From Generated Content to Career Capability:

Chinese International Students' Use of GenAI for Career Guidance in UK Higher Education

Author: Ao Yu

St Catherine's College
University of Oxford

Master of Science in
Education (Digital and Social Change), 2026

Abstract

Grounded in Sen's capability approach, this study examines how Chinese international students enrolled in or recently graduated from taught master's programmes at UK universities use generative artificial intelligence (GenAI) as an informal career-guidance resource; how personal, social and institutional conditions shape its use; and how students perceive its influence on career capabilities and career-related actions. Reflexive thematic analysis of ten semi-structured interviews revealed that participants associated GenAI with task initiation, drafting, interview rehearsal and low-exposure exploration. However, its perceived contribution to substantive career capability remained bounded by immigration eligibility, recruitment timing, curriculum fit, employer recognition and vacancy supply. Crucially, the credibility of GenAI-generated career content depended on distributed verification across official sources, recruitment platforms, practitioners and students' own knowledge; unequal access to time, occupational knowledge, institutional portals and mobilisable contacts meant that participants faced unequal verification resources and burdens. The study develops conditional conversion and verification inequality as two study-specific concepts to help explain how a widely accessible interface can yield uneven practical value. The findings suggest that GenAI can serve as a preparatory component of hybrid career guidance when supported by destination-specific information, explicit verification processes and accountable human judgement.

Acknowledgements

At this moment, outside my window, trains continue to pass through Oxford Railway Station, carrying with them the sounds and stories of countless journeys. I know that this farewell is not easy. Yet, perhaps this is what makes every place, every encounter, and every moment of growth so meaningful.

I would like to express my deepest gratitude to my supervisor, Professor Jeremy Knox, for his invaluable guidance, patience, and encouragement throughout this journey.

I am grateful to my cohort and all the faculty members of the MSc Education (Digital and Social Change) programme, whose teaching and conversations have broadened my understanding of education, technology, and society.

My sincere appreciation goes to all the participants who generously shared their experiences, reflections, and perspectives with me. Without their openness and trust, this research would not have been possible. Their voices are at the heart of this dissertation.

I would also like to thank my family and friends for their unwavering love, support, and companionship throughout this journey. To my parents, thank you for always believing in me, encouraging me to explore the world, and providing me with the courage to pursue my dreams. To my friends, thank you for being there through moments of uncertainty, pressure, laughter, and joy.

Leaving often means turning away from familiar faces and places. Yet I have always believed that every farewell makes space for another encounter. Somewhere, sometime, our paths may cross again. The memories, friendships, and lessons I have gained from Oxford will continue to accompany me as I move towards the next chapter of my journey.

Finally, I would like to thank everyone who has taken the time to read this dissertation.

Thank you for everything…

Table of Contents

Abstract 2
Acknowledgements 3
List of Tables 5
List of Figures 6
Chapter 1: Introduction 7
Research Context and Problem 7
Research Gap 8
Research Aim and Questions 9
Conceptual Lens, Design and Structure 9
Chapter 2: Literature Review 11
Career Guidance in a Transnational Study-to-Work Transition 11
GenAI as a Career Resource 14
From Generated Output to Feasible Career Action 19
Theoretical Framework: The Capability Approach 20
Research Questions and Chapter Synthesis 25
Chapter 3: Methodology 27
Critical Realism 27
Research Design 28
Data Collection 29
Sampling and Recruitment 30
Data Analysis 33
Reflexivity 38
Ethical Considerations 39
Chapter Summary 39
Chapter 4: Findings and Discussion 40
Introduction and Analytical Model 40
Low-Friction Activation: Creating an Accessible First Step in Career Work 41
Co-Produced Personalisation: Tailoring Depended on Participant Knowledge 45
Distributed Verification: How Generated Content Acquired Situated Credibility 49
Perceived Influence on Career-Related Actions Under Persistent Constraint 54
Cross-Theme Synthesis and Direct Answers to the Research Questions 58
Chapter Conclusion 60
Chapter 5: Conclusion 61
Study Overview 61
Contributions and Implications 62
Limitations and Future Research 63
Reference List 66
Appendix A. Letter of the Research Ethical Approval 78
Appendix B. Participant Information Sheet 79
Appendix C. Participant Consent Form 83
Appendix D. Interview Guidance 85
Appendix E. Analytical Coding Framework and Cross-Theme Evidence Matrix 89
Appendix F. Cross-Case Evidence and Literature Contribution Matrix 91

List of Tables

Table 3.1. Participant Profiles 32
Table 3.2. Analytical Procedure and Evidential Decisions 33
Table 4.1. Analytical Themes, Mechanisms, Evidential Boundaries, and Research-Question Alignment 40
Table 4.2. Source Roles in Distributed Verification 50
Table E.1. Analytical Coding Framework and Evidence Labels 89
Table E.2. Case-by-Theme Evidence and Conversion Conditions 90
Table F.1. Cross-Case Evidence and Literature Contribution Matrix 91

List of Figures

Figure 3.1. Capability-informed framework for GenAI-generated career content 36

Chapter 1: Introduction

Research Context and Problem

International students form a substantial part of UK postgraduate education. In 2024/25, students with a non-EU permanent address accounted for 65% of full-time postgraduate enrolments. Across all levels and modes of study, students with a permanent address in China constituted the second-largest non-EU country cohort, with 143,200 enrolments (Higher Education Statistics Agency, 2026, Figures 8 and 11). Chinese international students on taught master's programmes at UK universities face a compressed, transnational study-to-work transition in which academic demands coincide with decisions about destination, sector, occupation, recruitment timing and immigration eligibility. The value of UK-acquired qualifications and experience also varies across labour markets and depends on how relevant gatekeepers recognise them (Tomlinson, 2017; Wang et al., 2025; Zhao & Cox, 2022).

These conditions create informational and interpretive guidance needs. Students require current information about vacancies, visa sponsorship and selection criteria; however, they must also judge how such information relates to their experience, values and feasible routes. Careers services and personal networks can support that judgement, yet destination fit, timing, disclosure costs and postgraduation access vary (Hillman, 2021; Wang et al., 2025). Career aims can also change during study and span employment in the UK, China or elsewhere, as well as familiar and cross-field transitions (Consoli, 2024; Donald et al., 2025; Xu, 2021; Zhang et al., 2024). The group is treated as a heterogeneous, bounded case, which should not be interpreted as representative of international students overall.

Against this background, generative artificial intelligence (GenAI) has become an accessible, on-demand resource for career exploration, application drafting and recruitment preparation. Conversational GenAI systems generate new content in response to prompts and allow iterative revision (Feuerriegel et al., 2024; Miao & Holmes, 2023). Career-focused studies report uses in brainstorming, role identification, professional writing and interview rehearsal (Flink et al., 2025; Han et al., 2025; Loidl et al., 2026; Sarıköse et al., 2025). While this immediacy can support unfinished or low-exposure questions, it does not provide the contextual judgement, professional accountability or institutional authority that practitioner-led guidance can offer (OECD, 2024a; Tanner & Samson, 2026).

However, nominally similar access does not ensure equivalent practical value. Generated content may be fluent yet inaccurate, poorly contextualised or infeasible: An application

statement may misrepresent experience, while an occupational route may conflict with curriculum requirements, immigration rules or employer criteria (Loidl et al., 2026; Soltani & Tomlinson, 2024; Wang et al., 2025). Students also differ in the knowledge, time, institutional access and relationships available to check such claims, even when interface access is nominally similar (Baek et al., 2024; Hargittai, 2002; Hillman, 2021; van Deursen & van Dijk, 2014). The central problem is one of conversion: Under what conditions does a widely accessible GenAI resource become credible guidance connected to a valued and feasible career option?

Research Gap

Existing research links GenAI to brainstorming, drafting, rehearsal, confidence and perceived employability (Flink et al., 2025; Han et al., 2025; Loidl et al., 2026; Sarıköse et al., 2025). However, much of this evidence originates from domestic undergraduate samples, with brief or bounded interactions and self-reported outcomes. Direct evidence on how Chinese international postgraduate students use GenAI in relation to UK and overseas labour-market plans remains limited (Pandya & Wang, 2024; Sarmurzin et al., 2026; Wang et al., 2025). Experimental findings also vary by task structure: Settings with observable feedback can support performance gains, whereas open-ended tasks place greater demands on domain knowledge and error detection (Brynjolfsson et al., 2025; Kim & Moon, 2024). The evidence therefore supports task assistance more readily than sustained decision-making or substantive career opportunities.

More importantly, a process gap remains between receiving generated content and taking career-related action. Surveys of access, trust or perceived value reveal little about what prompts doubt, which sources students consult or how they respond when sources disagree. An outcome gap also persists because confidence, polished language and task completion do not establish that a valued option has become realistically pursuable (Pandya & Wang, 2024; Sarmurzin et al., 2026).

Addressing these linked gaps requires evidence preserving the sequence and context of use: the initial need; the generated response; the sources consulted for assessment and subsequent action; delay or inactivity. Reconstructing this process can distinguish nominal access from effective conversion and show whether evaluative work falls on students with unequal resources. It can also clarify which claims require checking, who can provide authoritative judgement and whether that support is accessible.

Research Aim and Questions

This study aims to examine how Chinese international students currently enrolled in or recently graduated from UK-taught master's programmes use GenAI as an informal career-guidance resource, how personal, social, and institutional conversion factors shape that use, and how students perceive GenAI as influencing their career capabilities and career-related actions. It addresses three research questions:

RQ1. What career-guidance needs do students seek to address through generative AI, and what circumstances give rise to these needs?

RQ2. How do personal, social and institutional conversion factors shape students' interpretation, evaluation and use of GenAI-generated career content?

RQ3. In what ways, if any, do students perceive GenAI to influence their career capabilities and the career-related actions they take?

Together, these questions follow the analytical progression from career-guidance need through interpretation and evaluation to perceived influence on opportunity and action. The phrase 'if any' in RQ3 keeps stable capabilities and unchanged action patterns within the analytical scope and avoids presuming a beneficial effect.

Conceptual Lens, Design and Structure

To address the research questions, the study adopts Sen's capability approach as its primary conceptual and evaluative lens (Robeyns, 2017; Sen, 1999). The framework distinguishes access to a generated resource from substantive freedom to pursue a valued option and directs attention to the personal, social and institutional conversion factors mediating this relationship. In this study, a career capability is a career option that a student values and can realistically choose and pursue; clearer representation and preparatory action remain analytically distinct.

Building on the capability approach's account of conversion factors and digital-inequality research that distinguishes access from effective use (Hargittai, 2002; van Deursen & van Dijk, 2014), this study develops two study-specific sensitising concepts. Conditional conversion describes how knowledge, relationships and institutional conditions shape the practical value of similar GenAI outputs. Verification inequality describes unequal access to the resources and routes required for verification, together with unequal verification burdens in establishing their provenance, contextual fit, authenticity and feasibility.

The study used a small-scale, cross-sectional qualitative design based on ten semi-structured interviews with current students and recent graduates who had used GenAI for career-related purposes. Interviews were used to reconstruct situated use episodes, and a paired vignette completed by participants elicited reasoning about contrasting conditions. Reflexive thematic analysis identified patterns across cases; capability approach rules distinguished resources, action and opportunity; and critical realism bounded explanatory claims. The design prioritises process and variation, making no claims about population-level prevalence or causal effects.

In the following sections, the dissertation first develops the transnational employability context and evaluates evidence on GenAI career guidance before explaining the research design and evidential strategy. The findings and discussion are then integrated around low-friction activation, co-produced personalisation, distributed verification and preparatory action under persistent constraint. The conclusion synthesises the answers, contributions, limitations and implications.

Chapter 2: Literature Review

Introduction

Students use generative artificial intelligence (GenAI) to explore occupations, draft applications, and rehearse recruitment encounters. Studies report easier brainstorming and revision, together with gains in self-assessed task confidence and perceived employability (Flink et al., 2025; Loidl et al., 2026; Sarıköse et al., 2025). Reported gains in brainstorming, revision, confidence, and perceived employability support task initiation and practice, while accuracy, labour-market fit, and practical feasibility require separate evidence. Access to the sector knowledge, trusted relationships, and institutional support needed to check generated claims also varies (Baek et al., 2024; Hillman, 2021; Wang et al., 2025).

This chapter examines how GenAI-generated career content enters the career transition of Chinese international students enrolled in or recently graduated from UK-taught master's programmes The chapter first explains why opportunity in a transnational study-to-work transition depends on recognition, relationships, time, and institutional rules. It then evaluates evidence on GenAI career use by task, outcome, and evidential strength, identifies what remains unknown between generated output and feasible career action, explains the choice and use of the Capability Approach, and draws the argument into three research questions.

Career Guidance in a Transnational Study-to-Work Transition

Employability as a Contextual Process

Graduate employability depends on more than a list of skills. Research locates employability in the relationship among career resources, identity, social position, gatekeepers, and labour-market context (Akkermans et al., 2024; Delva et al., 2021; Holmes, 2013; Tomlinson, 2017). A qualification or skill acquires career value when it is presented credibly, recognised within a field, and connected to a route the graduate can pursue. A contextual account evaluates both graduate resources and the recognition and opportunities available in a particular labour market.

Evidence from Chinese students, alumni, employers, and UK educators shows that students and gatekeepers may interpret the same UK-acquired skills differently (Zhao & Cox, 2022). The multi-stakeholder design supports a claim of misaligned expectations, but its network-based recruitment and limited examination of employer practice provide little evidence that

any resource improves employment. The relevant point is recognition: skills matter only when an audience treats them as credible and relevant.

Research in other international-student settings supports the same relational account. Pham et al. (2024) show that human, social, cultural, identity, and psychological resources are mobilised together, although their Australian setting limits transferability to a one-year UK master's programme or a return-to-China transition. Soltani and Tomlinson (2024) show that professional language operates as participation in a field rather than as a decontextualised skill. While the studies do not observe later employment outcomes, they both suggest that resources require contextual mobilisation and recognition.

A UK systematic review identifies communication, work experience, networks, immigration conditions, and employer practices as recurring influences on international postgraduate employability, while finding little evidence tied to labour-market outcomes (Wang et al., 2025). Research on return to China specifies how those influences operate: returnees value transnational networks, communicative competence, and professional identity, but employment quality also depends on domestic job-search activity, employer interpretations of overseas education, and local institutions such as hukou rules (Gu & Schweisfurth, 2015; H. He et al., 2025; Kang, 2025; Liu & Ye, 2021). The designs range from self-report to field experiment and support distinct rather than pooled causal estimates. Nevertheless, they show that the career value of a UK master's degree is context-dependent, reflecting how students mobilise it and how employers and institutions recognise it.

Students' destinations and plans can also change during the programme. Studies of Chinese students in the UK identify varied educational, personal, mobility, and career aims (Consoli, 2024; Hattersley & Nicholson, 2024; Xu, 2021; Zhang et al., 2024), so nationality is an inadequate proxy for a single career intention. Diary evidence records changes in aspirations and career resources during international postgraduate study (Donald et al., 2025). Research on Korean doctoral mobility offers adjacent evidence that identity can be renegotiated across national academic settings, but its population and career stage differ from those examined here (Kim & Cho, 2022). Scholarship on career geography likewise shows that opportunity changes across spatial and institutional settings (Jöns et al., 2025). Career planning during a UK master's programme is transnational: students form and revise aims in relation to several national and professional settings.

Moving between national and professional settings creates a concrete interpretive task. Students must decide where a qualification will be recognised, which experience will carry weight, and how a possible route fits a longer career plan. Occupational information gains practical meaning through knowledge of the relevant audience, rules, and time horizon. Students need guidance on obtaining information and judging its significance for a specific decision.

Conditions That Shape Usable Career Guidance

The literature repeatedly identifies four conditions that shape whether career guidance is usable in this transition: professional language, access to knowledgeable people, time and legal or financial room to act, and destination-relevant institutional support. Professional language illustrates the first condition. Employers assess whether applicants can present their knowledge and experience according to the conventions of a field, with English proficiency forming one element (Soltani & Tomlinson, 2024). International-postgraduate studies place language alongside work experience, networks, visa rules, and employer expectations (Nachatar Singh et al., 2023; Pham et al., 2024; Wang et al., 2025). GenAI can help students phrase existing experience, but it cannot supply missing experience, and credible claims still require evidence beyond generated wording.

Knowledgeable people provide more than vacancy information. Alumni can explain unwritten recruitment conventions, employers can clarify whether a qualification is accepted, and advisers can challenge assumptions within a plan (Donald et al., 2025; Nguyen & Cheng, 2025; Zhang et al., 2024). These interactions help students identify what a role requires and which experiences can support a credible application. A GenAI response may help phrase a field-relevant claim, but only practitioners and gatekeepers can confirm local practice and confer recognition.

Time, money, and immigration status set a separate limit on action. Financial resources shape how long students can explore, where they can move, and which opportunities they can afford to pursue (OECD, 2024b; Xu, 2021). Visa rules, a one-year programme, and recruitment calendars can compress the period available for preparation and application (Donald et al., 2025; Nachatar Singh et al., 2023; Wang et al., 2025). Advice that is generally accurate may still be unusable for a student's intended destination or application deadline. Practical fit is part of guidance quality.

Careers services can supply professional challenge and contextual knowledge, but availability and fit are distinct. Huang and Turner (2018) found that Mainland Chinese students in UK universities valued self-understanding, confidence, and flexibility and were more aware of curriculum-embedded provision than central services. Later research likewise reports uneven demand for destination-specific sector knowledge, work-related learning, and language support (Hillman, 2021; Wang et al., 2025). Reviews find limited evidence linking attendance or satisfaction with stronger labour-market outcomes (Scandurra et al., 2024; Wang et al., 2025). The unresolved question is whether guidance helps a student identify a credible option and act within the relevant destination and timetable.

The combined limits of professional language, destination knowledge, time, legal eligibility, and service fit help explain why GenAI attracts students during a time-pressured transition. A conversational system lets a student ask an unfinished question, obtain an immediate draft, and revise it repeatedly between appointments (Flink et al., 2025; Loidl et al., 2026). It can organise information, draft language, and suggest sources to consult. Trusted relationships, legal eligibility, and employer recognition remain external sources of authority. Section 2.3 evaluates the tasks GenAI can support and the judgement required once content has been generated.

GenAI as a Career Resource

The Scope of GenAI Career Use

Digital career guidance includes websites, assessments, communication tools, data services, and AI-supported systems (OECD, 2024a). Research labelled as AI covers recommender systems, predictive models, and conversational agents, so findings from one system may not transfer readily to other forms of digital guidance (Pandya & Wang, 2024; Sarmurzin et al., 2026; Westman et al., 2021). GenAI differs because it composes content from the user's prompt. This review limits GenAI career use to content-generating conversational systems used for occupational exploration, application work, decision support, or recruitment rehearsal. This boundary focuses attention on the content that students must interpret once it has been generated.

The attraction of such systems is tied to the form of the interaction. Students can pose an unfinished question, request revisions, and move between exploration and drafting in one exchange. Studies of higher-education use record value in idea generation, writing support, and rapid responses (Baek et al., 2024; Chan & Hu, 2023; A. J. He et al., 2025; Holland &

Ciachir, 2025). Career-focused studies add occupational searches, role identification, application development, and interview practice (Flink et al., 2025; Han et al., 2025; Loidl et al., 2026; Sarıköse et al., 2025). The system can lower the threshold for starting career work. Whether it improves decision-making depends on the task and the outcome used to judge it.

Evidence From Bounded Career Tasks

Career-focused evidence is concentrated in occupational exploration, application drafting, interview rehearsal, and workplace communication (Flink et al., 2025; Han et al., 2025; Loidl et al., 2026; Sarıköse et al., 2025). Studies report task initiation, confidence, perceived employability, and engagement (Han et al., 2025; Loidl et al., 2026; Sarıköse et al., 2025). Yang et al. (2026) add an important qualification: AI reduced information deficits in their brief comparison, while human counselling produced stronger decision self-efficacy and decision-making efficiency. These divergent outcomes require assessment by delivery mode, task, and outcome rather than a single category of benefit.

Most studies use domestic undergraduate samples, specific disciplines, brief encounters, and self-reported outcomes (Flink et al., 2025; Han et al., 2025; Loidl et al., 2026; Sarıköse et al., 2025; Yang et al., 2026). Han et al.'s (2025) small experiment with a career chatbot shows how motivational prompts affected reported career difficulty during a structured interaction. Loidl et al.'s (2026) single-institution study identified the career tasks and concerns reported by participants. Neither study verifies output accuracy or subsequent employment outcomes. The evidence supports bounded task assistance; sustained decision-making, application success, and transferability to Chinese international master's students in the UK remain unexamined.

Evidence on whether GenAI narrows or widens performance differences depends on task structure. A large workplace rollout found larger productivity gains among novice customer-support agents within a monitored system supplied with organisational data and feedback (Brynjolfsson et al., 2025). By contrast, a randomised university experiment found stable general reading performance, a fall in mathematics performance, and greater difficulty among lower-skilled participants in identifying plausible errors (Kim & Moon, 2024). The first setting made performance observable and provided support for checking it; the second required users to detect errors in the domain in which their competence was being supplemented.

The contrast suggests a proposition for examination. Bounded tasks with observable feedback and institutional verification may support more even gains, whereas open-ended tasks may place greater demands on domain knowledge and checking. Transnational career guidance often involves open-ended questions of contextual fit, while drafting and rehearsal can form more bounded tasks. The study examines how task structure and verification support shaped participants' use.

The next section examines the judgement required when GenAI-generated career content informs an open-ended, high-stakes decision.

Judgement, Trust, and Authenticity

The limits of task evidence direct attention to what happens after a response appears. Higher-education studies connect use and perceived value with language background, income, paid access, institutional rules, and study context (Baek et al., 2024; Chan & Hu, 2023; A. J. He et al., 2025; Holland & Ciachir, 2025). Students report benefits in access and expression, together with concerns about accuracy, privacy, integrity, dependence, and unequal access. The interface may be common, but the reasons for use and the resources available to verify it vary.

Digital-inequality research provides a useful analogy for verification inequality. Internet users who already had access still differed in their ability to find information, the time required to complete tasks, and the purposes for which they used the internet (Hargittai, 2002; van Deursen & van Dijk, 2014). Those studies predate GenAI and offer indirect evidence about who can verify a career claim, but they demonstrate why access and effective use require separate assessment. This dissertation specifies a narrower phenomenon: unequal access to the knowledge, relationships, time, and institutional authority needed to evaluate generated career content in context.

Trust measures identify another part of the problem. Cross-sectional survey evidence models an association between AI trust and protean career orientation through intrinsic motivation, with job insecurity moderating the reported association (Jiang & Chen, 2026). Other surveys examine whether students believe AI contributes to learning or career development (A. J. He et al., 2025). These associations identify dispositions linked with use, but these studies do not examine the act of checking a recommendation. A positive perception can coexist with weak verification.

Human-automation research provides a sharper evaluative distinction. Lee and See's (2004) interdisciplinary review treats reliance as appropriate when it is calibrated to the automation's capabilities, the task, and the use context, rather than when trust is simply high. Because its evidence base predates large language models and concerns mainly other automated systems, it offers limited insight into students' use of conversational GenAI. It nevertheless exposes a measurement problem in the career literature: a trust score is an incomplete indicator of warranted reliance for a particular claim. The present study treats reported trust as a disposition and examines checking, revision, or rejection within a concrete episode.

Fluent output makes miscalibrated reliance difficult to detect. A preregistered experiment found gains in creative problem-solving and self-efficacy after ChatGPT use, while self-evaluation remained an imprecise guide to performance (Urban et al., 2024). In another experiment, participants followed AI-labelled advice when it conflicted with contextual information and their own assessment (Klingbeil et al., 2024). Both studies operated outside career decision-making settings, so they identify mechanisms that may arise in career use rather than a measured level of career harm. Their shared implication is precise: coherent language and user confidence provide limited evidence about whether advice is sound. Verification remains part of the career task.

Model-level and behavioural studies support the same conclusion. Farquhar et al. (2024) demonstrate that semantic uncertainty can identify a particular class of confabulation, while Kalai et al. (2026) show that evaluation incentives can reward guessing rather than abstention. Neither study establishes the error rate in career guidance. Steyvers et al. (2025) found that people overestimated the accuracy of default LLM explanations and that longer explanations increased confidence while correctness remained stable. Fluency and detail are weak credibility cues; students still need claim-specific checks against occupational, institutional, or biographical evidence.

Authenticity creates a second test. Baek et al. (2024) found that students who spoke English as an additional language reported frequent ChatGPT use for writing. In career-related contexts, students described ChatGPT as a tool for revising professional materials and rehearsing communication, while raising concerns about authenticity and overreliance (Loidl et al., 2026). Professional communication requires a claim grounded in the student's experience and accepted within the relevant field (Soltani & Tomlinson, 2024). A polished statement can hide weak evidence or generate claims that extend beyond the applicant's defensible experience. Accuracy asks whether the statement is true. Authenticity asks

whether it represents the applicant's experience. Career use requires both forms of judgement.

Human Support and Responsibility

Human guidance varies in authority and relevance. Career advisers and industry professionals can provide career-related support and labour-market information, while employers and alumni can clarify recruitment requirements and local conventions (Donald et al., 2025; Nguyen & Cheng, 2025; Zhao & Cox, 2022). Career counsellors can also help students interpret AI-generated information and potential pitfalls, but this contribution depends on practitioner competence, institutional remit, time, and student access (OECD, 2024b). International students report uneven relevance in careers provision, and intervention reviews find weak links between service participation and labour-market outcomes (Hillman, 2021; Scandurra et al., 2024; Wang et al., 2025). Yang et al. (2026) found changes in career decision self-efficacy and decision-making efficiency during brief human counselling. The short trial left advice accuracy, durability of change, and employment effects outside its design. Human input is thus a situated career resource whose contribution depends on who provides it, what authority that person holds, and whether support is available at the decision point.

Hybrid provision also has institutional prerequisites. Kettunen and Sampson's (2019) phenomenographic study drew on 52 career-development experts from 16 countries and identified four progressively broader implementation problems: access to technology, access to current information, practitioner competence, and integration with existing services. Expert conceptions gathered before widespread GenAI use offer only indirect evidence of student benefit. They nevertheless show that hybrid guidance is an organisational arrangement involving technology, information, practitioner competence, and service integration. Where information, competence, or integration is weak, the human layer may provide less verification than a hybrid model assumes.

Accountable people and sources remain essential for factual quality and contextual verification. Barger's (2025) Wizard-of-Oz study found that participants reported a working alliance with a system presented as AI, although the design tested source framing rather than autonomous model competence. Accordingly, engagement alone provides limited evidence of factual quality or effective error correction. Reviews likewise identify opacity, weak empathy, and unclear responsibility in AI career support (Pandya & Wang, 2024; Sarmurzin

et al., 2026; Westman et al., 2021). Hybrid guidance requires an explicit allocation of responsibility for checking and accountability; reported trust offers an incomplete indicator of warranted reliance.

Policy and sector guidance recommend hybrid provision in which students test GenAI-generated career content against official information and human expertise (Miao & Holmes, 2023; OECD, 2024a; Tanner & Samson, 2026). Recent reviews identify the allocation and evaluation of those roles as an unresolved gap (Sarmurzin et al., 2026; Tanner & Samson, 2026), while international-student research documents unequal access to tailored support and professional networks (Hillman, 2021; Wang et al., 2025). The key issue is who checks which claim, whose judgement carries authority, and who remains accountable when human resources are uneven. Hybrid guidance remains an arrangement for empirical examination whose value depends on the allocation of authority and access. These unresolved questions of checking authority and accountability lead to the research gaps considered next.

From Generated Output to Feasible Career Action

The bodies of literature reviewed here identify a shared gap in understanding how GenAI-generated career content is interpreted within an international postgraduate transition. Task studies identify support within defined encounters, while adjacent experiments identify risks in self-evaluation and reliance. International-student research supplies the recognition, relationships, time, and rules that shape career opportunity (Donald et al., 2025; Flink et al., 2025; Klingbeil et al., 2024; Urban et al., 2024; Wang et al., 2025). Three linked gaps concern the process, outcome, and population that an adequate account must integrate.

Interpretation and Verification

The process gap concerns what happens between receipt of GenAI-generated career content and career action. Surveys link use or trust with perception, motivation, and career orientation, while interventions measure information, confidence, anxiety, engagement, and short-term decision outcomes (A. J. He et al., 2025; Jiang & Chen, 2026; Sarıköse et al., 2025; Yang et al., 2026). Experimental and student-use studies identify miscalibrated self-evaluation, continued reliance despite doubt, error, dependence, and authenticity concerns (Baek et al., 2024; Holland & Ciachir, 2025; Klingbeil et al., 2024; Loidl et al., 2026; Urban et al., 2024). Current evidence leaves unresolved what prompts doubt, which sources students consult, how they assign authority, and how they respond when sources disagree. Tracing those decisions can explain acceptance, revision, delay, or rejection of a suggested route.

Career Opportunity as an Outcome

The outcome gap concerns the set of options from which feasible action can emerge. Career studies report access, satisfaction, confidence, perceived employability, motivation, intention, and short-term decision outcomes (Flink et al., 2025; A. J. He et al., 2025; Jiang & Chen, 2026; Loidl et al., 2026; Sarıköse et al., 2025; Yang et al., 2026), whereas international-student research shows that language practices, networks, time, and institutions affect the value of career resources (Donald et al., 2025; Nachatar Singh et al., 2023; Wang et al., 2025; Xu, 2021). AI career reviews have yet to connect language practices, networks, time, and institutional conditions to feasible options and subsequent action (Pandya & Wang, 2024; Sarmurzin et al., 2026). The gap calls for a framework that relates generated content, contextual conditions, opportunity, and action while preserving them as distinct outcomes.

Population and Context

The population gap is also an analytical opportunity. Research on Chinese students in UK higher education examines motivation, transition, career imagination, and the recognition of overseas study (Consoli, 2024; Hattersley & Nicholson, 2024; Xu, 2021; Zhang et al., 2024; Zhao & Cox, 2022). The literature reviewed here provides little direct evidence about how Chinese international master's students use GenAI in relation to UK and overseas labour-market plans (Pandya & Wang, 2024; Sarmurzin et al., 2026; Wang et al., 2025). Their varied aims, language practices, and prior education intersect with qualification recognition, immigration conditions, and a compressed programme calendar. These features make the group a revealing case for studying conversion: a suggestion can appear relevant at the occupational level yet prove infeasible because of sponsorship, recruitment timing, or employer recognition.

Theoretical Framework: The Capability Approach

Why the Capability Approach?

The study requires a framework that distinguishes access to a generated resource from a career option a student can genuinely choose and pursue. Previous studies such as Tomlinson's (2017) work on graduate-capital model offers a complementary account of employability resources: GenAI may provide occupational knowledge, suggested contacts, application conventions, a coherent professional narrative, or confidence (see also Pham et al., 2024). However, cataloguing capital can overstate practical value if possession or

presentation is treated as proof that a route is open. An AI-polished CV, a list of alumni, and greater confidence remain insufficient when sponsorship requirements, recruitment timing, or employer recognition constrain practical use.

Bourdieu's concepts of field, habitus, and capital is considered useful in explaining why the same overseas qualification, linguistic style, or network can be valued differently by UK and Chinese employers, and why classed histories shape which career futures feel natural or reachable (Delva et al., 2021; Xu, 2021). While this perspective exposes how advantage can be reproduced beneath apparently individual differences in prompting skill, it is less suited to the study's main evaluative task of distinguishing among possessing a resource, having substantive freedom, taking action, and making a valued refusal. Yet, Bourdieu helps explain how position and recognition affect the conversion of a resource into an opportunity.

For Sen (1999), capability is a person's real freedom to choose among and achieve the beings and doings they have reason to value. Resources are the goods, services and other means available to a person; functionings are the states and actions the person achieves. Robeyns (2017) locates conversion factors between resources and capabilities, drawing attention to the conditions that shape what a person can do with a resource. The study evaluates the valued and feasible career options that emerge when generated content is mediated by personal, social, and institutional conversion conditions, while treating AI-assisted capital and field position as inputs to conversion. Evaluating valued and feasible options directly addresses the preceding discussion of the gap between generated output and feasible career action.

Capability-based career-guidance research examines informed choice, sustainable work, justice, and well-being (Egdell & Robertson, 2021; Fuertes et al., 2021; Robertson & Picard, 2021). It asks whether people can access work and career paths they have reason to value, with the information and support required for informed choice. Applied to GenAI, the question becomes concrete: which option becomes available, to whom, under which conditions, and with what room for informed refusal?

Resources, Capability Sets, and Agency

A GenAI response is treated as a career resource (Sen, 1999; Robeyns, 2017). It may contain occupational information, a suggested route, draft language, or a rehearsal prompt. Its characteristics include specificity, currency, factual quality, source traceability, language form, and revision affordances. Conversational speed and revision can support exploration, while prompt sensitivity, opaque provenance, and error create work for the user (Feuerriegel

et al., 2024; Miao & Holmes, 2023). These characteristics belong to the resource; conversion factors describe the conditions through which a student can turn it into an opportunity.

A capability set comprises the valued and feasible options open to a person (Sen, 1999; Robeyns, 2017). In this dissertation, career capabilities refer to career options that students value and can realistically choose and pursue within a specific context (Robertson & Picard, 2021). They include genuine opportunities to explore a credible occupation, present experience with integrity, seek informed support, apply for a suitable role, or reject content that conflicts with a valued aim. Clearer understanding can support informed choice, but it remains a representational change until feasibility is demonstrated. Confidence and skill can contribute to career capabilities, while contextual feasibility determines whether an option is genuinely available.

Functionings are achieved states or actions realised from a capability set (Sen, 1999; Robeyns, 2017). Career-related actions such as submitting an application, contacting an alumnus, changing a target role, seeking further evidence, or rejecting a route can provide evidence of functionings when participants link them to valued career aims. Similar actions can arise from different option sets: an application may follow informed judgement or an unchecked claim, while delay may reflect constraint or an effort to preserve future choice. Participant accounts reveal the alternatives and reasons surrounding reported action.

Functionings provide partial evidence about how a student shaped the underlying aim and responded to constraint. Agency concerns the student's role in forming aims, evaluating options, and acting under constraint (Sen, 1999; Robeyns, 2017). Adaptive preferences may arise when sustained constraint reshapes what a person regards as attainable or worth seeking. The analysis takes participants' stated values as its starting point and considers adaptive preference formation by examining how an aim was formed, which alternatives were known, and whether revision reflected informed reconsideration or constrained access to alternative possibilities (Egdell & Robertson, 2021; Robeyns, 2017; Xu, 2021). This focus matters for international postgraduates whose career aims and opportunities can evolve over time and are shaped by access to resources, institutional support, and visa restrictions (Donald et al., 2025; Fakunle, 2021; Nguyen & Cheng, 2025; Xu, 2021).

Conversion Factors and the Analytical Model

Personal conversion factors shape whether a response becomes intelligible (Robeyns, 2017). AI literacy includes understanding, use, evaluation, and ethical judgement, while career knowledge, language practice, prior experience, trust, and motivation affect what a student notices in an answer (Jiang & Chen, 2026; Klingbeil et al., 2024; Loidl et al., 2026; Ng et al., 2021; Soltani & Tomlinson, 2024). The analysis will ask how personal knowledge, language, experience, trust, and motivation shape interpretation while recognising that responsibility for system error and structural constraints does not rest solely with the student.

Social conversion factors shape interpretation and recognition (Robeyns, 2017). Parents and peers can influence career aims, while alumni, advisers, employers, and professional communities can test factual claims, occupational fit, and the value of prior experience (Consoli, 2024; Donald et al., 2025; Nachatar Singh et al., 2023; Nguyen & Cheng, 2025; Pham et al., 2024; Zhang et al., 2024; Zhao & Cox, 2022). These relationships situate content within a field of practice and help the student judge which evidence and identity claims will be accepted. Verification is a social process as well as an information check.

Institutional and environmental conversion factors include careers provision, university rules, immigration conditions, employer practice, cost, time, and labour-market conditions (Baek et al., 2024; Hillman, 2021; Jöns et al., 2025; OECD, 2024a, 2024b; Scandurra et al., 2024; Wang et al., 2025; Xu, 2021). Platform design, model version, and output properties are treated separately as characteristics of the GenAI resource. An access tier or platform charge is coded as an institutional conversion condition only where the account links it to feasible use. GenAI may make a route visible during self-directed or classroom career work (Flink et al., 2025; Loidl et al., 2026). A route forms part of the student's capability set when it is valued and feasible under the relevant conditions.

Personal, social, and institutional conversion factors work in combination (Robeyns, 2017). A claim may become intelligible through language support, credible through contact with an alumnus, and feasible after a visa or recruitment rule is checked. The analysis asks three explicit questions: which response characteristic affected understanding, which person or document established credibility, and which rule or resource affected feasibility and action. These questions locate verification within the conversion process.

The concepts are sensitising rather than deterministic (Egdell & Robertson, 2021; Robeyns, 2017). GenAI use can involve generation, interpretation, checking, revision, action, delay, or

rejection in different orders, with pauses, reversals, and abandoned routes. The analysis can compare cases in which conversion appeared to proceed, stall, or change direction within a flexible sequence.

Conceptual Scope and Limits of the Framework

The first challenge is operationalisation. A capability set is counterfactual because it encompasses genuine alternatives, including unchosen options, and interviews provide indirect evidence of that set. Translating capability into indicators can create false precision or collapse freedom back into resources and outcomes (Burchardt & Vizard, 2011; Robeyns, 2017; Sugden, 1993). The framework therefore better suited to asking whether particular valued options are genuinely open than to reducing career freedom to an observed outcome.

The second challenge is selecting which capabilities count. Nussbaum (2011) defends a universal list of central capabilities and threshold entitlements as a basis for justice, whereas Sen (2004) resists fixing one list independently of purpose and public reason. A universal list would give the present capability analysis an explicit normative floor, but could pre-empt the career aims of participants whose valued lives and destinations differ. An entirely participant-led list creates the opposite risk: a reported preference may reflect limited information, repeated exclusion, or an option that has remained outside perceived reach. The study adopts a middle position. It limits the evaluative domain to career-related opportunities relevant to the research questions, begins with participants' stated values, and examines how those values were formed, what alternatives were known, and whether institutional constraints narrowed what could reasonably be imagined.

The third challenge is circularity. Inferring capability from functioning would treat an application as proof that applying was freely available; inferring it from confidence would treat a persuasive answer as proof of expanded freedom. The same application may follow informed choice, pressure, or an unchecked claim, while inactivity may represent exclusion, strategic delay, or informed refusal (Egdell & Robertson, 2021; Robeyns, 2017). Accordingly, this study does not automatically treat confidence, clearer understanding, or career-related action as evidence of capability expansion. These may be important outcomes, but they do not by themselves show that a valued career option has become realistically available under relevant conditions.

Evidence of real freedom requires more than resources or observed outcomes: a person may understand a route, act on advice or feel more confident while the route remains unavailable, unacceptable or unsustainable. A valued opportunity can also exist without immediate action. Capability language therefore remains provisional when feasibility, choice or the relevant structural conditions are unclear.

A fourth challenge is the individualising risk. Because the individual is the unit of evaluation, a capability analysis can inadvertently recast unsuccessful conversion as a deficit in AI literacy, confidence, or motivation. Career-guidance scholars warn that employability interventions can demand individual adaptation while institutional barriers remain intact (Egdell & Robertson, 2021; Fuertes et al., 2021). The analysis separates platform design as a GenAI resource characteristic from personal, social, and institutional conversion conditions. Institutional conditions include access tier or cost, careers services, recruitment calendars, visa and hukou rules, and employer recognition; social conditions include access to knowledgeable people. Responsibility for fabricated content and barriers to verification remains with the relevant platform and institutional arrangements rather than the participant alone.

Finally, conversion factors interact in open systems, so the framework is not a causal model designed to produce stable independent effects. Language knowledge may make a response intelligible, an alumnus may make it credible, and an employer or immigration rule may still leave the route infeasible. It serves as an evaluative and explanatory lens for considering how personal, social and institutional conditions can enable, constrain or redirect the conversion of resources into valued opportunities. These limits frame the research questions developed next.

Research Questions and Chapter Synthesis

The review identifies three connected problems for understanding GenAI use in the career transition of Chinese international master's students. First, career-guidance needs arise within a transnational setting in which professional recognition, destination-specific information, relationships, time and institutional rules shape what students can realistically pursue. Second, existing GenAI research provides evidence of support for bounded tasks but offers less insight into how students interpret, verify, revise or reject generated career content under different personal, social and institutional conditions. Third, commonly reported outcomes

such as confidence, engagement and task completion do not establish whether a valued career option has become more realistically pursuable.

The Capability Approach connects these problems by treating GenAI-generated career content as a resource whose practical value depends on its conversion under particular personal, social and institutional conditions. This framing leads to three research questions:
RQ1. What career-guidance needs do students seek to address through generative AI, and what circumstances give rise to these needs?
RQ2. How do personal, social and institutional conversion factors shape students' interpretation, evaluation and use of GenAI-generated career content?
RQ3. In what ways, if any, do students perceive GenAI to influence their career capabilities and the career-related actions they take?

Together, the questions move from the circumstances prompting GenAI use, through the processes of interpretation and evaluation, to perceived influence on action and opportunity. RQ1 establishes the needs and situations in which GenAI enters career guidance; RQ2 examines how generated content is converted, checked and used; and RQ3 evaluates the outcome claim by distinguishing preparatory or representational change from perceived change in a valued and realistically pursuable career option. These questions provide the analytical structure for the research design and evidence rules developed in Chapter 3.

Chapter 3: Methodology

Introduction

This chapter explains the study's research philosophy, design, data collection, sampling, analysis, reflexivity and ethics. Each methodological choice serves one of the research questions, leading to the claims to be presented in Chapter 4. General personal assessments, reported personal episodes, and hypothetical-vignette reasoning are distinguished to preserve the boundary between interpretation and reported behaviour.

Critical Realism

The study is underpinned by a critical realist perspective, accepting that there is a reality independent of knowledge, but that every account of that reality is partial and situated, and therefore meaningful – but partial – evidence of participants' experiences and interpretations (see Bhaskar, 1978; Fletcher, 2017; Maxwell, 2012; Sayer, 2000; Stutchbury, 2022). The interview transcripts have been interpreted as empirical evidence of (reported) experience and action, distinct from the actual events and the structures or causal powers that contribute to them; surrounding conditions remain inferential and require assessment against the available evidence.

Consistent with this critical-realist position, social life was treated as an open system in which several conditions may strengthen, weaken or block one another (Danermark et al., 2002; Sayer, 2000). GenAI may support similar tasks differently, as usability depends on students' knowledge, trusted support, and labour-market rules. The small sample supports analysis of conditional variation but cannot isolate single causes or support causal comparison. This open-system account aligns with the study's aim of explaining why participants associated similar guidance with action in some cases and with difficulty in others.

Abduction and retroduction – interpreting patterns through concepts and seeking concepts that make patterns possible (Danermark et al., 2002; Fletcher, 2017) – have been applied to advance explicit, challengeable explanations through case comparison, disconfirming evidence, and rival interpretations. There is a risk of overstating certainty, however (Kemp, 2005; Cruickshank, 2010; Stutchbury, 2022), which the study mitigates by separating participant report, reported action, and researcher inference, grounding explanations in identifiable evidence, and bounding claims to the sample circumstances. The choices made

here support context-specific explanations of variation while treating mechanisms as provisional.

The claims are, accordingly, deemed descriptive when they report what a participant said, perceived or did; interpretive when they identify patterns across accounts; explanatory when they propose conditions that could account for patterned variation. Contextual quotations and episode summaries support descriptive claims, cross-case comparison supports interpretive claims, and convergence across episodes, disconfirming cases and rival explanations supports explanatory claims. The analysis considers reported experience to be more reliable than inferred conditions. Terms such as shaped, enabled, constrained and made easier are used to express conditional contribution.

Research Design

The study used a small-scale, single-site, cross-sectional qualitative design based on ten individual semi-structured interviews. Following Kallio et al. (2016), semi-structured interviews were chosen to ensure coverage of the central topics while leaving room for unanticipated experiences and explanations.

Interviewees were interviewed individually to allow detailed reconstruction of personal episodes and to support privacy and confidentiality when discussing family- or immigration-related circumstances. Interviewees were asked about the same topics to enable subsequent comparison, and follow-up probes were chosen to clarify each participant's decisions in context.

The main unit of analysis was a situated GenAI use episode: a concrete occasion when a participant needed career guidance, considered available sources, used GenAI, evaluated the response and acted, delayed or remained inactive. RQ1 was addressed mainly through reported needs and their circumstances; RQ2 through conversion factors linked to interpretation, evaluation, and use; and RQ3 through participant-perceived influence on career capabilities and career-related actions. Drawing on Sen (1999) and Robeyns (2017), clearer understanding was coded as representational change unless feasibility was evidenced, while a completed action that the participant linked to a valued career aim was coded as functioning evidence. The ten-case personal-episode analysis included two pilot interviews; paired-vignette material was analysed separately from reported experience. The four-week fieldwork window was used to reduce temporal variation because GenAI tools and graduate-recruitment conditions can change quickly. The design supports short-term retrospective

interpretation, whereas a study of sustained career outcomes would require longitudinal evidence (Hermanowicz, 2013).

The design used a pre-prepared fictional mini-story (the vignette) as a common prompt across interviews. The vignette was chosen to elicit decision reasoning that general questions might miss and to allow sensitive circumstances to be considered through fictional characters (Bradbury-Jones et al., 2014; Hughes, 1998; Jenkins et al., 2010). Yilin and Mei were configured as contrasting packages of support and constraint across English confidence, career clarity, family expectations, access to career support, time before application deadlines and sponsorship arrangements, while the GenAI response was held constant. The vignette was used to examine reasoning about combinations rather than to estimate independent effects or actual behaviour (Hughes & Huby, 2004; Jenkins et al., 2010).

Episode reconstruction was selected because a general attitude question would merge the stages at which guidance can succeed or fail. Each account traced the trigger for use, prompt context, generated content, evaluation, supplementary sources, decision and subsequent action or inactivity. Follow-up questions inquired about timing and asked how the participant might have proceeded through another source. This sequence allowed the analysis to distinguish immediate production help from stronger claims about option feasibility or career capability. Conducting the interviews individually also protected participants' privacy around immigration, family and job-search concerns. Overall, the design prioritised depth and process over prevalence. Cross-sectional interviews offered the possibility of recovering participants' recent interpretations; behavioural logs, employer assessments and later outcomes were deemed beyond this study's scope.

Data Collection

The interview guide covered career context, the need that prompted GenAI use, one concrete use episode, evaluation and checking, action, delay or inactivity, the fictional case, and a final opportunity to add or correct relevant points. Questions used everyday language and requested concrete examples. Capability Approach terminology was not used in the interviews and was applied only during analysis to avoid ambiguity between technical and everyday meanings. The complete guide and paired-vignette script are provided in the appendices below.

Two pilot interviews were conducted under the same eligibility, consent and core elicitation procedures. Piloting produced three focused types of revision: question order was adjusted, ambiguous wording was simplified, and vignette probes were revised to distinguish possible, advisable and personally comparable action (Malmqvist et al., 2019). The research questions, eligibility criteria and personal-episode evidence remained stable. One pilot preceded the final vignette and contributed personal-episode material to the ten-case analysis while remaining outside vignette comparison. A pilot memo documented revisions and supported comparison between early and later interviews.

Interviews were conducted on Microsoft Teams over four weeks and lasted approximately 45–60 minutes. Online interviewing was chosen to reduce travel and scheduling barriers but relied on participants having a private Internet connection and at the cost of conveying fewer embodied cues (see Archibald et al., 2019). Participants could pause, skip a question or turn off their camera at will and, when a participant's response was unclear, the interviewer asked the participant to explain or elaborate.

Following the personal account, the interviewer introduced the pre-prepared fictional vignette and asked participants to compare how Yilin and Mei might interpret, evaluate and use the same GenAI response. Follow-up probes distinguished what might be possible, what might be advisable and whether the fictional case resembled the participant's experience.

Participants could speak Standard Chinese (Mandarin), English or switch between them. Recordings were automatically transcribed and manually checked against audio; clear recognition errors were corrected, and each participant's files were de-identified as P1–P10. Where audio records were too uncertain to permit direct quotation, their content was paraphrased. Analysis was conducted in the source language whenever possible. Standard Chinese extracts selected for reporting were retained alongside English translations, and a translation memo recorded culturally specific terms and final wording choices (following Squires, 2009; Temple & Young, 2004; van Nes et al., 2010). Quotation edits followed a simple convention: fillers and immediate repetition were removed sparingly; ellipses marked substantive omissions; square brackets marked supplied wording or clarification; and light grammatical cleaning preserved meaning and tone.

Sampling and Recruitment

Following Palinkas et al. (2015), criterion-based purposive sampling was used to recruit Chinese international students aged 18 or over who were enrolled in or had recently

graduated from a taught master's programme at a UK university. All ten participants were either current students or recent graduates of taught-master's programmes at the same UK university. The study was therefore single-site, although the institution was a contextual setting rather than the unit of analysis. Eligible participants had used GenAI at least once for career planning, job searching or professional development. These criteria ensured that every participant could describe the process examined by the research questions, while their shared institutional context supported meaningful comparison. A short screening form confirmed eligibility and monitored broad variation in discipline, career direction and reported GenAI use.

Recruitment notices circulated through postgraduate networks, social media and peer contacts, inviting interested students to contact the researcher directly. Consistent with Noy's (2008) account of chain referral, peer contacts extended recruitment beyond the researcher's immediate network. As the recruited sample developed, it was reviewed for disciplinary, career-direction and use variation. The sample is self-selected: every participant was already a GenAI user, and volunteers may have been more digitally confident, engaged with career planning or willing to discuss AI than other eligible students. The claims of this study are limited to variation among participating users and give less coverage to former users, people outside the active user sample, and students with restricted digital access.

Sampling sought explanatory contrasts within a bounded population rather than demographic representation. As the sample developed, it was reviewed for discipline, stage, target location, career direction and patterns of GenAI use, because these features could alter information needs and access to checking resources—that is, careers portals, occupational knowledge and contacts used to test generated claims. Enough variation was sought to enable comparisons such as continued access versus lost access to a careers portal, established versus limited occupational knowledge, and mobilisable versus difficult-to-activate contacts. Ten cases provided recurring evidence for central processes and enough difference to challenge simple user-skill explanations. Recruitment closed at the end of the four-week fieldwork window, after ten eligible interviews had been completed. The shared institutional context supported comparison but limits transferability to other UK universities and careers-service environments. Broader analytic transferability rests on readers' assessment of the reported context and mechanisms. In sum, the study offers analytic insights for similar transnational postgraduate settings, whereas prevalence estimates and subgroup comparisons would require a larger, purpose-designed sample.

The final analytic sample comprised ten interviews. Sample adequacy was judged through information power rather than an unsupported claim of saturation. As Malterud et al. (2016, p. 1754) state, "The larger information power the sample holds, the lower *N* is needed, and vice versa". The narrow aim, specific eligibility criteria, theoretically focused but open interview guide and dialogues supported intensive within- and cross-case analysis. Ten cases are adequate for the bounded explanatory task undertaken here, whereas estimating the distribution of experiences in the wider student population would require a larger representative sample. The self-selection constraint noted above remains part of that adequacy judgement, rather than being used only to motivate future sampling.

Table 3.1

Participant Profiles

ID	Broad disciplinary area	Broad career direction/Location	Main reported GenAI use
P1	Health and life sciences	Pharmaceutical or health-related work; geographically open	Exploration, applications and skills
P2	Comparative literature	Marketing, commercial or management roles; China	CV and application editing
P3	Education	Academic research or international policy; location undecided	Exploration and interview preparation
P4	Mathematical and computing sciences	Computing or AI-related technical work; geographically open	Learning, applications and interviews
P5	Environmental and social sciences	Sustainability, data or project work; UK and China	Application materials and option review
P6	Humanities and languages	Consulting or corporate graduate roles; China	Direction, applications and interviews
P7	Business and finance	Finance or investment-banking-related roles; China	Exploration, applications and interviews
P8	Business and management	Marketing and games-sector roles; UK preferred, China considered	CV, planning and interview preparation
P9	Quantitative and STEM-related studies	Data and analytics roles; UK	Job matching, research and applications
P10	Economics and finance	Financial-markets sales and trading; China	Industry research and exploration

Note. *Participants were current students or recent graduates of UK taught-master's programmes seeking graduate employment. Categories have been broadened to protect privacy.*

Data Analysis

The transcripts were examined through reflexive thematic analysis, with the Capability Approach supplying decision rules and critical realism guiding the later comparison of explanations. Reflexive thematic analysis was chosen for its ability to organise patterned meanings across cases while recognising that themes are produced through interpretation rather than discovered as fixed facts (as per Braun & Clarke, 2006, 2021). The Capability Approach distinguished generated resources, valued and feasible options, and actions linked to valued aims; critical realism supported cautious comparison of why reported outcomes differed. The three approaches informed interpretation rather than serving as a predetermined code list.

Analysis began in the language used by each participant. Recordings and transcripts were checked, each transcript was read as a whole, and a case memo written to summarise career context and GenAI use. Material was then tagged as general personal assessment (GPA), personal episode (PE) or vignette reasoning (VR). The GPA captured a participant's interpretation of recurring practice; a PE captured a reported event or action in a specific situation; and VR captured hypothetical reasoning about the fictional case. The three evidence types were kept separate to prevent a general attitude or hypothetical expectation from being reported as behaviour. Initial coding used concrete labels for needs, judgements and actions. Each personal episode was reconstructed as a sequence: guidance need, sources considered, GenAI response, interpretation, checking, adaptation and subsequent action, delay or inactivity. Coding remained open and data-attentive while recognising the theoretical influence on attention, because the research questions, interview guide and researcher's position had already shaped attention (see Braun & Clarke, 2021).

Table 3.2

Analytical Procedure and Evidential Decisions

Stage	What was done	Evidence rule
1. Familiarise	Checked transcripts against recordings, read each case in the source language and prepared a case memo.	Preserved context and recorded early questions or uncertainties.

Stage	What was done	Evidence rule
2. Open code	Coded general personal assessments (GPA), personal episodes (PE) and vignette reasoning (VR) separately with concrete, data-attentive labels.	Recognised theoretically informed interpretation and kept hypothetical material separate from reported behaviour.
3. Reconstruct episode	Traced need, support sources, GenAI response, interpretation, checking, adaptation and action, delay or inactivity within each case.	Kept each reported process connected to the circumstances in which it occurred.
4. Apply capability rules	Applied the resource, conversion-factor, capability and functioning decision rules to the open codes.	Required an explicitly valued and realistically pursuable option, participant-perceived change in realistic pursuability, and separate classification of representational change.
5. Compare cases and vignette	Coded condition → affected process → expected consequence, then compared cases in a matrix.	Kept vignette reasoning separate; examined configurations rather than isolated effects.
6. Review explanations	Developed themes, used abduction and retroduction, and checked negative cases and rival explanations.	Reported bounded explanations that remained proportionate to the interview evidence.

Analysis proceeded through linked within-case and cross-case stages. In the first stage, the complete transcript and case memo preserved chronology, values, constraints and the participant's own explanation. In the second, coding captured needs, resource characteristics, conversion conditions, checking practices, perceived influences and actions. In the third stage, the case-by-code matrix was constructed to compare similar episodes and to mark supporting, qualifying, and disconfirming evidence. Candidate themes were then tested for internal coherence, separation from adjacent themes and relevance to a research question. Capability labels were applied after this inductive and descriptive work, with a view to reducing the risk that theoretical categories would predetermine what counted as important. Analytic attention shifted iteratively between data, cases and concepts until each reported theme expressed a clear central claim and its limits. Table 3.2 records the procedural and evidential decisions used at these stages.

The final reporting structure followed the evidential demands of each question: RQ1 required recurring needs plus circumstances; RQ2 required an episode-level link between a condition and interpretation, evaluation or use; RQ3 required perceived influence on action or perceived change in realistic freedom to pursue a valued, feasible option. General assessments established participants' interpretations, personal episodes supplied the strongest evidence of sequence and action, and vignette responses suggested possible contextual processes. Claims were labelled according to the evidence type that they drew upon, in order to prevent a hypothetical comparison from standing in for experience and to prevent an isolated action from standing in for capability expansion. This is also why Chapter 4 sometimes reports a qualified conclusion even when participants described GenAI as useful.

The Capability Approach was applied in a second analytical pass using explicit decision rules. Access to GenAI and the receipt of generated content were classified as resources. Participant-perceived influence on career capability required an explicitly valued option, evidence that the option was realistically pursuable under relevant personal, social, and institutional conditions, and an account linking GenAI to a change in the participant's freedom or ability to pursue that option. Clearer understanding or a clearer account of a career route was coded as a change in how the route was understood and explained; evidence of feasibility was still required for a capability claim. A reported career-related action was treated as functioning evidence only when the participant reported completing it and linked it to a valued career aim; other actions remained reported behaviour, while inactivity remained an outcome for comparison. Conversion factors required an episode-level link to understanding, evaluation, checking, adaptation, or action. Practical familiarity with GenAI and awareness of its limitations were treated as episode-level indicators rather than independently measured competencies. Relevant evidence included prompting experience, recognition of unsupported claims, selection of an appropriate verification method, or evidence of ethical judgement. Model, interface, and output properties were treated as features of the GenAI resource; prompt language was treated as part of the interaction; and the access tier or cost was classified as institutional when it affected feasibility of use. These rules required different evidence for informative output, perceived influence on action and real opportunity.

The labels for resources, actions, functionings, conversion factors, and capabilities were assigned to specific episodes and checked against complete cases. A claim that an answer was useful remained classified as perceived usefulness unless the account identified the relevant

career option, established its feasibility, and described GenAI's perceived influence on the realistic pursuit of that option. Codes were revised when subsequent comparison indicated that an earlier label had overstated or overlooked part of the process. Theme development required a clear central idea, evidence across relevant cases, and an identifiable contribution to at least one research question. Episode-level coding and complete-case checking ensured that the theoretical concepts remained grounded in, and accountable to, the interview material.

Figure 3.1. Capability-informed framework for GenAI-generated career content

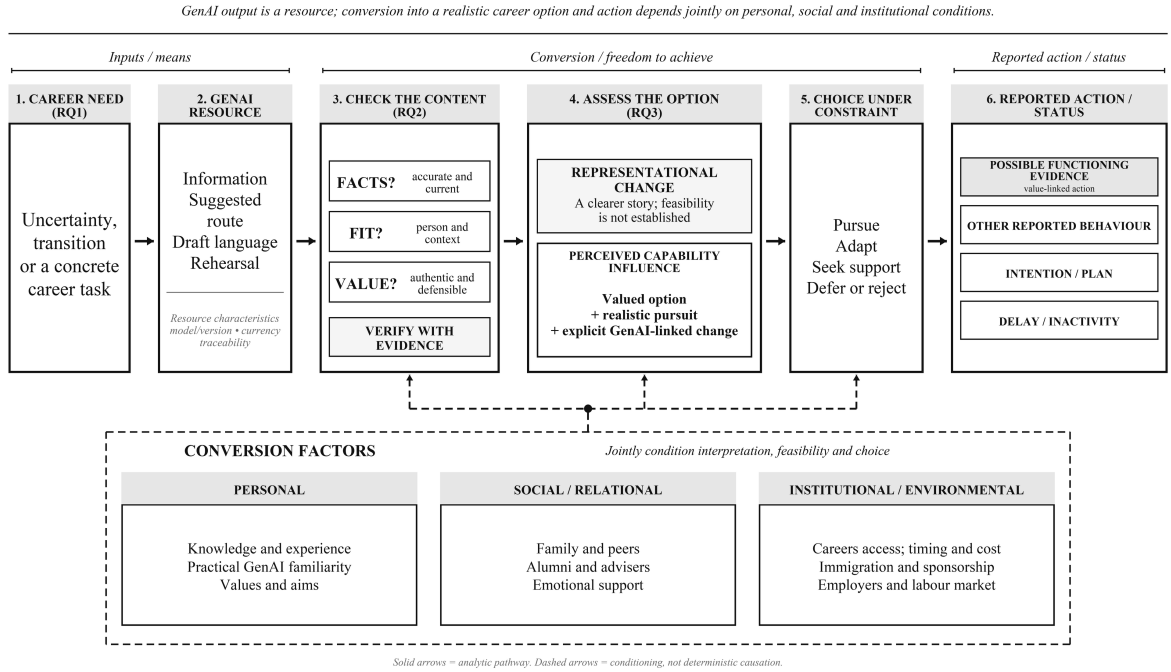

Note. *The figure adapts the Capability Approach's resource–conversion–capability–functioning logic to this study. Career-capability claims required a valued, realistically pursuable option and a participant-perceived GenAI-linked change; the complete capability set was not measured. Completed career-related action was treated as possible functioning evidence only when linked to a valued career aim, while intention, other reported behaviour and delay or inactivity were classified separately. Author's synthesis, informed by Sen (1999), Robeyns (2005, 2017), and Martins (2006).*

Vignette responses from nine participants were analysed separately as evidence of hypothetical reasoning. Each response was coded as a proposed relationship between a contextual condition, the aspect of guidance use it affected, and the anticipated consequence. Personal episodes and vignette material were coded independently before being compared. A proposed process also evident in lived accounts could refine an explanation, whereas patterns appearing only in the vignette remained hypothetical. Because Yilin and Mei differed across

several conditions, comparisons focused on configurations of support and constraint rather than an independent effect (Hughes & Huby, 2004; Jenkins et al., 2010).

Case memos and a case-by-code matrix supported comparison across the ten interviews and the development of candidate themes. The appendices provide the coding framework and case-by-theme matrix. Within-case interpretation preceded cross-case comparison to preserve extracts within their complete-case context. A theme required a coherent pattern relevant to at least one research question. Cross-case comparison examined similar resources under different conditions, while disconfirming cases were used to challenge emerging interpretations. Abduction reconsidered observed patterns through explanatory concepts, and retroduction examined the conditions that made a pattern or contrast possible (Fryer, 2022; Wiltshire & Ronkainen, 2021). Proposed explanations were checked against complete cases and plausible alternatives. Thus, the final claims concern contributing conditions and setting-specific processes.

The study used four procedures to strengthen the analysis. Contextualised extracts made the evidence underpinning each theme visible, while disconfirming cases and rival explanations helped narrow or revise claims. The audit trail documented changes to codes, theme boundaries, and rejected interpretations. Reflexive and translation memos documented how the researcher's position and language choices shaped the analysis; Standard Chinese extracts, presented alongside English translations, supported checks of wording and meaning. NVivo 15 was used for coding and retrieval; Excel was used for the comparison matrix; and Word was used for analytical memos. Interpretive decisions remained the researcher’s responsibility. Consistent with reflexive thematic analysis, quality was assessed through transparent analytical decisions, coherent themes, and claims proportionate to the evidence (Braun & Clarke, 2021).

Several checks addressed plausible alternative interpretations. Participants’ reported confidence was compared with any evidence of verification; polished wording was distinguished from biographical accuracy; clearer articulation of a career route was separated from employer recognition and legal eligibility; and intended action was distinguished from completed action. Complete-case review assessed whether an extract was consistent with the participant’s wider account. Disconfirming cases helped refine theme boundaries, while evidence-type labels prevented vignette predictions from being presented as lived outcomes. Translation decisions were checked against the source-language extract and accompanying memo. Together, these procedures supported interpretive credibility through transparent

reasoning rather than numerical agreement. As a single researcher conducted the analysis, consistency was supported through documented decision rules, repeated review, and an audit trail. Independent replication, longitudinal follow-up, and evidence from employers or advisers could test these interpretations using different forms of evidence.

The study used participant clarification during interviews rather than post hoc validation of cross-case themes. In reflexive thematic analysis, themes are contextualised interpretations developed across cases (Braun & Clarke, 2021). Returning decontextualised themes to participants could imply that an individual participant could validate an interpretation involving other participants, while also creating additional contact and anonymity concerns. Participants' meanings remained central through reconstruction of concrete episodes, contextualised quotation, attention to disconfirming cases, translation checks, complete-case review, and a transparent audit trail. During interviews, participants were able to clarify, correct, and add relevant information, while the researcher retained responsibility for the final analytical account.

Reflexivity

Reflexivity was an integral part of the method, as the interviews and resulting themes were shaped by interactions between participants and the researcher. I occupied a partial-insider position as a Chinese-speaking postgraduate researcher familiar with UK higher education. Shared linguistic and educational references facilitated rapport and enabled participants to use Mandarin, English, or code-switching. However, these shared references may also have encouraged me to regard familiar expressions or constraints as self-evident, while my institutional position may have encouraged participants to present themselves as responsible GenAI users (Dwyer & Buckle, 2009).

Several practices addressed the risks of assumed shared understanding and socially desirable presentation while acknowledging the researcher's situatedness. Clarification probes encouraged participants to explain familiar terms and reconstruct specific events. After each interview, a reflexive memo recorded moments of surprise, identification, potentially leading questions, and alternative interpretations. During analysis, these memos were reviewed alongside disconfirming cases and cross-case comparisons to assess whether an interpretation extended beyond the case in which it originated (Berger, 2015; Pillow, 2003). Reflexivity therefore made researcher subjectivity explicit and open to critical evaluation.

Ethical Considerations

Ethical approval was obtained before recruitment through the University of Oxford's departmental review process, which operates under CUREC governance. The study also followed British Educational Research Association guidance on voluntary participation, informed consent, confidentiality, and minimising harm (BERA, 2024; University of Oxford, 2026). Participants received an information sheet and written consent form explaining recording, withdrawal, and the use of de-identified quotations; both documents are reproduced in the methodological appendices. Consent was treated as an ongoing process, and participants could pause the interview, decline to answer a question, or end their participation at any time. Direct contact from interested students reduced potential pressure associated with peer-network recruitment.

Career discussions could involve sensitive family, financial, or immigration concerns. The interviewer acted solely in the role of researcher and collected de-identified interview accounts rather than credentials or identifiable application documents. Participants were labelled P1–P10, while distinctive details and combinations of characteristics were removed or broadened to reduce the risk of deductive disclosure within this small sample (Kaiser, 2009). Contact and consent records were stored separately from de-identified research data on encrypted, password-protected, university-approved storage. Data retention, withdrawal, access, and deletion procedures followed the approved participant documentation.

Chapter Summary

This single-site study uses ten semi-structured interviews to reconstruct situated episodes of GenAI use and compare variation in career needs, conversion factors, perceived influence on career capabilities, and career-related actions. Chapter 4 reports general personal assessments, lived episodes, and vignette reasoning as distinct types of evidence. Reflexive thematic analysis identifies cross-case patterns; Capability Approach decision rules distinguish resources, representational change, reported action, functioning evidence, and perceived influence on substantive opportunity; and critical-realist reasoning treats interview accounts as bounded forms of evidence. Together, these methodological choices provide the evidential structure for the analysis presented in Chapter 4.

Chapter 4: Findings and Discussion

Introduction and Analytical Model

Across the interviews, participants most often associated GenAI with starting a career task: they put an uncertain idea into words, produced a rough draft, broke work into steps, or began a search. Participants' accounts linked perceived changes in opportunity to prior knowledge, accessible people and sources, and the recruitment, immigration, and labour-market rules they faced. The interviews provide strong evidence that participants perceived preparation to be easier, but much weaker evidence that GenAI expanded the set of careers they could realistically choose and pursue.

Findings and discussion are integrated. Within each theme, interview evidence is analysed through the Capability Approach and related to Chapter 2 and the relevant research question. Generated content is treated as a resource. Career capabilities are valued options that participants could realistically pursue; clearer understanding alone is classified as a representational change. A career-related action becomes a possible functioning when the participant links it to a valued career aim. Participant-perceived influence on capability requires an explicit association between GenAI and a change in the realistic freedom or ability to pursue an option (Egdell & Robertson, 2021; Robertson & Picard, 2021; Robeyns, 2005, 2017; Sen, 1999).

Table 4.1. Analytical Themes, Mechanisms, Evidential Boundaries, and Research-Question Alignment

Theme	Empirical pattern	Analytical mechanism	Evidential boundary	RQ
Low-friction activation	Privacy, immediacy, task decomposition, and blank-page support	Cognitive externalisation, reduced interpersonal exposure, and temporal availability	Evidence of guidance quality or opportunity change remains limited	RQ1
Co-produced personalisation	Specificity increased with detailed biography, aims, and constraints	Audience-oriented co-production across user, model, verifier, and anticipated audience	The novice inequality pattern is participant-theorised and vignette-supported	RQ2
Distributed verification	Sources were allocated different credibility roles by task and claim	Distributed epistemic labour, authority and accountability	Unequal resources and burdens were observed; employment effects remain open	RQ2
Bounded preparatory actions under constraint	Searching, revising, rehearsing, and testing options	Preparatory action and clearer representation under conversion conditions	Evidence of capability-level influence remained unsubstantiated; external feasibility remained decisive	RQ3

The four themes describe connected analytical tasks rather than a causal or temporal sequence followed in every case: starting, tailoring, checking, and choosing action, delay, or inactivity. Participants could return to an earlier task, repeat a task, or stop; checking could prompt a new query, while an application could expose missing preparation and change the assessment of an option. Participants associated GenAI with easier initiation more often than with easier verification or improved feasibility.

The evidence labels also delimit the claims. A personal episode records reported experience or action; a general personal assessment captures an interpretation of recurring practice; and vignette reasoning records responses to hypothetical conditions. Quotations were spoken in Mandarin, unless they contained an English term, and were translated by the author. Fillers and immediate repetitions were removed sparingly; ellipses mark substantive omissions, and square brackets mark supplied wording or clarification. Light grammatical editing preserved meaning and tone. Idioms were translated for meaning, with analytically useful images retained and explained.

Low-Friction Activation: Creating an Accessible First Step in Career Work

Exploring Uncertain Ideas With Low Social Exposure

Participants began career work through three distinct routes. Cognitive externalisation turned a diffuse task into words, a draft, or a sequence of steps. Participants described reduced interpersonal exposure as supporting tentative articulation before social scrutiny. Temporal availability removed the need to wait or coordinate when work stalled. Each route helped a participant take a first preparatory action, but each addressed a different obstacle. P3 illustrated reduced interpersonal exposure by describing possibilities that felt too speculative to disclose:

> "Some ideas were rather far-fetched, and I felt embarrassed to tell other people—for example, that I wanted to start a business or open a shop. Asking someone for an opinion takes courage because you are exposing your thoughts. AI is like a friend who keeps your secrets." (P3, general personal assessment)

The relevant privacy was relational. P3 could test a possible future in private before discussing it with another person. By contrast, P10, who was already oriented towards financial-market roles and could consult industry contacts, used GenAI mainly for rapid information gathering and drafting. P1 similarly valued speed at an early exploratory stage.

The comparison across P3, P10, and P1 shows that apparently similar AI-first behaviour arose from different circumstances: low social exposure for one participant and the convenience of repetition or search compression for the others.

Emotional support varied across participants. P2 drew a clear boundary between practical assistance and reciprocal understanding:

> "Compared with AI, I would definitely choose to talk to a person... I use it for very practical things, such as editing my CV." (P2, general personal assessment)

For P2, low-friction access was useful when the task had a visible object that the system could edit. Emotional uncertainty called for reciprocal human understanding. P3's image of a 'friend who keeps your secrets' served a different function: it expressed relational privacy, allowing an unfinished idea to be articulated before exposure to social judgement.

P10 provided a disconfirming case by rejecting simulated warmth while continuing to use GenAI:

> "It tries very hard to make itself sound like a real person, but the artificial feeling is strong... I get [emotional support] more from real people." (P10, general personal assessment)

P10 nevertheless used GenAI for rapid industry scanning and drafting. Across the cases, reassurance was optional and unevenly credible. In P5's account, reassuring language helped the participant move towards a revisable draft. For P8, reassurance accompanied task decomposition: the model organised several demands into smaller steps, producing material that the participant could appraise and act upon. P2, P9, and P10 evaluated warmth as emotional support and found generic, repetitive, or artificial responses to be of limited value. Reassurance aided activation when it bridged anxiety and action; participants seeking recognition, challenge, or emotional reciprocity preferred people.

> "When several things pile up, I send it everything I need to do. It organises them and reassures me, which eases my anxiety a little... If I tell it that I am anxious, it breaks things into smaller pieces so that I can deal with them step by step." (P8, general personal assessment)

Moving Beyond the Blank Page

For other participants, the threshold was the blank page rather than disclosure. P5 described wanting to apply for a demanding role but being unable to begin:

> "I usually tell AI that I must apply for this role today... It first reassures me... I use AI as a way to initiate the task. It is much easier if it produces something rough first and I gradually revise it." (P5, personal episode)

P5 used the generated draft as a revisable object: the intimidating blank page became material that could be rejected, reorganised, and rewritten. The reported change was a movement from avoidance to active revision. Submission status remained unverified, so the evidence supports task initiation rather than an application outcome.

P6, who was moving from linguistics towards commercial work, described a related use when procrastination made the task feel too large:

> "Sometimes I procrastinate and feel that the task is too difficult to continue. At those times I [ask] AI how to break it into the smallest possible steps so that it is easier to carry out." (P6, general personal assessment)

P8 likewise converted diffuse anxiety into a shorter plan. Across the three accounts, P5 used GenAI to produce a rough draft, P6 used it to sequence a large task, and P8 used it to categorise a problem. A blank template or checklist requires the user to define the problem in advance. GenAI instead accepted fragments in the participant's own words, proposed a provisional structure, and responded to follow-up questions. Participants could then revise an imperfect object while retaining responsibility for the evidence, wording, and execution (Baek et al., 2024; Miao & Holmes, 2023).

Activation as a Bounded Contribution

Participants particularly valued immediacy when formal career provision poorly matched their needs. P2, who was preparing to return to China, described the problem in destination-specific terms:

> "The vacancy emails I received were mostly unsuitable for my subject and more academic... Most of the roles were also based in the UK. For someone planning to return to China, I felt that the university could offer very little." (P2, general personal assessment)

The service was available, but its dominant geographical and occupational orientation poorly matched P2's plan to return to China. GenAI provided an immediate space in which to organise questions and application language, while Xiaohongshu, employer pages, and employed friends supplied current knowledge about Chinese recruitment that lay beyond the generated response.

P9's mismatch arose from access rather than destination:

> P9 reported losing portal access after graduation and shifting more of their job-search activity to AI because it produced the quickest result.

Read together, the cases show different relationships between relevance and entitlement. P2 retained access to a service that was poorly suited to a return-to-China search; P9 lost access to a previously available portal. GenAI became salient in both cases, but for different institutional reasons.

Temporal availability reduced waiting and coordination. A scheduled appointment operated on a different timetable from the moment a cover letter stalled; a vacancy email still required sorting; and repeated minor questions were easier to pursue through an interface that could be reopened immediately. P2 retained institutional access but found its UK-facing content poorly suited to a return-to-China search. P9 lost portal access after graduation, while P10 valued rapid scanning despite having access to industry contacts. Immediacy became salient through destination mismatch, loss of entitlement, or a preference for repeatable searching, depending on the case.

Friends and advisers offered situated judgement, emotional reciprocity, and accountability, although accessing them required disclosure and coordination. Templates and checklists offered fixed rather than conversational support. GenAI combined immediate revision with interaction, yet supplied limited situated accountability. Participants used it alongside people, careers services, and recruitment platforms, assigning task initiation, judgement, and verification to different sources. Digital career-guidance research similarly links repeated self-service with concerns about quality, equity, and professional support (OECD, 2024a; Pandya & Wang, 2024; Tanner & Samson, 2026).

Summary: Different Frictions and the Same First Move

The findings support research linking general-purpose GenAI to rapid assistance (Chan & Hu, 2023; Holland & Ciachir, 2025) and career-specific use to drafting and task support

(Flink et al., 2025; Loidl et al., 2026). They distinguish among cognitive, social, and temporal obstacles and qualify evidence of an AI working alliance because P10 continued using GenAI after rejecting its simulated warmth (Barger, 2025). The cases answer RQ1 by showing how different circumstances produced the shared outcome of beginning a preparatory task. Theme 2 examines how participants made that initial output personally relevant.

Co-Produced Personalisation: Tailoring Depended on Participant Knowledge

The interviews revealed four graded forms of personalisation, each requiring different forms of knowledge and a different relationship to the intended audience. Linguistic personalisation reformulated user-supplied material into an anticipated professional register. Biographical personalisation selected and organised experience that the participant could truthfully own. Occupational personalisation mapped that experience onto the criteria of a role or application genre. Destination-specific personalisation aligned the narrative with current recruitment timing, institutional rules, and tacit expectations in a particular labour market. External sources became increasingly important as tailoring moved from wording towards occupational and destination fit.

The four-layer relationship between prior knowledge and tailoring supports a conditional personalisation paradox, best treated as a participant-theorised mechanism. In discussions of the vignette, participants reasoned that a novice could gain broad orientation from a general prompt yet lack the biography, occupational vocabulary, and evaluative criteria required for fine-grained tailoring. Personal episodes partly illustrated the mechanism: P5, P6, and P10 achieved stronger tailoring after supplying structure and evidence, while P2 and P4 repaired fabricated claims using knowledge they already possessed. These episodes establish the resource demands involved in production and repair; the unequal novice pattern itself remains a vignette-supported proposition (Wang et al., 2025).

Biographical and Linguistic Input

The strongest examples occurred in application writing. P5 maintained an ‘experience bank’, wrote the substance before using GenAI, and compared outputs from different models. The participant described the upstream authorship requirement directly:

> “I make sure that everything I want to say is already there, even if it is poorly organised. The more detail I write myself, the better. It can summarise what I write, but I do not let it extend the experience.” (P5, personal episode)

P5's account reveals three forms of labour. The participant supplied biographical material, the model compressed it into professionally legible language, and, analytically, the anticipated interviewer functioned as an accountability standard because retained claims had to remain explainable under questioning. 'Accountability standard' is therefore a researcher-derived formulation rather than P5's own term:

> "I would not include something I had never done, because I would not be able to answer questions about it in an interview. If the model said that I had done something when I had only participated in or observed it, I would still include it if I felt able to explain it and answer questions." (P5, personal episode)

Personalisation was therefore oriented towards an anticipated audience before that audience encountered the application.

P6 described a closely related process:

> "The most useful part is that I already have some ideas about my internship experience and write them down first in relatively immature language. It then identifies quite precisely how these points connect with the role I want and presents them in a [more professional] way." (P6, personal episode)

Here, personalisation exceeded Mandarin-to-English translation. GenAI connected biographical evidence with an occupation and the genre conventions of a CV. P10 made the upstream authorship equally explicit:

> "I first have a broad structure and feed that structure to it... It generates a version, and then I adjust what it has produced." (P10, personal episode)

Across P6, P9, and P10, participants supplied events, aims, and exclusions; GenAI proposed links between biography and role; and participants tested their truthfulness and defensibility. Employer or assessor expectations shaped what participants asked the model to foreground. Employer recognition lay beyond the interview evidence, so the accounts support a claim of audience-oriented drafting.

The allocation of biography, model wording, and anticipated audience judgement extends accounts of international-student employability that emphasise making experience recognisable within context rather than merely accumulating skills (Pham et al., 2024; Soltani & Tomlinson, 2024). Linguistic and biographical tailoring were comparatively strong because participants could anchor them in evidence they possessed. Occupational and

destination relevance remained subject to application conventions, current labour-market information, and employer interpretation (Gedrimiene et al., 2024).

When the Model Fabricated the Applicant

The same fluency could turn personalisation into misrepresentation. P2, who was preparing applications for employment in China, used GenAI to make their experience sound more action-oriented but found that it inserted an invented achievement:

> "Sometimes it generates figures... I correct those details to match what really happened." (P2, personal episode)

P2 could repair the invented metric because it contradicted a remembered achievement. The same repair rule appeared in different forms elsewhere: P4 removed a skill inferred from an experience because it fell outside the participant's actual learning; P7 repeatedly reduced exaggerated wording; and P8 removed a recognisably synthetic tone. These variations do more than record four corrections. They identify the shared evidential test that the model lacked: whether the generated applicant remained anchored in experience that the participant could truthfully own and defend.

Hallucination refers here to generated content that conflicts with available source material or cannot be supported by relevant factual knowledge (Ji et al., 2023). On this definition, P2's invented metric and P3's fabricated vacancies were hallucinations. Stale recruitment information, generic wording, and contestable recommendations instead raised distinct problems of currency, fit, and judgement. P4 discussed hallucination as a wider limitation of GenAI:

"It can help with analysis, but it still has the problem of hallucination. The industry treats this as an important research issue, and at present there is no better solution." (P4, general personal assessment)

Unlike P2 and P3, P4 offered a general assessment rather than a separate observed incident. The account helps explain why generated career analysis remained provisional and required independent judgement.

Participant knowledge made repair possible: P2 knew the actual metric, P4 knew which skill was absent, and P5 knew what could be defended. Self-knowledge served as both a production input and an authenticity control. This dual role arose from an asymmetry between participant knowledge and model fluency. Consistent with Bender et al.'s (2021) account of ungrounded

language generation, GenAI could transform supplied material into plausible professional claims without independently establishing whether those claims corresponded to lived experience. Participants therefore had to reconnect each polished claim to biographical evidence and decide whether they could defend it to an employer. Biographical knowledge thus operated as a personal conversion factor: the model supplied fluency, while the applicant retained responsibility for truthfulness and defensibility.

Orientation, Mirroring, and the Limits of Tailoring

Limited prior knowledge could nevertheless support orientation. P3's cross-border tax query shows what that modest contribution looked like in a personal episode:

> "Once it told me that a particular tax might exist, I would search for it myself. It gave me a point of entry." (P3, personal episode)

The answer created searchable categories while leaving legal resolution to further inquiry. P3 turned an unknown unknown into a searchable question and preserved the distinction between orientation and authority. P6 described a related benefit when consulting and management trainee roles were difficult to compare with a linguistics background. P7 explained the graded relationship through the vignette:

> "When you are completely unclear, a broad question may only help you find a direction. When you already understand the information more clearly, you can use it to refine a particular experience or what you want." (P7, vignette reasoning)

P7's vignette articulates the hypothesised novice pathway: broad output can create searchable categories, whereas role-specific tailoring requires criteria that the user already possesses or can obtain. P3's tax episode partly illustrates the orientation stage, while the application episodes of P5, P6, and P10 illustrate the more resource-intensive tailoring stage. Occupational tailoring was stronger when a job description, transferable evidence, and an ability to reject assumptions were available; destination-specific tailoring additionally required current external sources.

Tailoring could also become preference mirroring. P8 described how the interaction itself could narrow disagreement:

> P8 observed that repeated preference-laden prompts steered the model towards existing views and reduced the likelihood of challenge (general personal assessment).

Greater personal relevance sometimes coincided with less independent challenge. P5 similarly found that the model elaborated within an existing scenario, while P6 described advice as responsive to the position taken in the exchange. A personally agreeable response could offer limited guidance and required a separate assessment of realism, value, and selection requirements.

The layered nature of personalisation is particularly important for transitions between the UK and China. P2 could use GenAI to improve application language but still needed local knowledge of Chinese recruitment calendars and role-specific expectations. Zhao and Cox (2022) similarly identify mismatches between UK study, Chinese recruitment cycles, and the forms of experience recognised by different stakeholders. Participants used GenAI to make narratives more portable at the linguistic level; establishing destination fit required local knowledge external to GenAI.

Specificity alone was a weak proxy for personalisation quality. The same precise sentence could be anchored in biography, inferred from a job description, or produced through fluent completion. Quality depended on a traceable chain: truthful user evidence, modelled relations, participant verification, and audience-relevant criteria. The cases illustrate three breaks in the chain: genericity, preference mirroring, and fabrication.

Summary: Personalisation as Audience-Oriented Conversion

The findings support literature that treats employability as contextual recognition (Pham et al., 2024; Soltani & Tomlinson, 2024) and extend personalisation research by identifying four layers and an anticipated audience. For RQ2, relevance emerged from the allocation of work: users supplied biography, GenAI proposed relations, users authenticated them, and employers or institutions controlled recognition. The participant-theorised paradox follows because fine-grained help demands more input and evaluative knowledge. Theme 3 asks how this tailored content acquired warranted credibility.

Distributed Verification: How Generated Content Acquired Situated Credibility

Participants evaluated trust claim by claim, according to the question, its stakes, and the available evidence. Distributed verification names the resulting claim-specific allocation of checking and authority. The analysis asks who performed each check, which source had authority, and who remained accountable for use.

Three judgements structured participants' checking. Factual verification asked whether a claim was true, accurate, and current: whether a vacancy existed, a visa rule applied, or a numerical achievement was real. Contextual appraisal asked whether a claim fitted a particular participant, employer, or destination: whether a recruitment timeline applied in China or a broad degree category concealed specific selection requirements. Evaluative judgement asked whether an option served a valued aim and whether generated language remained authentic and defensible. The three judgements overlap in practice, but naming them prevents 'verification' from becoming a catch-all term.

Table 4.2. Source Roles in Distributed Verification

Source	Main role	Principal limitation
GenAI	Synthesis, wording, rehearsal, and task decomposition	Opaque provenance; stale or fabricated content
Official employer or government pages	Vacancies, formal rules, and dates	Limited tacit knowledge or individual fit
Recruitment and social platforms	Timeliness, emerging opportunities, and lived accounts	Fragmentation and commercial incentives
Advisers, alumni, and practitioners	Tacit practice, challenge, and recognition	Unequal access and variable expertise
Participant knowledge	Biography, authenticity, and values	Uneven occupational and destination knowledge

Table 4.2 maps each source to its principal role and limitation in verification. Generative synthesis was readily available, while provenance recovery and the selection of authoritative sources remained the user's responsibility (Sundin, 2025).

Factual Verification

P3 expressed the claim-type contrast most clearly:

> P3 treated concrete vacancy claims as requiring external search, while using GenAI more readily for initial comparison and subjective analysis (general personal assessment).

The participant continued to use GenAI for interview questions, comparisons, and initial framing, but raised the evidential threshold for claims that should leave an external trace. P8 described the next move directly:

> "A good way is to open the official website and look at it... I mainly conduct a second search myself." (P8, general personal assessment)

P8's confidence followed the second search rather than the fluency of the initial answer. The official page made the generated claim inspectable against an accountable source; broader career fit required separate contextual and evaluative appraisal.

Fabricated, stale or unsourced content increased participants' checking workload. P2's earlier academic use shaped present caution:

> "It produced fake references and sometimes brought in information from several years ago... Going back and forth like that feels like a waste of time." (P2, general personal assessment)

Fake references represented fabricated content, while older information could be accurate but stale. Both placed the work of establishing truth or currency back on the participant. P1 still experienced comparison and correction as faster overall, showing that the net value depended on claim type, familiarity, and the availability of an authoritative check.

Recent research helps explain why fluent text was a poor credibility cue. Farquhar et al. (2024) treat confabulation as a specific uncertainty problem within a wider range of low-quality output. Kalai et al. (2026) show how common training and evaluation incentives can reward guessing rather than abstention. In behavioural experiments, users overestimated the accuracy of default LLM explanations, and longer explanations increased confidence while correctness remained unchanged (Steyvers et al., 2025). Although these studies address domains beyond career guidance, they clarify why confident language could increase the need for verification.

Contextual Appraisal

A factually plausible statement could still be unusable in context. P2 contrasted GenAI's synthesis with the timeliness and specificity of Chinese platform posts:

> "The greatest difference between AI and human sources is timeliness. A person's post can refer to a specific role and is backed by someone who has actually done it, whereas AI may provide last year's information or something very general... Human information is also fragmented... AI integrates information better, but human sources are more useful for specific details." (P2, general personal assessment)

P2's account shows how sources complemented rather than replaced one another. Xiaohongshu and employed friends supplied current recruitment timing, while GenAI organised and phrased material. P7 and P8 similarly used recruitment platforms to discover

roles and employer pages to confirm formal details. Platform accounts could be partial, commercial, or outdated; official pages could confirm that a vacancy existed; and practitioners could explain what a team valued in practice. Each source answered a narrower question.

P4's later application for a role in a suggested sector demonstrates why contextual appraisal exceeded fact-checking. The occupation existed and was broadly adjacent to a mathematics-and-computing background, but the suggestion assumed a level of preparation in statistics and probability that differed from the participant's curriculum. Contextual appraisal joined two questions: 'Is this a real career?' and 'Does this route fit the participant's curriculum and the selection practices actually encountered?' AI literacy research stresses knowing when a system's scope has been exceeded and where to seek further evidence (Ng et al., 2021). The interview evidence extends technical AI-literacy models to situated career judgement: participants had to recognise which context the model had omitted.

Authenticity and Evaluative Judgement

Evaluative judgement concerned both accuracy and judgements about what should be pursued and represented. P5's interview-defensibility rule is the clearest developed case: generated wording remained only where the underlying experience could still be owned under questioning. P2 applied the same principle by retaining a professional sentence structure while deleting an invented metric. P6's comparison of model outputs shows a further variation:

> "I used ChatGPT and [DeepSeek] together to revise my CV in Chinese... I chose the version with plainer language because the other one went too far; I would have felt uncomfortable using it. The plainer version still got directly to the point." (P6, personal episode)

P1 provided a disconfirming case because a plausible recommendation failed to direct action:

> P1 described applying even when GenAI judged fit as weak, because a broader application strategy kept the option open (general personal assessment).

P1 treated the model's judgement as one input and kept the option open. P10 similarly used GenAI to analyse alternatives while retaining final decision-making authority, and P7 reserved consequential choices for personal and human judgement. The disconfirming case

shows selective reliance in action: assistance with wording, rehearsal, or comparison coexisted with retained authority over which valued route to pursue.

Experimental work shows that people can follow AI advice even when contextual information weighs against it and that verification costs and cognitive effort affect overreliance (Klingbeil et al., 2024; Vasconcelos et al., 2023). The interviews provide a complementary qualitative account of resistance: participants used biography, anticipated accountability, and valued aims as grounds for rejecting fluent content. The contribution is to treat authenticity as an evaluative achievement rather than an automatic effect of personalisation.

Platform–User Responsibility Asymmetry and Verification Inequality

Participants described an asymmetry between what the platform influenced and what users had to verify. They reported lower production effort than verification effort; the interface presented options and applicant narratives, while students retained accountability for using and defending them. Research on educational platforms frames this redistribution of agency and labour (Perrotta et al., 2021). P10 could mobilise industry contacts, whereas P4 recognised the need for insider knowledge but struggled to activate the relationship; P6 faced similar uncertainty when seeking alumni in unfamiliar commercial fields. Institutional access also varied: P9 lost access to a careers portal after graduation, while P2 could use destination-specific recruitment platforms. Verification inequality therefore refers to unequal access to the resources and routes required for verification, together with unequal verification burdens in establishing truth, contextual fit, authenticity, and feasibility. Interface access and access to verification resources remained distinct.

Participants' vignette reasoning supplied a further interpretive layer. P5 contrasted an experienced fictional student, who could treat generic advice as an execution aid, with a novice:

> P5's vignette contrasted an experienced student using AI for efficient execution with a novice who could start quickly yet faced a greater burden in checking accuracy (vignette reasoning).

P7 similarly reasoned that richer knowledge enabled more granular questions, while P10 thought that an urgent, less-prepared student should consult a person. These vignette responses support a participant-theorised account of greater evaluative difficulty for a novice, while personal episodes document actual differences in portal access, contacts, and

occupational knowledge. Variation in student AI literacy and evidence that nominal access can coexist with unequal benefit provide relevant context (Hornberger et al., 2023; OECD, 2024a, 2024b). Together, these accounts support variation in verification resources and burdens at the level of reported evaluation. The term ‘verification inequality’ is used here for the unequal distribution of those resources and burdens, while the direction and magnitude of any employment-outcome inequality remain open.

Summary: Credibility, Responsibility, and Inequality

The evidence supports recommendations for hybrid guidance (Miao & Holmes, 2023; Sarmurzin et al., 2026) and provides a more specific account of how responsibility was allocated. GenAI synthesised and worded material; students recovered sources and tested contextual fit; official pages and practitioners supplied different forms of authority; and students remained accountable for action. RQ2 shows that unequal knowledge, time, institutional access, and contacts made this checking work easier for some participants than for others. Theme 4 examines whether a checked resource changed action or opened a feasible option.

Perceived Influence on Career-Related Actions Under Persistent Constraint

The clearest perceived influences concerned task-level actions: searching, revising, rehearsing, applying, and testing. An action became a possible functioning when the participant linked it to a valued career aim; other activity remained reported behaviour, while plans and generated text remained intentions or resources. Capability-influence claims additionally required a valued and realistically feasible option and an explicit participant association between GenAI and a change in the realistic ability to pursue it (Egdell & Robertson, 2021; Robertson & Picard, 2021).

Participants most consistently perceived GenAI to influence bounded preparation and the clarity of possible routes. Evidence of substantive career-capability influence remained limited. The recurring contribution was the immediate, revisable organisation of fragments from participants and other sources, with structure, context, and accountability distributed across GenAI, job descriptions, web searches, and people.

Preparatory Action

P1 provided a clear resource-to-action chain:

> "I ask it, 'If I want this role, which skills am I currently missing?' Based on its answer, I conduct a second round of searches and then work through the skills and experience I need to build." (P1, personal episode)

The GenAI response was a resource. P1's further search and preparation, P4's application and selection test, P2's correction of a fabricated metric, and P3's follow-up tax search were reported preparatory actions. They were treated as possible bounded functionings only where the participant linked them to a valued career aim. Evidence of action was strongest where the task had a bounded object that the participant could inspect.

Reported next steps varied in evidential strength. P5 used a draft to overcome the blank page, while submission status remained unclear; P8 described high-volume application automation as a future intention. P7's contingency plan remained unused after employment circumstances improved. This case illustrates informed inactivity following a change in circumstances but leaves option protection or expansion unresolved.

> "It might make small changes to my CV and then submit applications directly for me; apparently it can do that automatically. I have seen online examples of other people using it in this way... For mass applications, I think it is something I could try." (P8, future intention)

Narratable Possibilities and Exercisable Options

P6's account shows a representational change. Moving from linguistics towards commercial work, the participant recalled a reframing of consulting:

> "The methodology remains the same while the content changes. Consulting is more like the research work I did before, except that the object changes from linguistic data to information about industries." (P6, personal episode)

The reframing made an unfamiliar transition easier to explain. P6 kept consulting and management trainee roles under consideration, while employer recognition and selection eligibility remained unobserved. The evidence supports clearer route articulation, while capability expansion remained unsubstantiated (Nguyen & Cheng, 2025; Pham et al., 2024).

P4's attempted transition exposed the difference:

> P4 reported that the application exposed the preparation in statistics and probability required by the suggested sector and changed the assessment from plausible adjacency to poor selection fit (personal episode).

The recommendation connected P4's broad degree profile with a superficially adjacent occupation, while the application exposed assumptions about preparation in statistics and probability. P6's account showed a different result: GenAI helped articulate continuity across fields, while practical feasibility remained untested. A narratable route may remain constrained by curriculum and selection conditions.

The contrast between P6 and P4 shows why the Capability Approach matters. Generating or redescribing an option can support imagination and preparation. Capability-level influence requires a valued route that remains realistically available and an explicit participant association between GenAI and a change in freedom to pursue it. Participants' accounts linked career guidance to preparation, while employer recognition and structural filters set boundaries on whether the route could be pursued (Fuertes et al., 2021; Robertson & Picard, 2021).

Participants associated problems in generated content with weaker preparation. Fabricated metrics and vacancies required extra checking; stale information raised timing concerns; preference mirroring reduced challenge; and P4's application exposed a weakly grounded inference from degree title to occupational preparation. Downstream effects on wider exploration remained unknown. Legal, institutional, and labour-market constraints operated alongside those content-level problems.

Structural Filters and Selective Action

Visa sponsorship provided the clearest direct structural boundary. P9, a recent graduate seeking work in the UK, stated:

> P9 limited applications to sponsoring employers because immigration rules left minimal room for negotiation (general personal assessment).

A vacancy for which the participant was ineligible remained infeasible regardless of improvements in wording or search speed. Other constraints operated differently. P2's destination decision involved living costs, cultural fit, and proximity to family. P5 faced immigration tightening, sector contraction, and uncertain recruitment cycles. P7's choice set

depended on current vacancy supply, while P4 and P8 found relevant insider knowledge or networking difficult to obtain. Participants treated legal, economic, temporal, and relational constraints as constitutive filters for worthwhile action.

Structural conditions entered at different points in the sequence. Sponsorship could remove an option before application work began. A selection exercise could reveal a curriculum gap only after a recommendation was tested. Family proximity or living costs could leave an occupation technically available but make it incompatible with the life a participant valued. Conversion depended on four concrete conditions: legal eligibility, practical preparation, access to knowledgeable people, and consistency with the participant's valued life. GenAI could help a student articulate each problem. Changing those constraints required decisions or action beyond the GenAI response.

The constraints also gave inaction several meanings. P7's contingency plan remained unused after circumstances improved, representing an informed decision not to act. P4's application revised the assessed feasibility of a route after selection demands became visible, while P10 retained consequential choice after using GenAI for information. A capability analysis separates informed refusal, valued choice, strategic delay, and structurally constrained progression instead of reading every absence of action as failure.

Inequality entered at two connected levels. The platform–user responsibility asymmetry concentrated verification and contextualisation on students, while differences among students shaped their capacity to complete those checks. P9's lost portal, P4's difficulty approaching insiders, and P6's limited ability to activate an alumni route constrained checking; P2's local platforms and P10's industry contacts supported it. Beyond verification, visa rules, recruitment timing, and employer recognition structured which warranted options became feasible (Nguyen & Cheng, 2025; Zhao & Cox, 2022).

Evidence from other task domains shows why the distributional direction depends on context. In Kim and Moon's (2024) controlled laboratory experiment, ChatGPT significantly reduced mean performance on mathematical problem-solving tasks; the authors attributed this result to participants with weaker mathematical problem-solving skills being less able to distinguish plausible but incorrect ('hallucinated') responses from correct ones. By contrast, novice customer-support workers achieved larger gains within a bounded system that provided organisational knowledge and feedback (Brynjolfsson et al., 2025). The contrast directs

attention to task boundedness, domain knowledge, and verification support when interpreting possible distributional effects.

P9 summarised the gap between access and capability:

> "It helps, but I cannot say that it makes me much more certain or suddenly more confident. Everyone has this tool, and if everyone knows how to use it, using it does not make my own job search much more certain. Even so, it helps to some extent." (P9, general personal assessment)

P9 separated instrumental value from relative advantage: broad availability made access a shared resource rather than a differentiator, while subjective judgement continued to shape certainty.

Summary: Preparatory Gains and Substantive Capability

The findings support contextual accounts of employability and the Capability Approach (Akkermans et al., 2024; Egdell & Robertson, 2021). Participants perceived clear effects on preparatory actions and some help in understanding and explaining possible routes. The interviews provided limited evidence that GenAI changed their real freedom to pursue a valued option. Fabrication, preference mirroring, and superficial occupational matching could reduce the reliability of preparation. External recognition, relationships, legal eligibility, and available opportunities remained decisive.

Cross-Theme Synthesis and Direct Answers to the Research Questions

Across the themes, participants associated GenAI with easier task initiation, but they still had to establish provenance, contextual fit, and practical feasibility. Personal knowledge, access to other people, institutional services, and labour-market rules shaped the burden of checking and conversion. Conditional conversion captures the resulting contrast: lower initiation effort coexisted with demanding evaluation and a largely unchanged set of feasible options. Appendix E displays supporting, boundary, and disconfirming evidence across all ten participants.

Appendix F (Table F.1) provides a cross-case synthesis linking developed and disconfirming cases to Chapter 2 and recent, claim-specific research, showing where the findings support, extend, or challenge prior accounts and how the underlying mechanisms inform each research question.

RQ1—What career guidance needs do students seek to address through generative AI, and what circumstances give rise to these needs?

Participants sought initial orientation, task decomposition, professional wording, application and interview rehearsal, live vacancy information, and low-exposure exploration of uncertain possibilities. These needs arose from blank-page friction, anxiety, uncertainty, destination mismatch, loss of portal access, urgency, and a desire for repeatable help. Participants used GenAI for first moves through cognitive externalisation, reduced interpersonal exposure, and temporal availability. The three activation pathways explain why use could coexist with competent human guidance and why artificial warmth was optional rather than a shared basis for activation.

RQ2—How do personal, social and institutional conversion factors shape students' interpretation, evaluation and use of GenAI-generated career content?

Personal conversion factors included biographical and occupational knowledge, AI literacy, values, and the capacity to recognise fabrication or preference mirroring. They shaped what students could disclose, prompt, authenticate, and defend. Social factors included the interpersonal exposure involved in asking for help, access to emotional reciprocity, and the ability to mobilise friends, alumni, or practitioners. Institutional factors included available time, deadlines, access to careers portals, destination-specific recruitment infrastructure, selection criteria, and immigration rules. Together, these factors shaped which layer of personalisation became possible and which generated claims were subjected to further checking.

Evaluation was distributed by claim type. Official pages supplied formal authority; platforms supplied timeliness; practitioners supplied tacit context; participant knowledge supplied biographical truth and values; and GenAI supplied synthesis and wording. Richer knowledge, more time, institutional access, and contacts lowered the cost of provenance recovery and contextualisation. Personal episodes document unequal portal access, local information, occupational knowledge, and mobilisable contacts. Vignette reasoning adds a participant-theorised expectation of greater prompting and checking difficulty for a less-prepared student. Therefore, the study identifies unequal verification resources and burdens, while employment outcomes and the distributional direction remain open (Bankins et al., 2024; OECD, 2024b).

RQ3—In what ways, if any, do students perceive GenAI to influence their career capabilities and the career-related actions they take?

Participants perceived GenAI to influence career work most clearly through bounded preparation: searching, correcting, drafting, rehearsing, applying, and testing suggestions. P6 also perceived a clearer account of a possible cross-field transition. P4's application exposed previously hidden selection requirements, while P9's sponsorship requirement constrained options independently of preparation. The accounts support a clear action-level influence and some representational influence; evidence of capability-level influence remained unsubstantiated.

Participants associated fabricated or stale content, preference mirroring, generic destination advice, and superficial occupational matching with less reliable preparation. They described visa sponsorship, recruitment timing, family proximity, vacancy supply, curriculum, and employer recognition as setting boundaries on whether a checked option remained feasible and valued. P7 perceived a stable option set after leaving a contingency plan unused, while P9 restricted applications to roles offering sponsorship despite GenAI-assisted preparation. Therefore, RQ3 is answered conditionally: participants perceived stronger influence on preparation than on substantive freedom, while capability-level influence remained unsubstantiated.

Chapter Conclusion

Participants used GenAI chiefly for orientation, task decomposition, professional wording, rehearsal, and low-exposure exploration. Personal knowledge, relationships, available time, and institutional access to services and authoritative sources shaped personalisation and verification. Participants perceived the clearest influence on preparatory actions and clearer representations of possible routes; evidence of capability-level influence remained unsubstantiated. The conclusion is therefore bounded: easier generation coexisted with variable accuracy, demanding verification, and limited change in substantive freedom. Retrospective accounts support short-term perceived influence. Long-term capability change requires longitudinal research that follows students into graduate employment, comparing those who continue using GenAI with those who discontinue it, across transitions into different national labour markets. Repeated interviews could trace how verification practices, institutional access, and feasible career options change after graduation.

Chapter 5: Conclusion

Study Overview

This dissertation examines how Chinese international students enrolled in or recently graduated from UK taught master's programmes used GenAI as an informal career-guidance resource. It asks what needs prompted this use; how personal, social and institutional conversion factors shaped students' interpretation and evaluation of generated career content; and how students perceived that GenAI influenced their career capabilities and career-related actions. Ten semi-structured interviews were used to reconstruct situated episodes of GenAI use, while a paired vignette completed by participants provided additional evidence of how students reasoned about contrasting conditions. The material was analysed through reflexive thematic analysis, capability approach decision rules and critical-realist reasoning.

The literature review brought together research on transnational study-to-work transitions, GenAI career guidance and its risks, and the conditions linking generated content to feasible career action. Sen's Capability Approach provided the evaluative lens, distinguishing resources, conversion factors, capabilities and functionings so that task support was not assumed to be expanded opportunity.

The finding was that GenAI reduced friction at the point of entry into career work. Participants used it to articulate uncertain ideas, break down difficult tasks into manageable steps, produce rough drafts, organise searches and rehearse recruitment encounters. These functions were particularly valuable under conditions of blank-page anxiety, time pressure, low interpersonal exposure, destination mismatch in formal careers provision or loss of institutional access. Therefore, GenAI helped participants begin or resume preparatory work. However, easier initiation alone did not resolve the career problem prompting the interaction.

The practical value of generated content emerged through a wider process of conversion. Participants contributed biographical evidence, occupational knowledge, preferences and career aims. Friends, alumni and practitioners supplied occupational judgement and tacit knowledge. Official websites and recruitment platforms provided formal authority, current vacancies and destination-specific information. Immigration rules, recruitment calendars, curriculum fit and employer recognition then determined whether a proposed route remained feasible. Personalisation was consequently co-produced through interactions among

participants, social contacts, institutional resources and the model. Credibility was established through verification across several sources.

This process also revealed an inequality extending beyond interface access. Participants differed in the time, occupational knowledge, institutional access and mobilisable relationships available to verify the claims generated. Some could consult practitioners or compare content with destination-specific recruitment platforms; others recognised the need for such information but struggled to access the relevant people or services.Thus, access to the same conversational interface led to unequal verification pathways.

Across the three research questions, the central difference was between assistance with a career task and the expansion of a substantive career opportunity. Participants associated GenAI most clearly with preparatory actions such as drafting, revising, rehearsing, applying and checking suggestions. In some cases, it also facilitated understanding and articulation of an unfamiliar route. These changes did not independently establish an expanded career capability. Under the capability approach, a career capability required a valued option that the participant could realistically choose and pursue. A route that became more intelligible and easier to articulate may remain unavailable due to curriculum mismatch, sponsorship requirements, recruitment timing, vacancy supply or employer recognition.

Thus, the overall conclusion is conditional. GenAI functioned as an immediate and revisable career resource whose perceived influence was strongest in task initiation, preparation and the representation of possible routes. Its conversion into credible guidance and a feasible career options depended on the participants' knowledge, access to appropriate people and sources, and the institutional and labour-market conditions surrounding each decision. The study supports clear participant-perceived influence on preparation and some representational change, while capability-level influence remained unsubstantiated within the interview evidence.

Contributions and Implications

The study contributes an evaluative framework for distinguishing different claims about the value of GenAI in career guidance. Generated content constitutes a resource. A clearer understanding or more coherent account of a career route constitutes representational change. An action connected to a valued career aim may provide evidence of bounded functioning. A

career capability requires stronger evidence that a valued option has become realistically pursuable. These distinctions prevent polished language, confidence, perceived usefulness or completed activity from being automatically treated as evidence of expanded career freedom. They also address a limitation of existing research, which frequently evaluates access, engagement, confidence or short-term task performance without tracing whether generated content contributes to a feasible option.

The study-specific sensitising concepts of conditional conversion and verification inequality further explain why similar GenAI resources acquired differing practical value. Conditional conversion directs attention to the combination of knowledge, relationships and institutional conditions through which generated content becomes usable, remains preparatory or fails to support action. Verification inequality identifies unequal access to the knowledge, time, relationships, institutional channels and authoritative sources required to establish provenance, contextual fit, authenticity and feasibility. Together, these concepts shift the analytical focus towards who can make GenAI outputs credible and actionable, which resources enable this process, and the conditions under which it occurs.

These findings also have implications for career-guidance practice. They support positioning GenAI as a preparatory resource within hybrid guidance, alongside professional and institutional support. A defensible workflow would distinguish generation from subsequent forms of judgement. Factual claims about vacancies, immigration eligibility, formal deadlines and published selection criteria should be checked against accountable official sources. Careers practitioners, alumni and other knowledgeable professionals can provide occupational context, challenge assumptions and explain tacit recruitment practices. Students remain the principal authority on their own biographical experience, values and on whether the generated application language is truthful and defensible.

The findings further suggest that hybrid provision requires organisational support. Advising students simply to ‘check with a human’ assumes that relevant expertise is identifiable, available and socially accessible. Destination-specific information, transparent routes to professional advice and continued access to careers resources after graduation could reduce some of the verification burdens identified in the interviews. Human expertise is likely to be particularly valuable where the task require contextual judgement, emotional reciprocity, high-stakes decisions or the interpretation of institutional rules that the generated content

cannot resolve authoritatively. GenAI can then retain its comparative strengths in immediacy, iterative drafting, task decomposition and rehearsal.

Limitations and Future Research

The conclusions are bounded by the study design. The sample was small and self-selecting, and all the participants had already used GenAI for career-related purposes. The study therefore explains variation among participating users but cannot estimate the prevalence of these experiences or represent students who avoid, discontinue or have restricted access to GenAI. The institutional setting and specific transition examined also limit transferability to universities with different careers provisions, programme structures or student populations.

The interviews were retrospective and cross-sectional. They provide evidence of participants' experiences, interpretations and reported actions; however, they do not directly capture prompts, complete model outputs, application quality, employer responses or longer-term employment outcomes. Participants may also attribute change to the most visible resource while overlooking other influences. The paired vignette supported structured comparison, although it combined several conditions and elicited hypothetical reasoning rather than actual behaviour. It could reveal how participants interpreted contrasting configurations of support and constraint; however, it could not isolate the effect of one factor or demonstrate how a less-prepared student would act. The rapid development of GenAI systems and changes in recruitment and immigration contexts also mean that the specific practices described here may evolve.

Future research should extend the population and evidence base. Comparative studies could include active, occasional, former and non-users across degree levels, universities, disciplines, employment destinations and immigration regimes. Such research could examine whether the patterns captured by conditional conversion and verification inequality recur where institutional support, platform access and labour-market infrastructures differ.

Longitudinal and multisource designs could follow students through a complete recruitment cycle, tracing how generated content enters applications, interviews, offers, revised aims or informed withdrawal. Repeated interviews could be combined with voluntarily shared prompts and outputs, records of checking, application materials and perspectives from careers practitioners or employers. This would provide stronger evidence of whether preparatory

action develops into sustained functioning or a change in the set of career options students can realistically pursue.

Intervention research could also examine specific forms of hybrid support. Studies might compare claim-classification prompts, links to official sources, structured practitioner reviews, destination-specific recruitment resources, alumni support and continued postgraduation access to university careers services. Outcomes should distinguish task completion, confidence, representational clarity, judgement quality, career-related action and substantive opportunity rather than combining them into a single measure of benefit.

The broader contribution of this dissertation is to change the evaluative question asked of GenAI career guidance. The relevant question is not simply whether GenAI can produce fluent career content or encourage students to act. It is whether students can establish that the content is truthful, contextually appropriate, consistent with a valued aim and connected to an option they can realistically pursue. GenAI widened access to a first move in career work more consistently than it widened access to substantive career opportunities. Between generated content and career capability lies the social and institutional work of making an option credible, valued and feasible.

Reference List

Akkermans, J., Le Blanc, P., van der Heijden, B. I. J. M., & De Vos, A. (2024). Toward a contextualized perspective of employability development. *European Journal of Work and Organizational Psychology, 33*(1), 1–10. https://doi.org/10.1080/1359432X.2023.2291763

Archibald, M. M., Ambagtsheer, R. C., Casey, M. G., & Lawless, M. (2019). Using Zoom videoconferencing for qualitative data collection: Perceptions and experiences of researchers and participants. *International Journal of Qualitative Methods, 18*, 1–8. https://doi.org/10.1177/1609406919874596

Baek, C., Tate, T. P., & Warschauer, M. (2024). "ChatGPT seems too good to be true": College students' use and perceptions of generative AI. *Computers and Education: Artificial Intelligence, 7*, Article 100294. https://doi.org/10.1016/j.caeai.2024.100294

Bankins, S., Jooss, S., Restubog, S. L. D., Marrone, M., Ocampo, A. C., & Shoss, M. (2024). Navigating career stages in the age of artificial intelligence: A systematic interdisciplinary review and agenda for future research. *Journal of Vocational Behavior, 153*, Article 104011. https://doi.org/10.1016/j.jvb.2024.104011

Barger, A. S. (2025). Artificial intelligence vs. human coaches: Examining the development of working alliance in a single session. *Frontiers in Psychology, 15*, Article 1364054. https://doi.org/10.3389/fpsyg.2024.1364054

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (pp. 610–623). Association for Computing Machinery. https://doi.org/10.1145/3442188.3445922

Berger, R. (2015). Now I see it, now I don't: Researcher's position and reflexivity in qualitative research. *Qualitative Research, 15*(2), 219–234. https://doi.org/10.1177/1468794112468475

Bhaskar, R. (1978). *A realist theory of science* (2nd ed.). Harvester Press.

Bradbury-Jones, C., Taylor, J., & Herber, O. R. (2014). Vignette development and administration: A framework for protecting research participants. *International Journal of Social Research Methodology, 17*(4), 427–440. https://doi.org/10.1080/13645579.2012.750833

Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology, 3*(2), 77–101. https://doi.org/10.1191/1478088706qp063oa

Braun, V., & Clarke, V. (2021). One size fits all? What counts as quality practice in (reflexive) thematic analysis? *Qualitative Research in Psychology, 18*(3), 328–352. https://doi.org/10.1080/14780887.2020.1769238

British Educational Research Association. (2024). *Ethical guidelines for educational research* (5th ed.). https://www.bera.ac.uk/publication/ethical-guidelines-for-educational-research-fifth-edition-2024

Brynjolfsson, E., Li, D., & Raymond, L. R. (2025). Generative AI at work. *The Quarterly Journal of Economics, 140*(2), 889–942. https://doi.org/10.1093/qje/qjae044

Burchardt, T., & Vizard, P. (2011). 'Operationalizing' the capability approach as a basis for equality and human rights monitoring in twenty-first-century Britain. *Journal of Human Development and Capabilities, 12*(1), 91–119. https://doi.org/10.1080/19452829.2011.541790

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*, Article 43. https://doi.org/10.1186/s41239-023-00411-8

Consoli, S. (2024). What motivates Chinese students to study in the UK? A fresh perspective through a 'small-lens'. *Higher Education, 88*(4), 1589–1610. https://doi.org/10.1007/s10734-024-01184-3

Cruickshank, J. (2010). Knowing social reality: A critique of Bhaskar and Archer's attempt to derive a social ontology from lay knowledge. *Philosophy of the Social Sciences, 40*(4), 579–602. https://doi.org/10.1177/0048393109340664

Danermark, B., Ekström, M., Jakobsen, L., & Karlsson, J. C. (2002). *Explaining society: Critical realism in the social sciences*. Routledge. https://doi.org/10.4324/9780203996249

Delva, J., Forrier, A., & De Cuyper, N. (2021). Integrating agency and structure in employability: Bourdieu's theory of practice. *Journal of Vocational Behavior, 127*, Article 103579. https://doi.org/10.1016/j.jvb.2021.103579

Donald, W. E., Mouratidou, M., Hughes, H. P. N., & Padgett, R. C. (2025). Employability, aspirations, and career resources: Reflective diary insights from postgraduate international students. *Higher Education, Skills and Work-Based Learning, 15*(3), 652–667. https://doi.org/10.1108/HESWBL-06-2024-0171

Dwyer, S. C., & Buckle, J. L. (2009). The space between: On being an insider-outsider in qualitative research. *International Journal of Qualitative Methods, 8*(1), 54–63. https://doi.org/10.1177/160940690900800105

Egdell, V., & Robertson, P. J. (2021). A critique of the capability approach's potential for application to career guidance. *International Journal for Educational and Vocational Guidance, 21*(2), 447–463. https://doi.org/10.1007/s10775-020-09445-0

Fakunle, O. (2021). International students' perspective on developing employability during study abroad. *Higher Education Quarterly, 75*(4), 575–590. https://doi.org/10.1111/hequ.12308

Farquhar, S., Kossen, J., Kuhn, L., & Gal, Y. (2024). Detecting hallucinations in large language models using semantic entropy. *Nature, 630*, 625–630. https://doi.org/10.1038/s41586-024-07421-0

Feuerriegel, S., Hartmann, J., Janiesch, C., & Zschech, P. (2024). Generative AI. *Business & Information Systems Engineering, 66*(1), 111–126. https://doi.org/10.1007/s12599-023-00834-7

Fletcher, A. J. (2017). Applying critical realism in qualitative research: Methodology meets method. *International Journal of Social Research Methodology, 20*(2), 181–194. https://doi.org/10.1080/13645579.2016.1144401

Flink, N. A., Black, H. G., & Dingus, R. (2025). Explore, identify, and align: Utilizing generative AI as a career coach. *Marketing Education Review, 35*(2), 109–119. https://doi.org/10.1080/10528008.2024.2430266

Fryer, T. (2022). A critical realist approach to thematic analysis: Producing causal explanations. *Journal of Critical Realism, 21*(4), 365–384. https://doi.org/10.1080/14767430.2022.2076776

Fuertes, V., McQuaid, R., & Robertson, P. J. (2021). Career-first: An approach to sustainable labour market integration. *International Journal for Educational and Vocational Guidance, 21*, 429–446. https://doi.org/10.1007/s10775-020-09451-2

Gedrimiene, E., Celik, I., Kaasila, A., Mäkitalo, K., & Muukkonen, H. (2024). Artificial intelligence (AI)-enhanced learning analytics (LA) for supporting career decisions: Advantages and challenges from user perspective. *Education and Information Technologies, 29*(1), 297–322. https://doi.org/10.1007/s10639-023-12277-4

Gu, Q., & Schweisfurth, M. (2015). Transnational connections, competences and identities: Experiences of Chinese international students after their return 'home'. *British Educational Research Journal, 41*(6), 947–970. https://doi.org/10.1002/berj.3175

Han, H., Park, B., & Seo, K. (2025). A self-determination theory-based career counseling chatbot: Motivational interactions to address career decision-making difficulties and enhance engagement. In *Extended abstracts of the CHI Conference on Human Factors in Computing Systems* (Article 48, pp. 1–9). Association for Computing Machinery. https://doi.org/10.1145/3706599.3720286

Hargittai, E. (2002). Second-level digital divide: Differences in people's online skills. *First Monday, 7*(4). https://doi.org/10.5210/fm.v7i4.942

Hattersley, D. J., & Nicholson, J. (2024). Unpacking the value sought by Chinese international students in UK higher education. *Studies in Higher Education, 49*(12), 2376–2395. https://doi.org/10.1080/03075079.2024.2306347

He, A. J., Zhang, Z., Anand, P., & McMinn, S. (2025). Embracing generative artificial intelligence tools in higher education: A survey study at the Hong Kong University of Science and Technology. *Journal of Asian Public Policy, 18*(2), 352–376. https://doi.org/10.1080/17516234.2024.2447195

He, H., An, L., & Lin, J. (2025). “Haigui” xueli bu zai chixiang?—Laizi jianli toudi shiyan de zhengju [Are overseas-returnee credentials no longer valued? Evidence from a résumé-submission experiment]. *China Economic Quarterly*, 25(4), 909–925. https://doi.org/10.13821/j.cnki.ceq.2025.04.06

Hermanowicz, J. C. (2013). The longitudinal qualitative interview. *Qualitative Sociology, 36*(2), 189–208. https://doi.org/10.1007/s11133-013-9247-7

Higher Education Statistics Agency. (2026). *Higher education student statistics: UK, 2024/25* (Statistical Bulletin SB273). https://www.hesa.ac.uk/news/27-01-2026/sb273-higher-education-student-statistics/location

Hillman, N. (Ed.). (2021). *Paying more for less? Careers and employability support for international students at UK universities* (HEPI Report 143). Higher Education Policy Institute. https://www.hepi.ac.uk/reports/paying-more-for-less-careers-and-employability-support-for-international-students-at-uk-universities/

Holland, A., & Ciachir, C. (2025). A qualitative study of students' lived experience and perceptions of using ChatGPT: Immediacy, equity and integrity. *Interactive Learning Environments, 33*(1), 483–494. https://doi.org/10.1080/10494820.2024.2350655

Holmes, L. (2013). Competing perspectives on graduate employability: Possession, position or process? *Studies in Higher Education, 38*(4), 538–554. https://doi.org/10.1080/03075079.2011.587140

Hornberger, M., Bewersdorff, A., & Nerdel, C. (2023). What do university students know about artificial intelligence? Development and validation of an AI literacy test. *Computers and Education: Artificial Intelligence, 5*, Article 100165. https://doi.org/10.1016/j.caeai.2023.100165

Huang, R., & Turner, R. (2018). International experience, universities support and graduate employability: Perceptions of Chinese international students studying in UK universities. *Journal of Education and Work, 31*(2), 175–189. https://doi.org/10.1080/13639080.2018.1436751

Hughes, R. (1998). Considering the vignette technique and its application to a study of drug injecting and HIV risk and safer behaviour. *Sociology of Health & Illness, 20*(3), 381–400. https://doi.org/10.1111/1467-9566.00107

Hughes, R., & Huby, M. (2004). The construction and interpretation of vignettes in social research. *Social Work and Social Sciences Review, 11*(1), 36–51. https://doi.org/10.1921/swssr.v11i1.428

Jenkins, N., Bloor, M., Fischer, J., Berney, L., & Neale, J. (2010). Putting it in context: The use of vignettes in qualitative interviewing. *Qualitative Research, 10*(2), 175–198. https://doi.org/10.1177/1468794109356737

Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y. J., Madotto, A., & Fung, P. (2023). Survey of hallucination in natural language generation. *ACM Computing Surveys, 55*(12), Article 248. https://doi.org/10.1145/3571730

Jiang, J., & Chen, J. (2026). AI trust and protean career among university students: The mediating role of intrinsic motivation and the moderating role of job insecurity. *Frontiers in Psychology, 16*, Article 1749655. https://doi.org/10.3389/fpsyg.2025.1749655

Jöns, H., Okorie, C. A., Deakin-Smith, H., & Esson, J. (2025). Critical geographies of professional careers: Contributions to the spatial turn in the social sciences. *Globalisation, Societies and Education, 23*(3), 589–603. https://doi.org/10.1080/14767724.2024.2441281

Kaiser, K. (2009). Protecting respondent confidentiality in qualitative research. *Qualitative Health Research, 19*(11), 1632–1641. https://doi.org/10.1177/1049732309350879

Kalai, A. T., Nachum, O., Vempala, S. S., & Zhang, E. (2026). Evaluating large language models for accuracy incentivizes hallucinations. *Nature, 653*, 1047–1051. https://doi.org/10.1038/s41586-026-10549-w

Kallio, H., Pietilä, A.-M., Johnson, M., & Kangasniemi, M. (2016). Systematic methodological review: Developing a framework for a qualitative semi-structured interview guide. *Journal of Advanced Nursing, 72*(12), 2954–2965. https://doi.org/10.1111/jan.13031

Kang, E. (2025). The influence of hukou precarisation on employers' recruitment of Chinese international student returnees in Shanghai. *Compare: A Journal of Comparative and International Education.* Advance online publication. https://doi.org/10.1080/03057925.2025.2571463

Kemp, S. (2005). Critical realism and the limits of philosophy. *European Journal of Social Theory, 8*(2), 171–191. https://doi.org/10.1177/1368431005051762

Kettunen, J., & Sampson, J. P., Jr. (2019). Challenges in implementing ICT in career services: Perspectives from career development experts. *International Journal for Educational and Vocational Guidance, 19*(1), 1–18. https://doi.org/10.1007/s10775-018-9365-6

Kim, D. G., & Moon, A. (2024). From helping hand to stumbling block: The ChatGPT paradox in competency experiment. *Applied Economics Letters, 32*(18), 2669–2673. https://doi.org/10.1080/13504851.2024.2337330

Kim, M., & Cho, E. (2022). Lost in transition: A two-year collaborative autoethnography of South Korean doctoral students' development and identity negotiation. *Journal of International Students, 12*(S2), 50–67. https://doi.org/10.32674/jis.v12iS2.4338

Klingbeil, A., Grützner, C., & Schreck, P. (2024). Trust and reliance on AI—an experimental study on the extent and costs of overreliance on AI. *Computers in Human Behavior, 160*, Article 108352. https://doi.org/10.1016/j.chb.2024.108352

Lee, J. D., & See, K. A. (2004). Trust in automation: Designing for appropriate reliance. *Human Factors, 46*(1), 50–80. https://doi.org/10.1518/hfes.46.1.50_30392

Liu, Y., & Ye, Y. (2021). Factors influencing Chinese returnees' employment quality: Empirical analysis based on ordinal logistic regression and ISM model. *Journal of East China Normal University (Educational Sciences), 39*(12), 42–58. https://doi.org/10.16382/j.cnki.1000-5560.2021.12.004

Loidl, V., Cao, Y., & Huang, W. D. (2026). How do undergraduate students use ChatGPT for career development? An explanatory mixed-methods study of employability-related activities. *International Journal of AI in Pedagogy, Innovation, and Learning Futures, 2*(1). https://doi.org/10.46787/ijaipil.v2i1.7284

Malmqvist, J., Hellberg, K., Möllås, G., Rose, R., & Shevlin, M. (2019). Conducting the pilot study: A neglected part of the research process? Methodological findings supporting the importance of piloting in qualitative research studies. *International Journal of Qualitative Methods, 18*, 1–11. https://doi.org/10.1177/1609406919878341

Malterud, K., Siersma, V. D., & Guassora, A. D. (2016). Sample size in qualitative interview studies: Guided by information power. *Qualitative Health Research, 26*(13), 1753–1760. https://doi.org/10.1177/1049732315617444

Martins, N. (2006). Capabilities as causal powers. *Cambridge Journal of Economics, 30*(5), 671–685. https://doi.org/10.1093/cje/bel012

Maxwell, J. A. (2012). *A realist approach for qualitative research*. SAGE.

Miao, F., & Holmes, W. (2023). *Guidance for generative AI in education and research*. UNESCO. https://doi.org/10.54675/EWZM9535

Nachatar Singh, J. K., Holmes, H.-L., & Gupta, S. (2023). South Asian postgraduate international students' employability barriers: A qualitative study from Australia and the United Kingdom. *British Journal of Educational Studies, 71*(4), 373–391. https://doi.org/10.1080/00071005.2023.2197989

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*, Article 100041. https://doi.org/10.1016/j.caeai.2021.100041

Nguyen, H., & Cheng, M. (2025). Understanding international students' agency in developing employability: Case study of a post-1992 university in the United Kingdom. *British Educational Research Journal.* Advance online publication. https://doi.org/10.1002/berj.70018

Noy, C. (2008). Sampling knowledge: The hermeneutics of snowball sampling in qualitative research. *International Journal of Social Research Methodology, 11*(4), 327–344. https://doi.org/10.1080/13645570701401305

Nussbaum, M. C. (2011). *Creating capabilities: The human development approach*. Belknap Press of Harvard University Press. https://doi.org/10.4159/harvard.9780674061200

OECD. (2024a). *Digital technologies in career guidance for youth: Opportunities and challenges* (OECD Education Policy Perspectives, No. 113). https://doi.org/10.1787/c9ab23da-en

OECD. (2024b). *Challenging social inequality through career guidance: Insights from international data and practice*. https://doi.org/10.1787/619667e2-en

Palinkas, L. A., Horwitz, S. M., Green, C. A., Wisdom, J. P., Duan, N., & Hoagwood, K. (2015). Purposeful sampling for qualitative data collection and analysis in mixed method implementation research. *Administration and Policy in Mental Health and Mental Health Services Research, 42*(5), 533–544. https://doi.org/10.1007/s10488-013-0528-y

Pandya, S. S., & Wang, J. (2024). Artificial intelligence in career development: A scoping review. *Human Resource Development International, 27*(3), 324–344. https://doi.org/10.1080/13678868.2024.2336881

Perrotta, C., Gulson, K. N., Williamson, B., & Witzenberger, K. (2021). Automation, APIs and the distributed labour of platform pedagogies in Google Classroom. *Critical Studies in Education, 62*(1), 97–113. https://doi.org/10.1080/17508487.2020.1855597

Pham, T., Soltani, B., & Nachatar Singh, J. K. (2024). Employability capitals as essential resources for employment obtainment and career sustainability of international graduates. *Journal of Further and Higher Education, 48*(4), 436–448. https://doi.org/10.1080/0309877X.2024.2344771

Pillow, W. S. (2003). Confession, catharsis, or cure? Rethinking the uses of reflexivity as methodological power in qualitative research. *International Journal of Qualitative Studies in Education, 16*(2), 175–196. https://doi.org/10.1080/0951839032000060635

Robertson, P. J., & Picard, F. (2021). An introduction to the special section on the capability approach to career guidance. *International Journal for Educational and Vocational Guidance, 21*(2), 395–404. https://doi.org/10.1007/s10775-021-09462-7

Robeyns, I. (2005). The capability approach: A theoretical survey. *Journal of Human Development, 6*(1), 93–117. https://doi.org/10.1080/146498805200034266

Robeyns, I. (2017). *Wellbeing, freedom and social justice: The capability approach re-examined*. Open Book Publishers. https://doi.org/10.11647/OBP.0130

Sarıköse, S., Şengül, T., Uncu, B., Kirkland-Kyhn, H., & Kaya, N. (2025). AI-guided interview simulation to improve employability and reduce anxiety in final-year nursing and midwifery students: A quasi-experimental study. *Clinical Simulation in Nursing, 109*, Article 101852. https://doi.org/10.1016/j.ecns.2025.101852

Sarmurzin, Y., Azhenov, A., Kudysheva, A., Ualikhanova, B., Saparbekova, I., & Besbayev, A. (2026). Implementation of AI in career counselling for university students: A systematic review. *Frontiers in Education, 11*, Article 1787689. https://doi.org/10.3389/feduc.2026.1787689

Sayer, A. (2000). *Realism and social science*. SAGE. https://doi.org/10.4135/9781446218730

Scandurra, R., Kelly, D., Fusaro, S., Cefalo, R., & Hermannsson, K. (2024). Do employability programmes in higher education improve skills and labour market outcomes? A systematic review of academic literature. *Studies in Higher Education, 49*(8), 1381–1396. https://doi.org/10.1080/03075079.2023.2265425

Sen, A. (1999). *Development as freedom*. Oxford University Press.

Sen, A. (2004). Capabilities, lists, and public reason: Continuing the conversation. *Feminist Economics, 10*(3), 77–80. https://doi.org/10.1080/1354570042000315163

Soltani, B., & Tomlinson, M. (2024). International students' employability: A language socialization perspective. *Education + Training, 66*(2/3), 213–232. https://doi.org/10.1108/ET-07-2022-0265

Squires, A. (2009). Methodological challenges in cross-language qualitative research: A research review. *International Journal of Nursing Studies, 46*(2), 277–287. https://doi.org/10.1016/j.ijnurstu.2008.08.006

Steyvers, M., Tejeda, H., Kumar, A., Belem, C., Karny, S., Hu, X., Mayer, L. W., & Smyth, P. (2025). What large language models know and what people think they know. *Nature Machine Intelligence, 7*, 221–231. https://doi.org/10.1038/s42256-024-00976-7

Stutchbury, K. (2022). Critical realism: An explanatory framework for small-scale qualitative studies or an 'unhelpful edifice'? *International Journal of Research & Method in Education, 45*(2), 113–128. https://doi.org/10.1080/1743727X.2021.1966623

Sugden, R. (1993). Welfare, resources, and capabilities: A review of Inequality Reexamined by Amartya Sen. *Journal of Economic Literature, 31*(4), 1947–1962. https://www.jstor.org/stable/2728332

Sundin, O. (2025). Theorising notions of searching, (re)sources and evaluation in the light of generative AI. *Information Research, 30*(CoLIS), 291–302. https://doi.org/10.47989/ir30CoLIS52258

Tanner, E., & Samson, R. (2026). *Navigating the future: A landscape review of AI in career guidance for young people*. Ada Lovelace Institute; Nuffield Foundation. https://www.adalovelaceinstitute.org/report/navigating-the-future/

Temple, B., & Young, A. (2004). Qualitative research and translation dilemmas. *Qualitative Research, 4*(2), 161–178. https://doi.org/10.1177/1468794104044430

Tomlinson, M. (2017). Forms of graduate capital and their relationship to graduate employability. *Education + Training, 59*(4), 338–352. https://doi.org/10.1108/ET-05-2016-0090

University of Oxford. (2026). *Research ethics policy*. https://governance.admin.ox.ac.uk/sitefiles/research-ethics-policy-4.0-ric-approved-09feb2026.pdf

Urban, M., Děchtěrenko, F., Lukavský, J., Hrabalová, V., Švácha, F., Brom, C., & Urban, K. (2024). ChatGPT improves creative problem-solving performance in university students: An experimental study. *Computers & Education, 215*, Article 105031. https://doi.org/10.1016/j.compedu.2024.105031

van Deursen, A. J. A. M., & van Dijk, J. A. G. M. (2014). The digital divide shifts to differences in usage. *New Media & Society, 16*(3), 507–526. https://doi.org/10.1177/1461444813487959

van Nes, F., Abma, T., Jonsson, H., & Deeg, D. (2010). Language differences in qualitative research: Is meaning lost in translation? *European Journal of Ageing, 7*(4), 313–316. https://doi.org/10.1007/s10433-010-0168-y

Vasconcelos, H., Jörke, M., Grunde-McLaughlin, M., Gerstenberg, T., Bernstein, M. S., & Krishna, R. (2023). Explanations can reduce overreliance on AI systems during decision-making. *Proceedings of the ACM on Human-Computer Interaction, 7*(CSCW1), Article 129, 1–38. https://doi.org/10.1145/3579605

Wang, F., Huang, R., Wu, J., & Lim, W. M. (2025). Employability of international postgraduate taught students in the UK: A systematic literature review using PRISMA. *British Journal of Guidance & Counselling.* Advance online publication. https://doi.org/10.1080/03069885.2025.2496173

Wang, K., Qu, Y.-Y., & Wong, S.-P. (2025). Exploring college students' utilization of generative AI for career information seeking: An integrated model with PLS-SEM and FsQCA approach. *Education and Information Technologies, 30*(14), 20071–20098. https://doi.org/10.1007/s10639-025-13569-7

Westman, S., Kauttonen, J., Klemetti, A., Korhonen, N., Manninen, M., Mononen, A., Niittymäki, S., & Paananen, H. (2021). Artificial intelligence for career guidance—Current requirements and prospects for the future. *IAFOR Journal of Education, 9*(4), 43–62. https://doi.org/10.22492/ije.9.4.03

Wiltshire, G., & Ronkainen, N. (2021). A realist approach to thematic analysis: Making sense of qualitative data through experiential, inferential and dispositional themes. *Journal of Critical Realism, 20*(2), 159–180. https://doi.org/10.1080/14767430.2021.1894909

Xu, C. L. (2021). Time, class and privilege in career imagination: Exploring study-to-work transition of Chinese international students in UK universities through a Bourdieusian lens. *Time & Society, 30*(1), 5–29. https://doi.org/10.1177/0961463X20951333

Yang, W., Xie, Y., & Wu, X. (2026). Can AI be a good counselor? Comparing the effectiveness of AI and human career counseling. *Technology in Society, 86*, Article 103301. https://doi.org/10.1016/j.techsoc.2026.103301

Zhang, J., Elliot, D. L., & Makara, K. A. (2024). Likened to "a boiled egg": A new framework for understanding Chinese postgraduate taught students' transitional experience. *Journal of International Students, 14*(3), 42–60. https://doi.org/10.32674/jis.v14i3.6046

Zhao, X., & Cox, A. (2022). Chinese students' study in the UK and employability: The views of Chinese employers, students and alumni, and UK teachers. *Journal of Education and Work, 35*(4), 422–440. https://doi.org/10.1080/13639080.2022.2073339

Appendix A. Letter of the Research Ethical Approval

Education (Educ) DREC
15 Norham Gardens, Oxford, OX2 6PY

Applicant: [redacted]
Principal Investigator: [redacted]
Department: Education

Study title: Dissertation for MSc Education (Digital and Social Change)
(Version: 1.0)

Ethics reference: Education (Educ) DREC - 2911419

Dear [redacted],

On behalf of the Committee, I confirm that the above research study described in the application and other supporting documentation submitted to the committee has been carefully considered on behalf of the Education (Educ) DREC in accordance with the University's regulations and policy for ethics approval of research involving human participants, human tissue and/or personal data. The opinion is as follows:

Opinion of Research Ethics Committee: Favourable Opinion

Subject to the following conditions:

Decision Date: 19 Mar 2026, 13:59

Opinion End Date: 21 Aug 2026

If favourable, insurance-provided indemnity arrangements will be in place between the decision date and opinion end date and you may now commence your study activities. Should you plan to continue the research beyond the end date above, it is your responsibility to ensure that you request, and receive, an extension (via amendment) from the committee for indemnity to remain in place. You may be required to provide a justification.

Please note the following:

Amendments: Should there be any subsequent changes to the reviewed study, applications for amendments can be made via the Oxford Ethics Application System (Worktribe Ethics).

Reports: Studies considered by OxTREC are expected to submit an *annual progress report* on each anniversary of study approval, until the study is completed. An end of study report is also required.

Audit: This study may be selected for audit at the discretion of the Research Governance, Ethics and Assurance Team.

Data safety: It is the responsibility of the PI to ensure that all data collected during the course of the study is stored and transferred safely and securely in accordance with University requirements. Further guidance and advice are available from the Research Data Team. Additional information is available at University Research Ethics (Please note that this link will not work in Protected View).

Appendix B. Participant Information Sheet

PI name: [redacted]
PI contact details: [redacted]
Researcher contact name and status: [redacted] **MSc student**
Oxford University email address: [redacted]

[**From Generated Content to Career Capability:**

Chinese International Students' Use of GenAI for Career Guidance in UK Higher Education]

PARTICIPANT INFORMATION SHEET

1. **Introductory paragraph**

You are being invited to take part in a research study. Before you decide, it is important for you to understand why the research is being done and what it will involve. Please take time to read the following information carefully and discuss it with others if you wish. Ask us if there is anything that is not clear or if you would like more information. Take time to decide whether you wish to take part.

2. **Why is this research being conducted?**

This research explores how higher education students use artificial intelligence (AI) tools, such as ChatGPT or similar platforms, for career guidance and career-related decision-making. The study aims to understand how students interpret AI-generated advice, how useful or trustworthy they find it, and how AI fits alongside more traditional sources of career support, such as university career services, tutors, peers, or family members.

The study also aims to understand how universities and career services present AI in official guidance materials. Overall, the research seeks to provide a clearer understanding of how AI is becoming part of students' career planning and what this may mean for students' agency, trust, and decision-making.

3. **Why have I been invited to take part?**

You have been invited because you are a higher education student aged 18 or over and may have experience of using, or be familiar with, AI tools for career-related purposes. The study will recruit approximately approximately 10–20 students for follow-up interviews.

Participants may come from different academic disciplines, including social sciences, humanities, business, and STEM-related subjects, in order to capture a range of student experiences.

4. **Do I have to take part?**

No. It is entirely up to you to decide whether to take part. Taking part is voluntary. If you decide to participate, you may withdraw at any time without giving a reason and without any negative consequences.

If you take part in the interview, you may request that your data be withdrawn until 1 June 2026 by contacting the researcher. After this date, the interview data will have been anonymised and incorporated into the analysis, and it may no longer be possible to identify and remove your individual contribution.

For the survey, because responses may be anonymised at the point of collection or shortly afterwards, it may not be possible to withdraw your individual response once it has been submitted.

5. What will happen if I take part in the research?

If you choose to take part, you may be involved in one or both of the following stages:

Stage 1: Online survey

You may be invited to complete an online survey. The survey will ask about whether you use AI tools for career-related purposes, which tools you use, and what kinds of questions or tasks you use them for. The survey should take around 10–15 minutes to complete.

Before beginning the survey, you will be asked to confirm that you have read this information sheet and consent to take part.

Stage 2: Interview

At the end of the survey, you may choose to indicate whether you are willing to take part in a follow-up interview. If selected and if you still wish to participate, you will be invited to a semi-structured interview lasting approximately 45–60 minutes.

The interview may take place online via Microsoft Teams or similar platform, or in person, depending on your location and availability. The interview will focus on your experiences of using AI for career guidance, how you interpret the advice given, whether you trust or question it, and whether it influences your career thinking or decisions.

With your consent, the interview will be audio-recorded so that the researcher can keep an accurate record of the conversation. If you do not wish to be recorded, you will not be interviewed.

You may pause or stop the interview at any time, and you may choose not to answer any question.

6. What are the possible disadvantages and risks in taking part?

There are no foreseeable significant risks associated with this research.

Some participants may feel mildly uncomfortable reflecting on career uncertainty, job pressure, or decision-making. You do not have to answer any question that you would prefer not to answer, and you may stop the survey or interview at any time.

The researcher will take reasonable steps to protect your confidentiality. However, because this is qualitative research, complete anonymity cannot always be absolutely guaranteed, especially if participants describe highly specific experiences. To reduce this risk, identifying details will be removed or altered where possible.

7. Are there any benefits in taking part?

There may be no direct or personal benefit to you from taking part in this research. However, some participants may find it useful to reflect on their experiences of using AI for career guidance.

It is hoped that this research will contribute to a better understanding of how students use AI in career planning and may help inform future university career support and guidance practices.

8. What information will be collected and why is the collection of this information relevant for achieving the research objectives?

The study will collect:

- survey responses about your use of AI for career-related purposes;
- interview data about your experiences, views, and reflections on AI-mediated career guidance;
- limited personal information necessary for arranging interviews, such as your name and email address, if you volunteer for follow-up contact.

This information is relevant because it will help the researcher understand how students use AI for career guidance, how they interpret and evaluate AI advice, and how AI may shape career-related thinking and decision-making.

The researcher and the research supervisor will have access to the research data. If transcription support is used, only authorised persons will have access to the relevant audio files or transcripts.

Identifiable data, such as consent forms and contact details, will be stored securely and separately from the main research data on University-approved secure storage systems. Audio recordings and transcripts will be anonymised as soon as practicable. Research data will be stored securely in accordance with University of Oxford data protection guidelines and retained for the period required by University policy.

If interviews are conducted with participants in China, data may be transferred securely to and stored in the United Kingdom for analysis as part of the University of Oxford research project. Identifiable data will be removed whenever possible, and data will be handled with appropriate safeguards.

9. Will the research be published? Could I be identified from any publications or other research outputs?

The findings from this research may be written up in a dissertation, academic publications, conference presentations, or other research outputs.

Short quotations from interviews may be used in these outputs. If quotations are used, they will normally be anonymised and any identifying details will be removed where possible. Your real name will not be used unless you have explicitly agreed otherwise.

A copy of the dissertation may be deposited in the Oxford University Research Archive.

10. Data Protection

The University of Oxford is the data controller with respect to your personal data, and as such will determine how your personal data is used in the research. The University will process your personal data for the purpose of the research outlined above. Research is a task that is performed in the public interest. Further information about your rights with respect to your personal data is available from the University's Information Compliance website at https://compliance.admin.ox.ac.uk/individual-rights.

11. Who has reviewed this research?

This research has received ethical approval from a subcommittee of the University of Oxford Central University Research Ethics Committee (CUREC).

12. Who do I contact if I have a concern about the research or I wish to complain?

If you have a concern about any aspect of this research, please contact:

[redacted]

[redacted]

or

[redacted]

We will do our best to answer your query and will acknowledge your concern within 10 working days and give you an indication of how it will be dealt with.

If you remain unhappy or wish to make a formal complaint, please contact the University of Oxford Research Governance, Ethics & Assurance (RGEA) team at [redacted].

13. Further Information and Contact Details

If you would like to discuss the research with someone beforehand (or if you have questions afterwards), please contact:

[redacted]
Department of Education
University of Oxford
University email: [redacted]

Supervisor: [redacted]
Department of Education
University of Oxford
University email: [redacted]

Appendix C. Participant Consent Form

PI name: [redacted]
PI contact details: [redacted]
Researcher contact name and status: [redacted]
Oxford University email address: [redacted]

Participant Consent Form

[From Generated Content to Career Capability:

Chinese International Students' Use of GenAI for Career Guidance in UK Higher Education]

Education (Educ) DREC - 2911419

Purpose of Study: This study explores how higher education students use AI tools for career guidance and career-related decision-making.

	Please initial each box if you agree with the statement
I confirm that I have read and understand the information sheet for the above research. I have had the opportunity to consider the information, ask questions and have had these answered satisfactorily.	
I understand that my participation is voluntary and that I am free to withdraw until 1 June 2026 without giving any reason.	
I understand who will have access to personal data provided, how the data will be stored and what will happen to the data at the end of the project.	
I understand that I will not be identifiable in the dissertation or any publications.	
I consent to being audio recorded.	
Use of quotations: Please indicate your preference (select *one* option): a) I do not wish to be quoted. **or** b) I agree to the use of quotations in research outputs if I am not identifiable.	a b

I understand how to raise a concern or make a complaint.	☐
I agree to take part.	☐

Name of participant	dd / mm / yyyy Date	Signature
Name of person taking consent	dd / mm / yyyy Date	Signature