process of gathering information about the world of work and about oneself in relation to it (Stumpf et al., 1983). In career construction theory, it is the core adapting behavior. Through it, people turn concern about their vocational future into concrete preparation for transitions (Savickas, 2005, 2013). Exploration predicts many valued outcomes, from decidedness to person–environment fit to employment quality. This is why its antecedents have drawn sustained attention in vocational and social-cognitive career research (Lent et al., 1994). We treat environment exploration and self-exploration as two independent constructs, not two halves of a single factor. Environment exploration is directed outward, toward occupations, employers, and labor-market opportunities. Self-exploration is directed inward, toward interests, abilities, values, and possible selves. The two are correlated but distinct. They rely on different cognitive and emotional operations, and may respond differently to the same antecedent. We verify their separability in Section 4.3.

2.3. The Resource Logic of Avoidance: Main Effect

Career exploration is resource-dependent. Scanning occupational information, comparing alternatives, and tolerating uncertainty about the future require sustained attention and cognitive effort. Reflecting on one's interests and limitations requires introspective capacity and the willingness to confront uncomfortable possibilities. These are effortful activities, not passive states. They draw on the same finite pool of self-regulatory resources that governs other goal-directed behavior (Carver & Scheier, 1998). This view is consistent with evidence that exploration contracts under cognitive load and emotional exhaustion. Students under strain explore less, not because their motivation disappears, but because their capacity does (Hobfoll, 1989).

Engaging with AI draws on the same pool. Effective use requires formulating prompts, evaluating outputs for accuracy and relevance, and integrating results into one's own work. These activities impose a real cognitive load, especially for users who are still learning the technology. For students who see AI as consequential for their careers, use can also carry an emotional load. It means monitoring one's standing relative to the technology and managing the anxiety that this generates (Greenhalgh & Rosenblatt, 1984). Avoiding AI sidesteps these expenditures. A student who does not use AI pays neither the cognitive costs of operating it nor the emotional costs of engaging with a technology framed as a threat. On average, then, avoidance should conserve resources that remain available for allocation elsewhere.

Career exploration is a natural beneficiary of whatever capacity is conserved, for two reasons. It is instrumental to the very future that makes AI worth worrying about, so conserved capacity is likely to flow toward it rather than toward unrelated pursuits. Unlike AI engagement, exploration is a domain in which avoiders face no conflict. Nothing about avoiding AI makes exploring one's options harder. This account does not claim that avoidance is costless. It claims that, averaged across the awareness spectrum, the costs of avoidance are lower than the costs of use. Section 2.5 refines this average into its conditional form.

H1. *Avoidant AI use is positively associated with (a) environment exploration and (b) self-exploration.*

2.4. The Threat-Appraisal Account: A Negative Linear Moderation

The consequences of avoidant AI use may depend on how students appraise AI. Threat-appraisal reasoning, which we operationalize as AI-shock awareness, gives a clear prediction about one such pattern. In the coping model of user adaptation, avoidance born of threat appraisal is an emotion-focused response. It manages distress rather than addressing its source (Beaudry & Pinsonneault, 2005; Bala & Venkatesh, 2016). Wang et al. (2025) showed that AI-shock awareness feeds avoidant AI use through this threat channel. It follows that the higher a student's awareness of AI's threat to their career, the more their avoidance is driven by threat rather than by habit or preference. Threat-driven avoidance differs from relaxed non-use. It comes with vigilance toward the avoided object, rumination about the future it represents, and the uncomfortable knowledge that one is steering clear of something that matters. Emotion-focused coping of this kind is most taxing, and least compatible with proactive behavior, when the threat is appraised as severe (Lazarus & Folkman, 1984). Evidence from the job insecurity literature points the same way. As perceived threat intensifies, avoidant and defensive responses tie more strongly to adverse work outcomes (Greenhalgh & Rosenblatt, 1984). Therefore, as AI-shock awareness increases, avoidance becomes increasingly threat-laden, and its positive association with career exploration should weaken.

H2. *AI-shock awareness negatively moderates the association between avoidant AI use and (a) environment exploration and (b) self-exploration, such that the association weakens as awareness increases.*

2.5. The Inverted-U Cost of Avoidance: A Curvilinear Moderation

Non-additive effects are not unique to our setting. In higher education, Zhou (2024) showed that a nonlinear teaching approach shifted learning outcomes only through an interaction with student attendance—an effect a purely linear specification would miss. The linear account, however, faces a complication it cannot absorb. Heightened awareness might mobilize rather than tax. A student acutely aware of AI's career impact may compensate for non-use by investing more heavily in non-AI exploration channels, turning avoidance into a springboard rather than a drain. The threat literature is mixed on whether threat energizes or paralyzes, and a single linear slope must assign one sign to the whole awareness continuum, so it cannot accommodate this tension. Self-regulation theory offers a resolution. The cost of an avoidant stance is governed not by threat intensity alone, but by the internal conflict the stance generates. The degree of conflict depends on how salient the avoided technology is, and salience is exactly what AI-shock awareness captures. Across the awareness continuum, the self-regulatory cost of avoidance follows an inverted-U.

At low awareness, avoidance is habitual. AI is barely present in the student's psychological landscape, and not using it requires no deliberation and generates no conflict. Well-practiced behaviors in stable contexts recur automatically, without engaging conscious decision-making (Ouellette & Wood, 1998). The stance should therefore cost almost nothing, leaving the resources that AI engagement would otherwise consume available for exploration. At moderate awareness, the stance becomes contested. The student now recognizes that AI matters for their future, yet continues to resist it. This resembles what Gollwitzer (1990) called an action crisis. Reminders of AI's relevance may repeatedly reactivate an unresolved decision, and the sustained effort of resistance could draw continuously on the resource pool that exploration also needs. Evidence on anxiety and cognition points the same way. Threat-related preoccupation diverts attentional control away from goal-directed processing and taxes the executive functions on which effortful

behavior depends (Eysenck et al., 2007). The avoidance-exploration link should be weakest here, not because avoidance has become threatening in itself, but because maintaining it may have become expensive. At high awareness, the conflict is likely resolved. A student who fully grasps AI's significance and nonetheless chooses to avoid it has converted resistance into commitment. Post-decisional, implemental mindsets close off deliberation and redirect cognition toward carrying out the chosen course (Gollwitzer, 1990; Heckhausen & Gollwitzer, 1987). With the decision made and the internal debate closed, the stance becomes self-sustaining and its regulatory cost should recede, releasing resources and allowing the avoidance-exploration link to recover. These mediating states, habitual ease, internal conflict, and resolved commitment, are inferred from the curvilinear pattern rather than directly measured, and Section 5.3 treats this as a limitation.

This inverted-U cost structure implies a U-shaped moderation. The positive association between avoidant AI use and career exploration should be reliable at low and high awareness and weakest at moderate awareness. The curvilinear account subsumes the linear one. Over the low-to-moderate range, rising awareness moves avoidance from habitual to contested, and the link weakens, just as the threat-appraisal account predicts. The accounts diverge only at the high end, where the linear account predicts continued weakening while the self-regulation account predicts recovery. The curvilinear hypothesis is therefore the riskier and more informative prediction. It would be contradicted by a monotonic pattern and by an inverted-U in the association itself.

H3. *AI-shock awareness exerts a U-shaped curvilinear moderation on the association between avoidant AI use and (a) environment exploration and (b) self-exploration, such that the association is stronger at low and high awareness and weaker at moderate awareness.*

Figure 1 depicts the full model. Gender, club management experience, STEM major, grade, and economic region serve as controls.

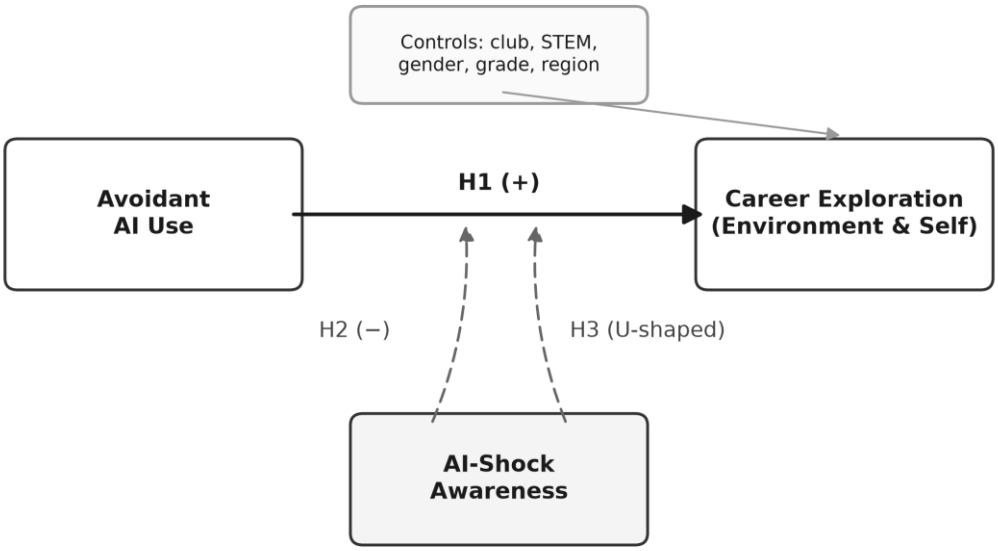


Figure 1. Theoretical model. H1: main effect of avoidant AI use on career exploration. H2: negative linear moderation by AI-shock awareness. H3: U-shaped curvilinear moderation by AI-shock awareness.

3. Materials and Methods

3.1. Participants and Procedure

Participants were 289 students recruited through snowball sampling from multiple universities across China. An online questionnaire was administered through the Wenjuanxing platform in spring 2026. Initial respondents were students in the research team's course and personal networks at several Chinese universities, who were invited to share the survey link with peers at their own and other institutions. Eligible participants were currently enrolled university students; participation was voluntary and anonymous. Of the 314 questionnaires returned, we excluded incomplete responses and invalid responses showing implausibly short completion times or invariant, patterned answering, yielding 289 valid responses. Because snowball sampling provides no known sampling frame, we do not report a conventional response rate; 92.0% denotes the proportion of returned questionnaires retained as valid (289 of 314). The exact breakdown of exclusions by criterion was not archived, and the survey did not record respondents' institutional affiliation, so the number and type of universities represented cannot be reported; we note both as reporting limitations. The final sample comprised 175 women (60.6%); respondents were predominantly in their third and fourth years of undergraduate study (grade level coded 1 = first-year undergraduate to 5 = graduate student; M = 3.59), 38.4% studied STEM majors, and 30.8% reported student-club management experience. Procedural remedies to reduce common method bias (Podsakoff et al., 2003) included anonymity assurance, no right-or-wrong-answer framing, separation of focal constructs with distinct instructions, and different scale anchors for predictor and outcome items where feasible. All continuous predictors were mean-centered before computing quadratic and interaction terms. This study was reviewed under the exemption policy of the Ethics Committee of the School of Management, Capital Normal University, which waived ethical review and approval because the study involved an anonymous, non-interventional online survey of adults and collected no personal, sensitive, or identifiable data. Before starting the questionnaire, participants read an information page describing the study purpose, voluntariness, anonymity, and data use; proceeding to the items constituted informed consent. This study was conducted in accordance with the principles of the Declaration of Helsinki.

3.2. Measures

Avoidant AI use (α = 0.858) was measured with four items adapted from Bala and Venkatesh (2016), assessing active withdrawal from AI (e.g., "I try to avoid using AI to complete learning or tasks"). Following Bala and Venkatesh (2016, Section 2.2.4), the construct is defined behaviorally as the degree to which a person tries not to use a technology when accomplishing tasks; the items therefore capture active withdrawal rather than the motives behind it (see Sections 2.2 and 5.3). AI-shock awareness (α = 0.925) was measured with four items adapted from the STARA awareness scale of Brougham and Haar (2018), capturing perceived threat of AI to one's future career (e.g., "I am worried about my future development because AI will replace workers"). Environment exploration (α = 0.892) and self-exploration (α = 0.900) were assessed with four items each, drawn from the Career Exploration Survey (Stumpf et al., 1983). All items used a 5-point Likert scale. Items originating from English-language scales were translated into Chinese and independently back-translated by bilingual researchers, with discrepancies resolved through discussion. The Chinese adaptations of the avoidant-use and AI-shock-awareness measures have demonstrated adequate psychometric properties in prior Chinese samples (e.g., Wang et al., 2025). Controls were chosen a priori and dummy-coded: gender (female = 1, male = 0), club management experience (yes = 1, no = 0; included because campus leadership roles may absorb time and attention relevant to career exploration), STEM major (1 = science,

engineering, medicine, agriculture, or related technical fields, 0 = otherwise; we note that this is a broad definition), and grade, entered as four dummy indicators for second-year, third-year, fourth-year, and graduate students (reference = first-year undergraduate). Economic region was controlled because geographic context can shape AI exposure and career exploration alike. Region was coded from the province implied by the survey platform's IP geolocation and entered as two dummy indicators for the central and western belts (reference = the eastern belt), treated as a proxy for geographic context (Tran et al., 2023; Anghel et al., 2024). Interpretations of control-variable associations are post hoc (Section 4.4). Composite reliability and average variance extracted are reported in Section 4.3.

3.3. Analytic Strategy

Harman's single-factor test was used to assess common method variance. Convergent and discriminant validity were examined via confirmatory factor analysis (CR, AVE, SRMR), the Fornell–Larcker criterion, and heterotrait–monotrait (HTMT) ratios. Hypotheses were tested with three-step hierarchical multiple regression for each outcome: Model 1 entered avoidant AI use and controls; Model 2 added the linear and squared terms of AI-shock awareness; Model 3 added the two interactions. Primary inference relied on BCa bootstrap 95% confidence intervals (5000 resamples; Efron & Tibshirani, 1993). Because Breusch–Pagan scores indicated heteroscedasticity, Model 3 was re-estimated with HC3 robust standard errors. CFA factor scores were used to address measurement error. Significant interactions were probed with percentile bootstrap simple-slope and Johnson–Neyman analyses (Hayes, 2018), and influence diagnostics and exclusion analyses were conducted to evaluate whether the curvilinear interaction depended on particular subgroups of respondents. Because bootstrap intervals are the primary inferential criterion, significance stars for the focal terms follow the BCa intervals (see the table notes).

4. Results

4.1. Common Method Variance

Harman's single-factor test indicated that the first unrotated factor accounted for 35.2% of the variance, below the 40–50% heuristic, and four factors had eigenvalues above 1.0. Harman's test is an insensitive diagnostic and cannot by itself rule out common method variance (Podsakoff et al., 2003), and the specificity of a curvilinear pattern is likewise not sufficient evidence against a common-method explanation. We therefore base our assessment on the procedural remedies described in Section 3.1 and the discriminant-validity evidence in Section 4.3, and we treat residual common method variance as a limitation of the single-source design (Section 5.3).

4.2. Descriptive Statistics and Correlations

Table 1 presents means, standard deviations, and intercorrelations. Avoidant AI use was positively correlated with environment exploration (r = 0.20) and self-exploration (r = 0.17). AI-shock awareness correlated positively with both (r = 0.26 and r = 0.17).

Table 1. Means, standard deviations, and intercorrelations.

Variable	M	SD	1	2	3	4	5	6	7
1. Env. exploration	3.63	0.88	—						
2. Self-exploration	3.84	0.79	0.59 ***	—					
3. Avoidant AI use	2.82	0.93	0.20 ***	0.17 **	—				
4. AI-shock awareness	2.82	1.06	0.26 ***	0.17 **	0.16 **	—			
5. Club mgmt	0.31	0.46	−0.22 ***	−0.15 *	0.01	−0.05	—		
6. STEM	0.38	0.49	0.08	−0.02	0.07	0.03	−0.08	—	

7. Female	0.61	0.49	0.00	0.01	−0.18 **	−0.02	−0.01	−0.27 ***	—

Note. N = 289. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

4.3. Measurement Model

A four-factor confirmatory model (Kline, 2016) fit the data well ($\chi^2(98) = 168.05$, CFI = 0.977, TLI = 0.972, RMSEA = 0.050, SRMR = 0.041). Table 2 shows that composite reliability exceeded 0.86 and AVE exceeded 0.61 for all constructs. The two-factor exploration model fit substantially better than a one-factor model ($\Delta\chi^2(1) = 309.28$, $p < 0.001$; latent r = 0.67), confirming the empirical distinctness of environment and self-exploration. Discriminant validity was further supported: for every construct, the square root of the AVE exceeded its correlations with the other constructs (Fornell–Larcker criterion), and all HTMT ratios were below 0.70 (maximum 0.661, between environment and self-exploration), within the 0.85 guideline (Henseler et al., 2015).

Table 2. Composite reliability and average variance extracted.

Construct	CR	AVE	Std. Loadings (Range)
Avoidant AI use	0.863	0.615	0.63–0.90
AI-shock awareness	0.926	0.758	0.79–0.92
Environment exploration	0.893	0.676	0.79–0.86
Self-exploration	0.902	0.698	0.78–0.89

Note. N = 289. Standardized estimates.

4.4. Hypothesis Testing

Tables 3 and 4 report the hierarchical regressions, with Cohen's f^2 as the effect-size measure (Cohen, 1988). VIFs ranged from 1.01 to 4.59 across the three models, with the largest values attaching to the imbalanced dummy controls (the fourth-year and club-management indicators), and all were below the conventional threshold of 5. Durbin–Watson statistics ranged from 1.89 to 2.00.

Table 3. Hierarchical regression predicting environment exploration.

Predictor	M1 b	(SE)	M2 b	(SE)	M3 b	(SE)
Avoidant AI use	0.203 ***	(0.056)	0.167 **	(0.055)	0.003	(0.078)
Club management	−0.426 ***	(0.109)	−0.407 ***	(0.106)	−0.396 ***	(0.105)
STEM major	0.142	(0.109)	0.125	(0.106)	0.127	(0.104)
Female	0.095	(0.111)	0.091	(0.109)	0.102	(0.108)
Grade 2 (sophomore)	−0.083	(0.252)	−0.194	(0.246)	−0.199	(0.243)
Grade 3 (junior)	−0.078	(0.238)	−0.050	(0.231)	−0.135	(0.230)
Grade 4 (senior)	0.050	(0.214)	0.070	(0.208)	0.021	(0.206)
Grade 5 (graduate)	0.018	(0.243)	0.062	(0.236)	−0.005	(0.234)
Central region	−0.092	(0.106)	−0.042	(0.104)	−0.032	(0.103)
Western region	−0.253	(0.239)	−0.202	(0.233)	−0.162	(0.230)
AI-shock awareness			0.192 ***	(0.048)	0.183 ***	(0.048)
AI-shock awareness²			0.042	(0.037)	0.034	(0.037)
Avoidant × awareness					−0.070	(0.042)
Avoidant × awareness²					0.095 **	(0.032)
R^2	0.105		0.161		0.189	
Adjusted R^2	0.073		0.124		0.148	
ΔR^2	0.105		0.055		0.028	
Cohen's f^2	—		0.066		0.035	
F	3.28		4.40		4.56	

Note. N = 289. BCa bootstrap 5000. Model 3: Avoidant × awareness² b = 0.095, BCa 95% CI [0.012, 0.186]. ** $p < 0.01$, *** $p < 0.001$.

Table 4. Hierarchical regression predicting self-exploration.

Predictor	M1 b	(SE)	M2 b	(SE)	M3 b	(SE)
Avoidant AI use	0.158 **	(0.050)	0.138 **	(0.050)	−0.061	(0.070)
Club management	−0.263 **	(0.099)	−0.254 **	(0.098)	−0.241 *	(0.095)
STEM major	−0.019	(0.098)	−0.030	(0.097)	−0.031	(0.094)
Female	0.016	(0.100)	0.021	(0.100)	0.033	(0.097)
Grade 2 (sophomore)	0.151	(0.227)	0.080	(0.226)	0.072	(0.219)
Grade 3 (junior)	0.356	(0.215)	0.373	(0.212)	0.259	(0.208)
Grade 4 (senior)	0.357	(0.193)	0.364	(0.191)	0.295	(0.186)
Grade 5 (graduate)	0.314	(0.219)	0.335	(0.217)	0.250	(0.212)
Central region	−0.209 *	(0.095)	−0.172	(0.096)	−0.165	(0.093)
Western region	−0.243	(0.215)	−0.208	(0.214)	−0.161	(0.208)
AI-shock awareness			0.106 *	(0.044)	0.095 *	(0.043)
AI-shock awareness²			0.044	(0.034)	0.036	(0.033)
Avoidant × awareness					−0.105 †	(0.038)
Avoidant × awareness²					0.116 ***	(0.029)
R^2	0.084		0.110		0.168	
Adjusted R^2	0.051		0.071		0.126	
ΔR^2	0.084		0.026		0.058	
Cohen's f^2	—		0.029		0.070	
F	2.55		2.84		3.95	

Note. N = 289. BCa bootstrap 5000. Model 3: Avoidant × awareness² b = 0.116, BCa 95% CI [0.038, 0.188]. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. For the focal terms, inference follows the BCa bootstrap 95% confidence intervals (Section 3.3). † The linear interaction for self-exploration shows marginal support: its standard-error-based $p = 0.006$, but its BCa 95% CI [−0.207, −0.004] has an upper bound close to zero, so it is interpreted cautiously.

H1 predicted a positive association between avoidant AI use and exploration. It was supported for both outcomes: avoidant AI use was positively associated with environment exploration (Model 1 b = 0.203, 95% BCa CI [0.069, 0.342]) and with self-exploration (Model 1 b = 0.158, 95% BCa CI [0.042, 0.286]), and the associations remained positive and significant when the moderator terms entered in Model 2 (environment b = 0.167, $p < 0.01$; self b = 0.138, $p < 0.01$). In Model 3, the main-effect coefficients were attenuated to non-significance (environment b = 0.003; self b = −0.061). This attenuation is expected rather than contradictory: with the quadratic moderation present, the main effect is the conditional association at mean awareness, the region where the U-shaped pattern places the weakest association (Table 5). H1 therefore holds as an average positive association across the awareness continuum.

Table 5. Conditional effects of avoidant AI use at levels of AI-shock awareness.

Outcome	Awareness Level	b	95% CI	Sig.
Environment	−1 SD	0.18	[−0.01, 0.37]	No
	Mean	0.00	[−0.16, 0.16]	No
	+1 SD	0.04	[−0.12, 0.20]	No
Self	−1 SD	0.18	[0.01, 0.33]	Yes
	Mean	−0.06	[−0.20, 0.08]	No
	+1 SD	−0.04	[−0.18, 0.11]	No

Note. Percentile bootstrap 5000. Sig. = confidence interval excludes zero.

H2 predicted a negative linear moderation. The direction of the linear interaction was consistent with this prediction in both models (environment b = −0.070; self b = −0.105), but the bootstrap confidence intervals told a mixed story: the environment interval [−0.181, 0.050] included zero, while the self-exploration interval [−0.207, −0.004] excluded zero only narrowly. H2 was therefore not supported for environment exploration and only marginally supported for self-exploration. This marginal result should be read alongside H3: a U-shaped moderation implies opposite slopes at low and high awareness, so a single linear slope is not a clean test of how the association varies with awareness.

H3 predicted a U-shaped curvilinear moderation. It was supported for self-exploration and partially supported for environment exploration. The quadratic interaction was positive and significant for environment exploration (b = 0.095, 95% BCa CI [0.012, 0.186]) and for self-exploration (b = 0.116, 95% BCa CI [0.038, 0.188]), over and above the linear term. A positive quadratic term with a negative linear term locates the weakest point of the association at intermediate awareness, the signature of the predicted U-shape. The pattern helps interpret the weakness of the linear term: a linear specification averages the opposing slopes at the two ends of the awareness continuum and consequently finds little. Supplementary exclusion analyses (Section 4.5) indicated that the pattern for self-exploration was carried by the full sample, whereas the high-awareness arm for environment exploration was sensitive to the small number of respondents in that region; we therefore regard the curvilinear pattern for environment exploration as supported, but its high-awareness rebound as preliminary (see also Sections 4.5 and 4.6).

Ancillary results for the control variables are worth noting. Club management experience was negatively associated with both environment exploration (b = −0.396, $p < 0.001$) and self-exploration (b = −0.241, $p < 0.05$). A plausible interpretation is that students who invest heavily in current campus leadership roles allocate attention and time to those responsibilities, leaving fewer resources for future-oriented career exploration. This role-conflict account is post hoc and should be tested directly in future work. In the full model (Model 3), STEM major, gender, grade, and geographic region were not significantly associated with either outcome.

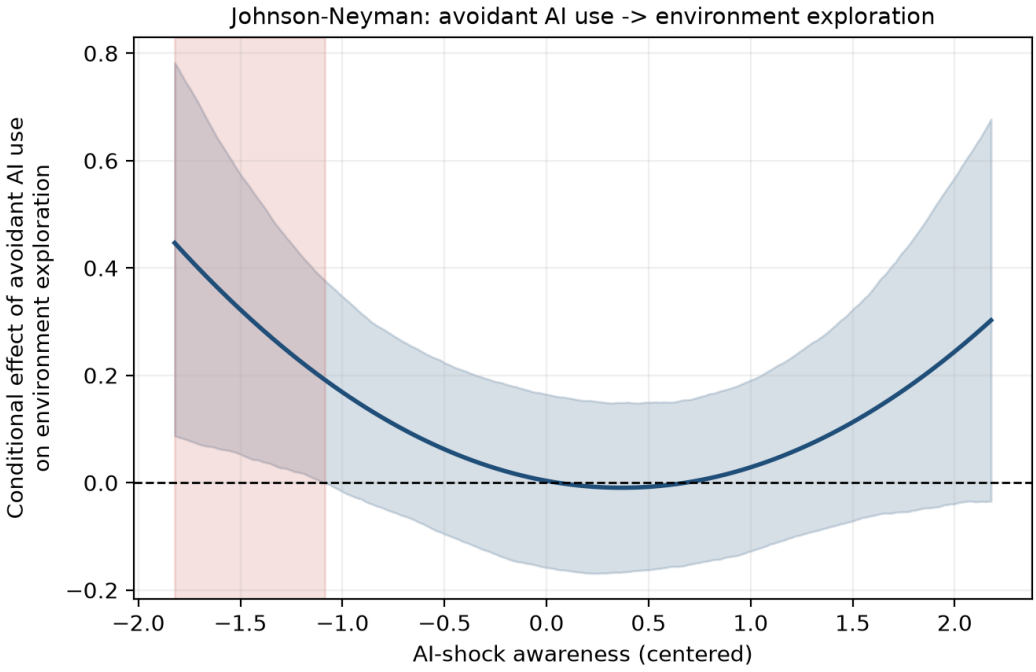


Figure 2. Conditional effect of avoidant AI use on environment exploration across AI-shock awareness (centered). The solid blue line denotes the conditional point estimate; the light blue band denotes the 95% bootstrap confidence interval (5,000 resamples); the dashed horizontal line denotes the zero reference. Pink shaded regions denote Johnson–Neyman significance regions, within which the confidence interval excludes zero.

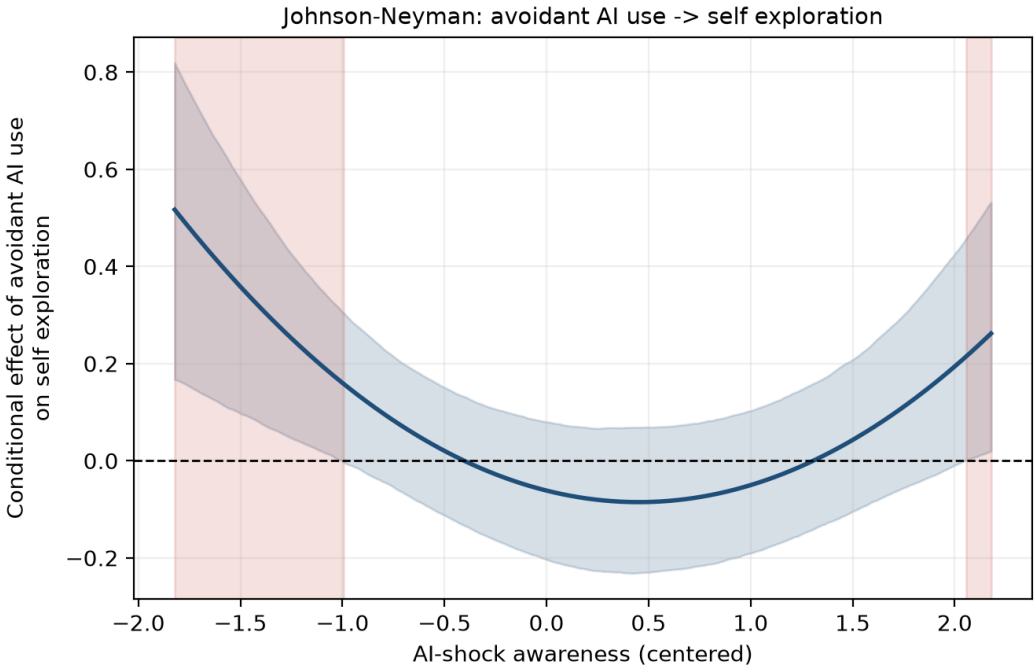


Figure 3. Conditional effect of avoidant AI use on self-exploration across AI-shock awareness (centered). The solid blue line denotes the conditional point estimate; the light blue band denotes the 95% bootstrap confidence interval (5,000 resamples); the dashed horizontal line denotes the zero reference. Pink shaded regions denote Johnson–Neyman significance regions, within which the confidence interval excludes zero.

4.5. Robustness Checks

Under HC3 robust standard errors, the quadratic interaction remained significant for self-exploration (b = 0.116, $p = 0.004$) and marginal for environment exploration (b = 0.095, $p = 0.049$); the bootstrap interval for environment exploration excluded zero. Re-estimating the model with CFA factor scores strengthened this quadratic interaction: the quadratic interaction term was b = 0.184 ($p = 0.007$) for environment exploration and b = 0.220 ($p = 0.001$) for self-exploration. A latent-variable structural model with a linear latent interaction confirmed the pattern (CFI = 0.947, TLI = 0.937, RMSEA = 0.065; main effects significant, linear interaction non-significant).

The upper tail of the moderator distribution was thin. AI-shock awareness averaged 2.82 (SD = 1.06) on the 5-point scale, and only 18 respondents (6.2%) scored at the scale maximum (5.00; centered values above 2.06, the Johnson–Neyman upper threshold reported in Section 4.6). The associations were least stable in this region. We therefore assessed whether the curvilinear pattern depended on that subgroup. Cook's distances were below conventional thresholds for severe influence (maxima 0.18 and 0.08), and excluding respondents with above-threshold values left both quadratic interactions significant. Excluding these 18 respondents told the two outcomes apart. For self-exploration, the quadratic interaction remained significant and was in fact stronger (b = 0.181, $p < 0.001$), indicating that the curvilinear pattern is carried by the sample as a whole rather than by the

upper tail. For environment exploration, the coefficient retained its sign but lost significance (b = 0.093, $p = 0.096$), partly a mechanical consequence of truncating the moderator range over which a quadratic effect is identified, but nonetheless grounds for treating the high-awareness rebound in that model as a preliminary trend awaiting confirmation in larger samples. Finally, all models were re-estimated without the two region dummies; every substantive conclusion was unchanged, and a full side-by-side comparison of the two specifications is provided in the Supplementary Material (Tables S1 and S2).

4.6. Simple Slopes and Johnson–Neyman Analysis

The Johnson–Neyman analysis, based on percentile bootstrap 95% confidence intervals for the simple slope, located the significant regions as follows (Figure 2, environment exploration; Figure 3, self-exploration). For self-exploration, the association was significant below a centered awareness of −0.99 and above 2.06, a two-tailed U-shape (Figure 3). This upper region contains the 18 respondents (6.2%) at the scale maximum, so estimates within it rest on a small subgroup and should be read with caution. For environment exploration, it was significant only below −1.09 (Figure 2); the high-awareness arm did not reach significance in the bootstrap probe, consistent with the exclusion analysis in Section 4.5. The association was non-significant at the mean in both models, as H3 predicts. Across all models, VIFs ranged from 1.01 to 4.59 and Durbin–Watson statistics from 1.89 to 2.00; Breusch–Pagan and Shapiro–Wilk tests indicated some heteroscedasticity and non-normality, which bootstrap-based inference accommodates.

5. Discussion

5.1. Findings in Dialogue with the Literature

Three findings summarize the results. First, avoidant AI use was positively associated with both environment and self-exploration on average; the association was reliable in Models 1 and 2 and, as the U-shape implies, near zero at mean awareness in Model 3. Second, the linear interaction with AI-shock awareness was negative in sign, as threat-appraisal reasoning predicts, but reliable only at the margin, and only for self-exploration, once estimation uncertainty was taken into account. Third, the quadratic interaction was positive and significant for both outcomes, although the high-awareness rebound was robust only for self-exploration (Sections 4.5 and 4.6). Simple average scores inevitably contain measurement error, which tends to dilute true effects. When we re-estimated the models using factor scores extracted from the confirmatory factor analysis, which separate the shared true variance of the items from measurement error, the quadratic interactions became stronger rather than weaker. This pattern indicates that the curvilinear finding is not an artifact of measurement noise. The pattern is consistent with the proposed cost account: the link was reliable where the account predicts avoidance to be cheap and undetectable in the moderate range where it predicts the highest self-regulatory cost. Because the mediating states were not measured, we offer this as an interpretation consistent with the data rather than as direct evidence of the mechanism (Section 5.3).

These findings give coping theory a concrete answer to the goodness-of-fit question. Lazarus and Folkman (1984) showed that avoidance is adaptive under some conditions and maladaptive under others, though they said less about which conditions produce which outcome. Our results identify one candidate condition: the appraised significance of the avoided stressor itself. A reading consistent with the data is that when the stressor is barely appraised, avoidance may be habitual and largely harmless to concurrent goal pursuit; that when it is appraised at full strength and nonetheless declined, avoidance may function as a resolved commitment that again leaves resources intact; and that in the middle, where the stressor is acknowledged but the response stays unsettled, avoidance

may turn costly. These states were inferred, not measured. The adaptiveness of avoidance therefore varies with the very appraisal that defines the coping situation. This brings the contingency principle closer to a testable form.

For the technology non-use literature, the findings both support and qualify the claim that non-use can be "active, meaningful, motivated, considered, structured, specific, nuanced, directed, and productive" (Satchell & Dourish, 2009). The support lies in the main effect. Students who actively steer clear of AI explore their career options more, not less. The qualification lies in the moderation. Non-use is not uniformly productive. Its productivity may depend on the self-regulatory cost of sustaining it, which our account places at a peak in the contested middle of the awareness continuum. Studies of technology adoption and resistance have largely treated non-use as an endpoint to be explained. Our results suggest treating it as a stance whose ongoing maintenance has downstream consequences for other valued behaviors.

These findings also widen the antecedent net of career exploration. The model's adapting process has usually been linked to resources such as adaptability, efficacy, and support (Savickas, 2005, 2013). We add a behavioral antecedent from the technology domain. How a student manages their stance toward AI spills over into how much career exploration they can sustain. This spillover makes sense once we view adapting behaviors as competitors for a shared resource pool. Career construction does not happen in a vacuum sealed off from other domains of self-regulation. It may be financed, so to speak, by whatever the student is not spending elsewhere, a resource-flow interpretation that our cross-sectional data cannot test directly.

Two interpretive caveats deserve emphasis. First, the cross-sectional design cannot rule out reverse causality: students who are proactive career explorers may deliberately limit AI use to preserve autonomous, critical thinking. We cannot adjudicate this alternative with the present data, and we discuss it further in Section 5.3. Second, the two arms of the U-shape are not equally firm. The low-awareness arm rests on a substantial share of the sample, and for self-exploration the curvilinear pattern survived the exclusion of respondents in the high-awareness region. For environment exploration, however, the high-awareness rebound depends on a comparatively small group of highly aware respondents and lost significance when they were excluded. That rebound should therefore be read as a preliminary trend awaiting confirmation in larger, more AI-exposed populations, not as a definitive conclusion.

5.2. Implications for Practice

The U-shaped pattern carries differentiated implications for career education and AI-literacy instruction in higher education (Ng et al., 2021). The common thread is that interventions should help students resolve their stance toward AI, whatever that resolution looks like, rather than pushing a uniform message about adoption. We offer these implications with two caveats. First, the three awareness zones discussed below are descriptive regions of a continuum, not empirically derived cut points; our analyses located significance regions but did not validate thresholds for classifying individual students. Second, the characterization of avoidance as habitual, contested, or resolved within these zones is theoretical and was not measured. The suggestions that follow are therefore directions for program design and evaluation, not validated protocols.

For students with low AI-shock awareness, avoidance may be habitual and conflict-free. Career educators can support these students with structured exploration resources delivered through traditional, non-AI channels, such as alumni panels, industry site visits, and informational interviews, rather than insisting on AI adoption that carries no perceived urgency for them. An AI-literacy module for this group could introduce AI tools

in low-stakes, career-relevant contexts, such as using a chatbot to generate a list of interview questions. Students who are open to AI can engage without feeling pressured, and those who prefer to abstain are not penalized.

For students with moderate awareness, avoidance may be the most psychologically costly, because it may generate ongoing internal conflict. These students know AI matters but are trying not to use it, and the sustained self-control would deplete resources they could otherwise invest in exploration. The priority here is conflict resolution. Instructors designing prompting or critical-thinking modules can structure activities that make the conflict explicit. For example, students might list three career tasks they currently avoid doing with AI, evaluate what they gain and lose by avoiding each, and then make a conscious choice to try one or to commit to abstaining. The point is not to convert avoiders into users, but to replace an anxious, unresolved stance with a deliberate one. Career counselors working with these students can use structured decision-making frameworks to help them articulate what specifically concerns them about AI and whether their avoidance is serving their goals.

For students with high awareness, avoidance may function as a resolved commitment. Career counselors should respect this choice while ensuring it does not translate into isolation from career development. Students with high AI anxiety, who perceive AI as a genuine threat, may benefit from counseling that distinguishes between threat perception and actual capability, helping them see that avoiding AI does not mean abandoning career planning. Those whose avoidance is principled, such as a commitment to human-centered work, should be supported with high-quality, non-AI exploration resources. In both cases, the counselor's role is to validate the stance while broadening the toolkit available within it.

Across all three zones, the practical message is that career education should not be organized around a single assumption about whether students should use AI. A well-designed AI-literacy curriculum might begin not with a skills tutorial, but with a structured reflection on what AI means to each student and how their current stance is serving or hindering their goals. This aligns with calls to build AI education around students' values and critical judgment rather than tool proficiency alone (Chan, 2023). Only after this reflection should skill-building follow, tailored to the student's chosen direction. This design principle applies equally to career guidance protocols, where the counselor's first task is to understand the student's AI stance and its self-regulatory cost before recommending a course of action.

5.3. Limitations and Future Directions

Several limitations qualify these conclusions. The cross-sectional, single-source design precludes causal claims; the associations may be reciprocal, and longitudinal or experimental designs are needed to establish directionality. All measures were self-reported. The procedural safeguards and Harman's single-factor test provide some reassurance, but neither can exclude common method variance in a single-source design, and the specificity of the curvilinear pattern is not by itself sufficient evidence against it. Future studies should incorporate multi-source or objective measures of AI use and career exploration where feasible. The sample was recruited through snowball sampling across Chinese universities. This approach reached students at multiple institutions, but it is a non-probability method prone to self-selection. Students willing to complete and forward a survey about AI may differ systematically from those who are not. The sample therefore cannot be treated as representative of Chinese university students in general. Replication with probability-based samples, and in other cultural and institutional contexts, is essential before broader claims about the role of AI avoidance in career development can be made. In

addition, because institutional affiliation was not recorded, we cannot report the number or type of universities reached; future administrations of the survey should log this information.

Respondents' region was inferred from the IP geolocation captured automatically by the survey platform rather than self-reported. Although IP-based geolocation is a commonly used proxy for respondent location in online studies (Tran et al., 2023; Anghel et al., 2024), it may mislocate respondents using virtual private networks or mobile carriers that route traffic through distant nodes; region should therefore be read as an approximate indicator of geographic context.

The self-regulatory cost mechanism was inferred from the curvilinear pattern but not directly measured. We did not assess students' internal conflict about AI use, their self-regulatory resource levels, or their experience of action crisis. Future work should measure these constructs and test whether they mediate the U-shaped moderation. Direct measurement of the proposed mediating states would provide stronger evidence for the self-regulation account than the curvilinear pattern alone can offer.

Third, our measure of avoidant AI use captures active withdrawal but not its motive. In Bala and Venkatesh's (2016) model, avoidance can arise from threat appraisal but also from low perceived controllability, including perceived lack of ability or resources. We did not measure AI literacy, usage frequency, or experience, so deliberate and incapacity-driven avoidance cannot be separated in our data. Future studies should measure AI competence and usage directly to separate deliberate from incapacity-driven avoidance.

The high-awareness arms of the U-shape relied on a small subset of respondents ($n = 18$ at the scale maximum, 6.2% of the sample), and the environment-exploration interaction was marginal under HC3 standard errors and non-significant when respondents in the high-awareness region were excluded. Replication with larger and more diverse samples is needed to confirm the rebound at high awareness and to increase the precision of the estimated thresholds. Club management experience showed a negative association with exploration, possibly because students in leadership roles allocate less attention to future-oriented exploration, a post hoc account that should be tested directly. We did not estimate a latent quadratic-moderation (LMS) model, though factor-score analysis addressed the measurement-error concern.

6. Conclusions

Among 289 Chinese university students, avoidant AI use was positively associated with both environment and self-exploration. A negative linear moderation by AI-shock awareness, predicted by threat-appraisal reasoning, matched the data in direction and, for self-exploration, only at the margin of reliability. What the data did support was a U-shaped contingency. The association held at low awareness and, for self-exploration, at high awareness; it was undetectable at moderate awareness. For environment exploration, the high-awareness rebound rested on a small subgroup and remains preliminary. This U-shaped pattern is consistent with the proposed self-regulatory cost account, in which avoidance is cheap where habitual or resolved and costly where contested, although that mechanism was inferred rather than measured. For career education, the implication is that helping students is not a matter of pushing them toward or away from AI. It is a matter of helping them settle their relationship with it, so that avoidance, if chosen, does not come at the cost of exploration.

Author Contributions: Conceptualization, B.Z. and Y.W.; methodology, B.Z., Y.W., and P.Z.; software, B.Z.; validation, B.Z.; formal analysis, B.Z.; investigation, Y.W.; data curation, Y.W.; writing—original draft preparation, B.Z.; writing—review and editing, B.Z., Y.W., and P.Z.; supervision, P.Z.; project administration, P.Z.; funding acquisition, P.Z. All authors have read and agreed to the published version of the manuscript.

Funding: The APC was funded by Open Access Support Team, Cardiff University, No. 2026-OA-0957.

Institutional Review Board Statement: Ethical review and approval were waived for this study by the Ethics Committee of the School of Management, Capital Normal University, under its Policy on Exemption from Ethical Approval, because the study involved an anonymous, non-interventional online survey of adults and collected no personal, sensitive, or identifiable data. This study was conducted in accordance with the principles of the Declaration of Helsinki.

Informed Consent Statement: Informed consent was obtained from all participants. Before the questionnaire began, an information page described the study purpose, the voluntariness of participation, anonymity, and data use; proceeding to the questionnaire constituted consent. No identifiable personal data were collected.

Data Availability Statement: The data are available on reasonable request from the corresponding author.

Acknowledgments: During the preparation of this manuscript, the authors used an artificial intelligence tool to refine language and improve readability. The authors reviewed and edited all output and take full responsibility for the publication's content.

Conflicts of Interest: The authors declare no conflicts of interest.

References

(Anghel et al., 2024) Anghel, E., Littenberg-Tobias, J., & Reich, J. (2024). Location in the multiverse of methods: Measuring online users' contexts. *International Journal of Social Research Methodology*, *27*(2), 159–172. https://doi.org/10.1080/13645579.2022.2125648.

(Bala & Venkatesh, 2016) Bala, H., & Venkatesh, V. (2016). Adaptation to information technology: A holistic nomological network from implementation to job outcomes. *Management Science*, *62*(1), 156–179. https://doi.org/10.1287/mnsc.2014.2111.

(Baumeister et al., 1998) Baumeister, R. F., Bratslavsky, E., Muraven, M., & Tice, D. M. (1998). Ego depletion: Is the active self a limited resource? *Journal of Personality and Social Psychology*, *74*(5), 1252–1265. https://doi.org/10.1037/0022-3514.74.5.1252.

(Beaudry & Pinsonneault, 2005) Beaudry, A., & Pinsonneault, A. (2005). Understanding user responses to information technology: A coping model of user adaptation. *MIS Quarterly*, *29*(3), 493–524. https://doi.org/10.2307/25148693.

(Brougham & Haar, 2018) Brougham, D., & Haar, J. (2018). Smart technology, artificial intelligence, robotics, and algorithms (STARA): Employees' perceptions of our future workplace. *Journal of Management & Organization*, *24*(2), 239–257. https://doi.org/10.1017/jmo.2016.55.

(Carver & Scheier, 1981) Carver, C. S., & Scheier, M. F. (1981). *Attention and self-regulation: A control-theory approach to human behavior*. Springer-Verlag.

(Carver & Scheier, 1998) Carver, C. S., & Scheier, M. F. (1998). *On the self-regulation of behavior*. Cambridge University Press.

(Chan, 2023) Chan, C. K. Y. (2023). A comprehensive AI policy education framework for university teaching and learning. *International Journal of Educational Technology in Higher Education*, *20*, 38. https://doi.org/10.1186/s41239-023-00408-3.

(Chan & Hu, 2023) Chan, C. K. Y., & Hu, W. (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, 1988) Cohen, J. (1988). *Statistical power analysis for the behavioral sciences* (2nd ed.). Lawrence Erlbaum Associates.

(Crompton & Burke, 2023) Crompton, H., & Burke, D. (2023). Artificial intelligence in higher education: The state of the field. *International Journal of Educational Technology in Higher Education*, *20*, 22. https://doi.org/10.1186/s41239-023-00392-8.

(Efron & Tibshirani, 1993) Efron, B., & Tibshirani, R. J. (1993). *An introduction to the bootstrap*. Chapman & Hall.

(Eysenck et al., 2007) Eysenck, M. W., Derakshan, N., Santos, R., & Calvo, M. G. (2007). Anxiety and cognitive performance: Attentional control theory. *Emotion*, *7*(2), 336–353. https://doi.org/10.1037/1528-3542.7.2.336.

(Gollwitzer, 1990) Gollwitzer, P. M. (1990). Action phases and mind-sets. In E. T. Higgins & R. M. Sorrentino (Eds.), *Handbook of motivation and cognition* (Vol. 2, pp. 53–92). Guilford Press.

(Greenhalgh & Rosenblatt, 1984) Greenhalgh, L., & Rosenblatt, Z. (1984). Job insecurity: Toward conceptual clarity. *Academy of Management Review*, *9*(3), 438–448. https://doi.org/10.2307/258284.

(Hayes, 2018) Hayes, A. F. (2018). *Introduction to mediation, moderation, and conditional process analysis* (2nd ed.). Guilford Press.

(Heckhausen & Gollwitzer, 1987) Heckhausen, H., & Gollwitzer, P. M. (1987). Thought contents and cognitive functioning in motivational versus volitional states of mind. *Motivation and Emotion*, *11*(2), 101–120. https://doi.org/10.1007/BF00992338.

(Henseler et al., 2015) Henseler, J., Ringle, C. M., & Sarstedt, M. (2015). A new criterion for assessing discriminant validity in variance-based structural equation modeling. *Journal of the Academy of Marketing Science*, *43*(1), 115–135. https://doi.org/10.1007/s11747-014-0403-8.

(Hobfoll, 1989) Hobfoll, S. E. (1989). Conservation of resources: A new attempt at conceptualizing stress. *American Psychologist*, *44*(3), 513–524. https://doi.org/10.1037/0003-066X.44.3.513.

(Kline, 2016) Kline, R. B. (2016). *Principles and practice of structural equation modeling* (4th ed.). Guilford Press.

(Lazarus & Folkman, 1984) Lazarus, R. S., & Folkman, S. (1984). *Stress, appraisal, and coping*. Springer.

(Lent et al., 1994) Lent, R. W., Brown, S. D., & Hackett, G. (1994). Toward a unifying social cognitive theory of career and academic interest, choice, and performance. *Journal of Vocational Behavior*, *45*(1), 79–122. https://doi.org/10.1006/jvbe.1994.1027.

(Ng et al., 2021) Ng, D. T. K., Leung, J. K. L., Chu, S. K. W., & Qiao, M. S. (2021). Conceptualizing AI literacy: An exploratory review. *Computers and Education: Artificial Intelligence*, *2*, 100041. https://doi.org/10.1016/j.caeai.2021.100041.

(Ouellette & Wood, 1998) Ouellette, J. A., & Wood, W. (1998). Habit and intention in everyday life: The multiple processes by which past behavior predicts future behavior. *Psychological Bulletin*, *124*(1), 54–74. https://doi.org/10.1037/0033-2909.124.1.54.

(Podsakoff et al., 2003) Podsakoff, P. M., MacKenzie, S. B., Lee, J. Y., & Podsakoff, N. P. (2003). Common method biases in behavioral research. *Journal of Applied Psychology*, *88*(5), 879–903. https://doi.org/10.1037/0021-9010.88.5.879.

(Ravselj et al., 2025) Ravselj, D., Kerzic, D., Tomazevic, N., Umek, L., Brezovar, N., Iahad, N. A., & Aristovnik, A. (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.

(Satchell & Dourish, 2009) Satchell, C., & Dourish, P. (2009). *Beyond the user: Use and non-use in HCI* (pp. 9–16). ACM. https://doi.org/10.1145/1738826.1738829.

(Savickas, 2005) Savickas, M. L. (2005). The theory and practice of career construction. In S. D. Brown & R. W. Lent (Eds.), *Career development and counseling* (pp. 42–70). Wiley.

(Savickas, 2013) Savickas, M. L. (2013). Career construction theory and practice. In R. W. Lent & S. D. Brown (Eds.), *Career development and counseling* (2nd ed., pp. 147–183). Wiley.

(Stumpf et al., 1983) Stumpf, S. A., Colarelli, S. M., & Hartman, K. (1983). Development of the Career Exploration Survey (CES). *Journal of Vocational Behavior*, *22*(2), 191–226. https://doi.org/10.1016/0001-8791(83)90028-3.

(Tlili et al., 2023) Tlili, A., Shehata, B., Adarkwah, M. A., Bozkurt, A., Hickey, D. T., Huang, R., & Agyemang, B. (2023). What if the devil is my guardian angel? ChatGPT as a case study of using chatbots in education. *Smart Learning Environments*, *10*(1), 15. https://doi.org/10.1186/s40561-023-00237-x.

(Tran et al., 2023) Tran, N. K., Welles, S. L., & Goldstein, N. D. (2023). Geolocation to identify online study-eligible gay, bisexual, and men who have sex with men in Philadelphia, Pennsylvania. *Epidemiology*, *34*(4), 462–466. https://doi.org/10.1097/EDE.0000000000001607.

(Wang et al., 2025) Wang, H., Kou, X., Zhao, D., & Gu, Y. (2025). The double-edged sword effect of AI disruption awareness on AI usage behavior. *Science and Technology Progress and Policy*, *42*(23), 1–11. (In Chinese). https://doi.org/10.6049/kjjbydc.D6202502010RJ.

(Zhou, 2024) Zhou, P. (2024). Make lectures match how we learn: The nonlinear teaching approach to economics. *Education Sciences*, *14*(5), 509. https://doi.org/10.3390/educsci14050509.

Supplementary Material

This supplementary table reports the results of a robustness check that isolates the contribution of the geographic-region controls. Because the two economic-belt dummies (central and western belts, reference = eastern belt) were added in a later revision step, we re-estimated every model under two specifications that differ only in whether those two dummies are included. Specification A holds the seven original dummy controls (gender, club management experience, STEM major, and four grade indicators). Specification B adds the two region dummies. All other details are identical across the two specifications: ordinary least squares, BCa bootstrap confidence intervals (seed = 42, 5000 resamples, 95%), and HC3 robust standard errors; N = 289.

Table S1. Focal coefficients under the two specifications

Term	Outcome	Without region b [BCa 95% CI]	With region b [BCa 95% CI]	Conclusion
H1 main effect	Environment exploration	0.202 [0.067, 0.342]	0.203 [0.069, 0.342]	Supported in both
H1 main effect	Self-exploration	0.161 [0.044, 0.292]	0.158 [0.042, 0.286]	Supported in both
H2 linear interaction	Environment exploration	-0.070 [-0.181, 0.049]	-0.070 [-0.181, 0.050]	Not supported in either
H2 linear interaction	Self-exploration	-0.101 [-0.201, 0.004]	-0.105 [-0.207, -0.004]	Marginal in both
H3 quadratic interaction	Environment exploration	0.097 [0.012, 0.187]	0.095 [0.012, 0.186]	Supported in both
H3 quadratic interaction	Self-exploration	0.120 [0.042, 0.192]	0.116 [0.038, 0.188]	Supported in both

Table S2. Diagnostics under the two specifications

Diagnostic	Without region	With region	Same conclusion?
HC3 robust SE, H3 quadratic (environment)	b = 0.097, p = .045	b = 0.095, p = .049	Yes
HC3 robust SE, H3 quadratic (self)	b = 0.120, p = .003	b = 0.116, p = .004	Yes
Excluding 18 high-awareness respondents, H3 (environment)	b = 0.092, p = .099	b = 0.093, p = .096	Yes
Excluding 18 high-awareness respondents, H3 (self)	b = 0.180, p < .001	b = 0.181, p < .001	Yes
CFA factor scores, H3 quadratic interaction (environment)	b = 0.191, p = .005	b = 0.184, p = .007	Yes
CFA factor scores, H3 quadratic interaction (self)	b = 0.232, p < .001	b = 0.220, p < .001	Yes
Simple slopes (low / mean / high awareness)	0.181 / -0.001 / 0.033 (env); 0.173 / -0.067 / -0.039 (self)	0.183 / 0.003 / 0.036 (env); 0.179 / -0.061 / -0.043 (self)	Yes
CFA fit, four factors (χ^2, df, CFI, TLI, RMSEA)	168.05, 98, .977, .972, .050	Identical	—
Cronbach's α (env / self / avoidant use / AI awareness)	.892 / .900 / .858 / .925	Identical	—
Cook's D maximum (env / self)	0.202 / 0.095	0.176 / 0.078	Yes
R^2 (env, Model 1 / Model 3)	.101 / .187	.105 / .189	—
R^2 (self, Model 1 / Model 3)	.067 / .158	.084 / .168	—

Reading. The two specifications agree on every substantive conclusion. H1 (the positive main effect) and H3 (the U-shaped quadratic moderation) hold for both outcomes in both specifications, with coefficients that are nearly identical. The only difference is the self-

exploration linear interaction (H2), which is marginal in both specifications: its BCa 95% interval moves from [-0.201, 0.004] without the region controls to [-0.207, -0.004] with them, a shift of -0.005 at the zero boundary. The region controls therefore absorb background variance without changing any substantive conclusion.