

Students' Non-Disclosure of Generative AI Use in Academic Assessment

Master thesis submitted to Delft University of Technology in partial
fulfilment of the requirements for the degree of

MASTER OF SCIENCE

Management of Technology

Faculty of Technology, Policy and Management

Delft University of Technology

Student name: Finn Timmermans
Student number: 6329934
Thesis defence date: 25 August 2026

Graduation committee

Prof.dr.mr.ir. N. Doorn, Ethics / Philosophy of Technology	Chair
O. Oviedo Trespalacios, Safety & Security Science	First supervisor
Dr. K.L.L. van Nunen, Safety & Security Science	Second supervisor

Preface

This thesis marks the completion of my Master of Science in Management of Technology at Delft University of Technology. By investigating why students do or do not disclose their use of generative AI in assessed work, I had the opportunity to explore a relevant and rapidly evolving topic at the intersection of technology, student behavior, academic integrity, and policy. The research process was both challenging and rewarding, and it provided me with valuable insights that I will carry forward into my future career.

I would like to express my sincere gratitude to my first supervisor, Oscar Oviedo Trespalacios, whose experience, guidance, and valuable insights supported me throughout the entire research process. I particularly appreciated our many thoughtful discussions, which encouraged me to reflect critically and played an important role in shaping and strengthening this thesis. I would also like to thank my second supervisor, Dr. Karolien van Nunen, for her valuable perspectives on AI policy and her support in refining the research, as well as Prof. dr. mr. ir. Neelke Doorn, for dedicating her time and expertise to the evaluation of this thesis. Finally, I would like to thank everyone who supported me throughout this journey. Completing this thesis has been a significant learning experience, and I am grateful for the opportunities, challenges, and insights it has brought.

I hope you enjoy reading this work.

Delft, the Netherlands

August 2026

Finn Timmermans

Executive Summary

Generative artificial intelligence (GenAI) has become embedded in student academic work, used for tasks ranging from generating ideas to refining written text. This has unsettled established assumptions about authorship in assessment, and institutions have responded by requiring students to declare their GenAI use. Such requirements are difficult to enforce, since GenAI use cannot be reliably detected. Whether they work therefore depends largely on students themselves, which makes it important to understand why students might withhold disclosure.

Existing research has mapped how institutions respond to GenAI through policy and guidance. Far less attention has been paid to the behavioral mechanisms behind students' own disclosure decisions. Recent work has called for explanations that go beyond psychological factors and also account for the enforcement environment students face. A second gap concerns the behavior itself. Most studies treat non-disclosure as binary, overlooking that students may acknowledge some GenAI use while withholding meaningful detail about it. A third gap is one of measurement, since the field lacks validated instruments for capturing non-disclosure in this more detailed form.

This study addressed both gaps in two ways. First, it integrated Deterrence Theory with the Theory of Planned Behavior (TPB), extending this model with academic self-efficacy and policy awareness. Second, it separated complete from partial non-disclosure, and past behavior from future intention. It posed the research question: *how do deterrence perceptions, TPB-based psycho-social factors, academic self-efficacy, and policy awareness explain students' non-disclosure of GenAI use in academic assessment contexts?* Deterrence Theory contributed perceived certainty of detection, severity of sanctions, and swiftness of consequences, while TPB contributed attitude, subjective norms, and perceived behavioral control. Academic self-efficacy captured confidence in completing assessed work without GenAI, and policy awareness students' understanding of institutional rules. The scope was perceptual, so the findings speak to how students interpret their governance environment rather than to the objective content of policy.

The study employed a quantitative, cross-sectional survey design to test the hypotheses deductively. The target population consisted of higher education students in the Netherlands, recruited through convenience and snowball sampling via university networks and social media. Data were collected over nine weeks through an anonymous, self-administered online questionnaire, offered in English and Dutch, containing 47 items across twelve constructs and five demographic questions. A total of 589 complete responses were received. After data cleaning, 431 valid responses were retained for analysis. The instrument was validated through exploratory factor and reliability analysis. The model was then tested using correlational analysis and hierarchical regression for each of the four non-disclosure outcomes, followed by an exploratory mediation analysis with bootstrapped confidence intervals.

Non-disclosure was explained mainly by psycho-social and individual factors rather than by the enforcement environment. Attitude was the strongest and most consistent predictor across all four outcomes, with subjective norms close behind. Perceived behavioral control, academic self-efficacy, and policy awareness mattered mainly for the partial forms of non-disclosure. Academic self-efficacy was negatively related to non-disclosure throughout, significantly so for the partial forms and marginally for the full forms, indicating that students who felt able to work without GenAI tended to conceal less. The deterrence perceptions performed less well and departed from theoretical expectations. Certainty and severity did not exert direct effects in the final models. Rather, they were associated with non-disclosure indirectly through attitude and subjective norms, indicating that students' evaluations of concealment and perceptions of its social acceptability were more influential than deterrence perceptions alone. Swiftiness ran opposite to prediction, relating positively to the intention outcomes, most strongly to full non-disclosure intention. The predictors also differed across the four outcomes, most clearly for the deterrence perceptions, which supported the decision to separate the forms of non-disclosure.

The findings should be read in light of several limitations. The cross-sectional design cannot establish causal order. This matters most for the mediation results, since the mediators and the behavior were measured at the same moment. The sample leaned toward master's students and women, limiting generalizability. GenAI use itself was captured only as a set of use types, not as the extent or intensity of that use. The main models also did not control for how heavily students relied on these tools. The observed relationships can therefore not be fully separated from underlying usage patterns. This is clearest for academic self-efficacy. Reliance on self-report is a final concern, given that non-disclosure is a sensitive behavior. Future research could address these issues in several ways. These include longitudinal or experimental designs, measuring actual GenAI use alongside disclosure behavior, and recruiting more representative samples across institutions. Self-report bias could also be reduced through techniques such as list experiments or the randomized response method.

The findings also carry clear implications for practice. Sanctions were not irrelevant, but their influence was indirect, uneven across the two forms of non-disclosure. Enforcement therefore works best when it is understood as a signal that shapes how students evaluate concealment, rather than as a mechanism that suppresses the behavior on its own. Measures that address students' attitudes are consequently the more promising point of intervention. Because detection remains unreliable, genuine accountability is better achieved through assessment design that makes independent mastery visible than through detection software or harsher penalties. Institutions should make disclosure feel safe rather than self-incriminating, for instance through greater transparency about how AI-assisted work is graded. Instructors can reinforce this by signaling that disclosure is expected rather than penalized, which shifts what students perceive as normal. Policy should also state clearly how much detail students are expected to provide about their GenAI use, since greater clarity reduces under-reporting. Finally, students who felt able to complete assessed work without GenAI concealed less, so strengthening that confidence tackles non-disclosure at its source. It gives institutions a supportive approach that works alongside rules and sanctions rather than in place of them.

List of Acronyms

Acronym	Definition
AI	Artificial intelligence
ASE	Academic self-efficacy
ATT	Attitude (toward non-disclosure)
CBS	Centraal Bureau voor de Statistiek (Statistics Netherlands)
CER	Certainty (of detection)
CI	Confidence interval
df	Degrees of freedom
DT	Deterrence Theory
EFA	Exploratory factor analysis
FDI	Full non-disclosure intention
GASE	General Academic Self-Efficacy scale
GenAI	Generative artificial intelligence
HEI	Higher education institution
IER	Insufficient effort responding
IRV	Intra-individual response variability
LLM	Large language model
OECD	Organization for Economic Co-operation and Development
PA	Policy awareness
PBC	Perceived behavioral control
PDI	Partial non-disclosure intention
PFD	Past full non-disclosure
PPD	Past partial non-disclosure
P-P plot	Probability-Probability plot
SD	Standard deviation
SER	Severity (of sanctions)
Sig.	Significance
SOS	Swiftness of sanction
SPSS	Statistical Package for the Social Sciences
STEM	Science, Technology, Engineering, and Mathematics
SUN	Subjective norms
TPB	Theory of Planned Behavior
VIF	Variance inflation factor
ZPRED	Standardized predicted value
ZRESID	Standardized residual

Contents

List of Tables.....	vii
List of Figures.....	ix
1. Introduction	1
1.1 Context and Focus of the Study	1
1.2 Conceptual Model and Research Questions	3
1.3 Scientific Relevance	5
1.4 Practical Relevance	6
1.5 Relevance to Management of Technology	7
1.6 Thesis Outline	7
2. Literature Review and Hypothesis Development	8
2.1. Theoretical Framework	8
2.2. Hypotheses Development	9
3. Methodology	22
3.1. Research Design.....	22
3.2. Sample	23
3.3. Measurement and Survey	27
3.4. Operational Conceptual Model	34
3.5. Data Analysis Approach.....	35
4. Results	38
4.1. Descriptive Statistics and Correlations.....	38
4.2. Prevalence of Full and Partial Non-Disclosure.....	40
4.3. Hierarchical Regression Analysis	41
4.4. Exploratory Mediation Analysis	47
4.5. Overview of Hypothesis Testing.....	47
4.6. Robustness and supplementary analyses	49
5. Discussion and Conclusion	52
5.1. Discussion.....	52
5.2. Answering the Central Research Question	58
5.3. Recommendations	58
5.4. Contributions to Theory and Practice.....	60
5.5. Limitations and Directions for Future Research	62

References	65
Appendix 1 – Survey Questions	72
Appendix 2 – Measurements of the Constructs.....	77
Appendix 3 - Survey Questions, Dutch translation	81
Appendix 4 – Sample Size Calculation.....	86
Appendix 5 – Data cleaning Process	87
Appendix 6 – Outlier Detection.....	107
Appendix 7 – Exploratory Factor Analysis (EFA)	112
Appendix 8 – Regression Assumption: Linearity and Normality	120
Appendix 9 – Regression Assumption: Homoscedasticity	124
Appendix 10 – Regression Assumption: Multicollinearity	126
Appendix 11 – Regression Assumption: Independence	127
Appendix 12 – Comparing Predictors Across the Forms of Non-Disclosure	128
Appendix 13 – Theory individual models.....	130
Appendix 14 – Exploratory Analysis of Indirect Effects.....	133
Appendix 15 – Robustness Check with Demographic Controls	140
Appendix 16 – Response Language.....	142
Appendix 17 – Direct Relationships, Policy Awareness and Deterrence Perceptions	145
Appendix 18 – Qualitative Analysis of the Open-Ended Comment Field.....	147
Appendix 19 – AI Disclosure statement	154

List of Tables

Table 1 – Descriptive statistics regarding demographic characteristics of the sample.....	25
Table 2 – Reported GenAI use cases (multiple responses allowed).....	25
Table 3 – Distribution of the number of reported GenAI use cases per respondent	26
Table 4 – Component structure of Model 1	30
Table 5 – Reliability and item-total statistics for the measurement scales	32
Table 6 – Descriptive statistics and intercorrelations of study variables.....	39
Table 7 – Hierarchical regression results for past full non-disclosure.....	44
Table 8 – Hierarchical regression results for past partial non-disclosure	44
Table 9 – Hierarchical regression results for full non-disclosure intention	45
Table 10 – Hierarchical regression results for partial non-disclosure intention.....	45
Table 11 – Hypothesis overview	48
Table 12 – Survey questions.....	72
Table 13 – Demographic questions	76
Table 14 – Survey questions translated to Dutch	81
Table 15 – Demographic questions translated to Dutch	85
Table 16 – Open-ended responses recategorization, Field of study	90
Table 17 – Field of study recategorization, justification.....	91
Table 18 – Open-ended responses recategorization, Level of study	93
Table 19 – Level of study recategorization, justification.....	94
Table 20 – Data cleaning process	95
Table 21 – Outlier detection overview.....	107
Table 22 – KMO and Bartlett's Test Results for EFA Model	112
Table 23 – Communalities, final solution, Model 1	113
Table 24 – Total variance explained, final solution, Model 1.....	114
Table 25 – Pattern matrix, final solution, Model 1	115
Table 26 - Pattern matrix, initial solution, Model 1.....	116
Table 27 – KMO and Bartlett's Test Results for EFA Model 2.	118
Table 28 – Communalities, final solution, Model 2	118

Table 29 – Total variance explained, final solution, Model 2.....	118
Table 30 – Pattern matrix, Final solution, Model 2.....	119
Table 31 – Pattern matrix, initial solution, Model 2.....	119
Table 32 – Skewness and kurtosis of standardized residuals.....	122
Table 33 – Kolmogorov-Smirnov and Shapiro-Wilk tests of normality	122
Table 34 – Collinearity statistics, non-disclosure models.....	126
Table 35 – Model summary and Durbin-Watson statistic, PFD model.....	127
Table 36 – Model summary and Durbin-Watson statistic, PPD model.....	127
Table 37 – Model summary and Durbin-Watson statistic, FDI model.....	127
Table 38 – Model summary and Durbin-Watson statistic, PDI model.....	127
Table 39 – Global test of coefficient equality across all four outcomes.....	128
Table 40 – Per-predictor tests of coefficient equality across the four outcomes.....	129
Table 41 – Averaged full-versus-partial contrast, by predictor.....	129
Table 42 – Difference-in-Differences: Full vs. Partial Across Past and Intention	129
Table 43 – Model fit for the deterrence, TPB, and combined regression models	131
Table 44 – Deterrence theory regression model.....	131
Table 45 – Theory of Planned Behavior regression model.....	132
Table 46 – Combined deterrence and TPB regression model	132
Table 47 – Indirect effects of certainty of sanctions on non-disclosure	136
Table 48 – Indirect effects of severity of sanctions on non-disclosure.....	136
Table 49 – Indirect effects of swiftness of sanctions on non-disclosure	137
Table 50 – Indirect effects of policy awareness on non-disclosure.....	138
Table 51 – Indirect effects of academic self-efficacy on non-disclosure.....	138
Table 52 – Change in explained variance when demographic controls were added	141
Table 53 – Demographic controls in the full model	141
Table 54 - Group Differences by Survey Language (Independent-Samples t-Tests).....	142
Table 55 – Group-wise regression, past behavior outcomes.....	143
Table 56 – Group-wise regression, intention outcomes.....	144
Table 57 – Simple linear regressions of policy awareness on each deterrence perception	145

Table 58 – Thematic framework for the open-ended comments.....	148
Table 59 – Distribution of themes across the 68 comments.....	149

List of Figures

Figure 1 – Conceptual Model	4
Figure 2 – Enhanced Conceptual Model.....	21
Figure 3 – Operational conceptual model	34
Figure 4 – Hayes’ (2013) PROCESS macro Model 4 for mediation.....	36
Figure 5 – G* Power 3.1, sample size calculation.....	86
Figure 6 – Boxplot PFD_mean.....	108
Figure 7 – Boxplot PPD_mean.....	108
Figure 8 – Boxplot FDI_mean.....	108
Figure 9 – Boxplot PDI_mean.....	109
Figure 10 – Boxplot CER_mean.....	109
Figure 11 – Boxplot SER_mean.....	109
Figure 12 – Boxplot SOS_mean.....	110
Figure 13 – Boxplot ATT_mean.....	110
Figure 14 – Boxplot SUN_mean	110
Figure 15 – Boxplot PBC_mean.....	111
Figure 16 – Boxplot PA_mean	111
Figure 17 – Boxplot ASE_mean.....	111
Figure 18 – Normal P-P plot and histogram of standardized residuals, PFD model.....	120
Figure 19 – Normal P-P plot and histogram of standardized residuals, PDI model.....	121
Figure 20 – Normal P-P plot and histogram of standardized residuals, FDI model.....	121
Figure 21 – Normal P-P plot and histogram of standardized residuals, PPD model.....	121
Figure 22 – Residual scatterplot, PFD model.....	124
Figure 23 – Residual scatterplot, PPD model.....	125
Figure 24 – Residual scatterplot, FDI model	125

Figure 25 – Residual scatterplot, PDI model	125
Figure 26 - Conceptual model, tested relationships, exploratory mediation analysis	133
Figure 27 – Conceptual model, showing which indirect pathways were supported	135

1. Introduction

This chapter introduces the context and the focus of the study. It then presents the conceptual model, the main research question, and the sub-questions. Finally, it describes the scientific and practical relevance of the study, its relevance to the Management of Technology programme, and the structure of the thesis.

1.1 Context and Focus of the Study

Generative artificial intelligence (GenAI), and in particular large language model (LLM) tools, has become deeply embedded in students' academic work in only a few years (Freeman, 2025; Ravšelj et al., 2025). Its adoption is now close to universal. Recent global figures indicate that roughly nine in ten students use AI in their studies (Digital Education Council, 2026). For many, generative AI has become a regular part of academic work, with over half of students reporting weekly use (Stephenson & Armstrong, 2026). These tools support tasks ranging from generating ideas to structuring assignments and refining written output (Moorhouse et al., 2023). This is not a local trend, but a broad shift in how students across countries and institutions approach their work (Ravšelj et al., 2025).

This shift matters because of the fundamental purpose of assessment. Assessments exist to establish whether students have mastered the skills and knowledge a course sets out to teach (Biggs et al., 2022). Valid assessment requires teachers to be able to judge the extent to which submitted work reflects students' own abilities. Generative AI challenges this by producing fluent text that is nearly indistinguishable from human-authored work (Hutson, 2024; Luo, 2024). As a result, teachers need to understand not only whether AI was used, but also how, for what purposes, and to what extent. Only then can the contribution of GenAI be separated from the student's own achievement (Coates et al., 2025). Seen this way, the central concern is not plagiarism alone, but the validity of assessment itself. This is why academic integrity has become a central theme in research on GenAI in higher education institutions (HEIs) (Bittle & El-Gayar, 2025; Coates et al., 2025; Ravšelj et al., 2025).

One way to establish authorship is to ask students to disclose their use of GenAI directly. Higher education institutions increasingly require students to report whether, how, and to what extent GenAI was used in assessed work (An et al., 2025; Moorhouse et al., 2023). In this form of self-disclosure, students retain control over what information is revealed about their use of GenAI. Institutions rely on this approach because external detection methods cannot provide certainty. AI detectors generate false positives and often perform poorly when distinguishing between human-written, AI-generated, and hybrid texts (Ardito, 2025; Deep et al., 2025). Human judgment offers little certainty either, as educators have been shown to identify AI-generated text unreliably and with considerable error (Fleckenstein et al., 2024; Kumar & Mindzak, 2024). Consequently, the success of this governance mechanism depends on the honesty of the very individuals it seeks to hold accountable.

The governance surrounding these requirements is far from uniform. In a review of the world's top-ranking universities, fewer than half had publicly available assessment guidelines that addressed generative AI (Moorhouse et al., 2023). Where guidelines did exist, they varied widely in emphasis, specificity, and clarity, leaving considerable room for interpretation (Moorhouse et al., 2023; Tan et al., 2024). In some institutions this effect is amplified by course-level variations as institutions often place responsibility on instructors to define GenAI rules for specific assessments (Adnin et al., 2025; Gonsalves, 2025). The result is a patchwork of governance in which GenAI use is regulated unevenly across institutions and courses (An et al., 2025; Moorhouse et al., 2023). Showing variation in whether they exist, how specific they are, what kinds of GenAI use they allow, what students are expected to disclose, and how consistently these expectations are applied in assessment (Moorhouse et al., 2023).

Within this evolving landscape, disclosure has emerged as the main governance mechanism (Gonsalves, 2025; Weaver, 2024). Many institutions and researchers frame disclosure as a transparency tool that clarifies the role of GenAI in assessed work (An et al., 2025; Schilke & Reimann, 2025). For students, however, transparency is not cost-free. Experimental evidence shows that disclosing AI use can lower trust in the person who discloses, compared with saying nothing (Schilke & Reimann, 2025). The problem is deepened by a lack of two-way transparency. Institutions are often not open with students about how AI-assisted work is graded, which further erodes students' trust in the system (Luo, 2025). A mechanism meant to create clarity can therefore induce negative outcomes that make it feel counterproductive.

Under these conditions, non-disclosure can become a deliberate choice rather than an oversight. In a case study by Gonsalves (2025), students frequently read disclosure as a form of self-incrimination rather than neutral transparency, which made withholding it a rational response to perceived risk. Adnin et al. (2025) reinforce this picture, showing that students actively manage whether, when, and how much they disclose. When they anticipate negative consequences, they often adopt strategies that make their use of GenAI less visible (Adnin et al., 2025). Seen this way, non-disclosure is better understood as a strategic adaptation to governance signals than as a simple act of dishonesty (Adnin et al., 2025; Gonsalves, 2025).

If students hide their GenAI use on purpose, explaining that choice means looking at two things: the rules and sanctions they face, and their own reasons for concealing. Deterrence Theory addresses the first, framing disclosure as a matter of rational choice. It assumes that students act as rational agents who weigh the costs and benefits of concealment, so their decision depends on the certainty of being caught, the harshness of punishment, and how quickly consequences follow (Gibbs, 1979). The Theory of Planned Behavior (TPB) addresses the second. It explains deliberate behavior through a person's attitudes, the social pressure they feel, and how much control they believe they have (Ajzen, 1991). Using both theories responds to recent calls in the disclosure literature to look beyond psychological factors alone and take the enforcement context into account (Gonsalves, 2025; Yusuf et al., 2025). Together, they help to showcase the intricacies of a behavior that is both affected by institutional pressure and individual choice.

These issues together point to the problem that motivates this study. Because GenAI use cannot be reliably detected, institutions increasingly rely on disclosure requirements to preserve transparency in assessment. Their effectiveness, however, depends on students choosing to disclose their use honestly. Yet many students either withhold disclosure entirely or reveal only part of their GenAI use, and the factors driving these decisions remain poorly understood. Two further gaps make the problem harder to address. First, non-disclosure is not a single act but a range of behaviors, from full concealment to the selective omission of detail, a distinction that existing research rarely draws. Second, the field lacks validated instruments for measuring non-disclosure, which makes the behavior difficult to study or monitor with precision. Institutions are therefore expected to protect assessment validity through a mechanism that depends on student behavior they cannot reliably monitor and do not yet fully understand. This study takes up that problem in the Dutch higher education context, drawing on one of the largest samples collected in this area. It conceptualizes and measures non-disclosure more precisely, and it examines the mechanisms that drive it.

1.2 Conceptual Model and Research Questions

This study develops a conceptual model to explain students' non-disclosure of GenAI use in assessed academic work. The model draws on Deterrence Theory and the Theory of Planned Behavior (TPB). It brings together governance-related and psycho-social mechanisms that may shape disclosure decisions, and extends them with individual-level differences between students (Gonsalves, 2025; Nguyen-Phuoc et al., 2024). A further feature of the model is how the outcome is defined. Non-disclosure is treated not as a single binary act, but as a set of related behaviors. The model distinguishes full from partial non-disclosure, and past behavior from future intention (developed further in Chapter 2).

From a deterrence perspective, disclosure decisions may depend on how students perceive the enforcement environment (Nguyen-Phuoc et al., 2024). Deterrence Theory centers on three perceptions: certainty of detection, severity of consequences, and swiftness of sanctions (Gibbs, 1979). From a TPB perspective, disclosure-related behavior is shaped by attitudes, subjective norms, and perceived behavioral control (Ajzen, 1991). The model extends these with two individual-level factors. The first is academic self-efficacy. It reflects students' confidence in completing academic work without the use of GenAI, and therefore in reaching the learning objectives an assessment is meant to certify. The second is policy awareness. It captures how well students believe they understand institutional rules and expectations regarding GenAI disclosure.

The outcome in the model is students' non-disclosure of GenAI use, examined across its four forms: past full non-disclosure, past partial non-disclosure, full non-disclosure intention, and partial non-disclosure intention. All eight predictors are modeled as direct effects on non-disclosure. Students may be less inclined to withhold disclosure when they perceive detection and sanctions as more certain, severe, or swift (Gibbs, 1979). Their attitudes, perceived social expectations, and sense of control shape how they then act (Ajzen, 1991). Academic self-efficacy and policy awareness enter as two further direct predictors. They capture,

respectively, students' confidence in working without GenAI and their understanding of the rules that apply.

The study also considers whether these effects may partly operate indirectly. Deterrence perceptions and policy awareness might first shape the psycho-social evaluations in the model, such as students' attitudes toward non-disclosure, which then shape behavior. The data are cross-sectional, so they cannot establish the temporal order that mediation implies. This possibility is therefore not stated as a formal expectation. It is examined separately, as an exploratory analysis. Its results show whether the observed pattern is consistent with an indirect pathway, not that one operates (Bryman, 2016).




Figure 1 – Conceptual Model

This study addresses the following research question:

How do deterrence perceptions, TPB-based psycho-social factors, academic self-efficacy, and policy awareness explain students' non-disclosure of GenAI use in academic assessment contexts?

In line with this research question and the conceptual model, the study addresses six sub-questions:

SQ1. How can students' non-disclosure of GenAI use be conceptualized and measured as more than a binary act, distinguishing full from partial non-disclosure?

SQ2. To what extent do the newly developed non-disclosure measures demonstrate construct validity and internal consistency?

SQ3. How do the perceived certainty, severity, and swiftness of sanctions relate to students' non-disclosure of GenAI use?

SQ4. How do attitude, subjective norms, and perceived behavioral control relate to students' non-disclosure of GenAI use?

SQ5. How do academic self-efficacy and policy awareness relate to students' non-disclosure of GenAI use?

SQ6. Which Theory of Planned Behavior constructs mediate the relationships between deterrence perceptions, policy awareness, academic self-efficacy, and non-disclosure?

1.3 Scientific Relevance

Research on GenAI in higher education has expanded rapidly (Bittle & El-Gayar, 2025; Yusuf et al., 2025). Within this growing literature, work on academic integrity has largely focused on descriptive analyses and institutional responses rather than on the student behavior those responses are meant to govern (Bittle & El-Gayar, 2025; Luo, 2024). Two gaps stand out within this literature. The first concerns how non-disclosure is conceptualized. The second concerns how its behavioral drivers are explained.

The first gap concerns the conceptualization of non-disclosure itself. Much of the existing literature treats disclosure as a binary outcome, separating only those who disclose from those who do not (Adnin et al., 2025; Gonsalves, 2025; Luo, 2025; Qu et al., 2025; Schilke & Reimann, 2025; Yusuf et al., 2025). This approach overlooks a potential middle ground, in which students may acknowledge some GenAI use while withholding important details about its extent, type, or location. This study therefore treats non-disclosure as more than a binary act. It distinguishes full from partial non-disclosure while measuring both past behavior and future intention. A related problem is measurement. The field currently lacks validated instruments for capturing non-disclosure in this more precise form. This study develops and validates such measures, and the distinction is elaborated in the literature review (see Chapter 2). Establishing this conceptual and measurement basis is a central contribution of the thesis.

The second gap concerns explanation rather than description. Prior work has largely mapped the policy landscape in descriptive terms, and clear disclosure requirements are in fact often uneven or absent (Moorhouse et al., 2023). The gap, however, is broader than whether such policies exist. Even where a disclosure requirement is in place, the behavioral mechanisms that determine how students respond to it remain largely untested in quantitative terms. This study focuses on precisely those settings, where disclosure is required, and examines the behavior that determines whether the requirement works. This matters because the relevant behavior is not GenAI use itself, but the non-disclosure of that use in assessment. Non-compliance with AI declarations can be widespread and often intentional, shaped by fear of consequences, inconsistent enforcement, and peer influence (Adnin et al., 2025; Gonsalves, 2025). Gonsalves (2025) argues that TPB should be integrated with other theories to capture the complexities of GenAI disclosure. Yusuf et al. (2025) reinforce this, calling for attention to external influences such as institutional policies, cultural norms, and peer pressure alongside psychological predictors.

This study addresses the second gap by integrating the Theory of Planned Behavior with Deterrence Theory. The model links non-disclosure to students' internal evaluations and to their perceptions of the enforcement environment (Nguyen-Phuoc et al., 2024). In doing so, it responds to recent calls for disclosure research that moves beyond purely psychological explanations (Gonsalves, 2025; Yusuf et al., 2025). Academic self-efficacy and policy

awareness extend the model further. They account for students' confidence in completing assessments without GenAI, and for their understanding of institutional GenAI rules. Together, these constructs help explain the more nuanced aspects of students' non-disclosure behavior in higher education.

1.4 Practical Relevance

As GenAI has become widely adopted in higher education, institutions have come under growing pressure to protect academic integrity through policy and assessment guidance (An et al., 2025; Luo, 2024). What is ultimately at stake is the purpose of assessment itself. Assessments are meant to certify that students have reached the intended learning objectives, and teachers can only judge this when they know how the work was produced (Biggs et al., 2022; Moorhouse et al., 2023). When students conceal their GenAI use, that judgment breaks down. The assessment can then no longer show whether the learning objectives were met. Reducing non-disclosure is therefore not an administrative concern, but a condition for assessment to do its job.

In practice, however, institutional responses are inconsistent and can work against this goal. Policy and guidance differ across institutions, courses, and assessments, both in whether they exist and in what students are expected to disclose (Adnin et al., 2025; An et al., 2025; Gonsalves, 2025; Moorhouse et al., 2023). Students and teachers also differ in how they view GenAI use and disclosure (Adnin et al., 2025; Barrett & Pack, 2023). At the same time, disclosure can carry a cost. It can lower trust in the student who discloses, and unclear grading practices further weaken students' trust in the system (Luo, 2025; Schilke & Reimann, 2025). Faced with these conditions, students may respond strategically, either hiding their GenAI use or managing what they reveal (Adnin et al., 2025; Qu et al., 2025). Enforcement and disclosure requirements alone are therefore unlikely to secure honest reporting.

This is where the study offers practical value. By identifying the factors that shape non-disclosure, it can help institutions direct their efforts where those efforts are most likely to work, rather than relying on detection and disclosure rules alone (Gonsalves, 2025; Nguyen-Phuoc et al., 2024). To do this, the study assesses how much each type of factor contributes to non-disclosure. It compares the explanatory weight of enforcement-based perceptions, psycho-social evaluations, and individual differences between students. This indicates which kind of lever institutions can most usefully prioritize, from stronger enforcement to changing how students view disclosure. Distinguishing full from partial non-disclosure sharpens this further, since under-reporting and complete concealment may call for different responses. Taken together, these insights can help higher education institutions design disclosure systems that are clearer, more credible, and better aligned with the goal of valid assessment (Gonsalves, 2025; Moorhouse et al., 2023).

1.5 Relevance to Management of Technology

The topic of this thesis sits at the core of the Management of Technology (MOT) programme, which studies how organizations adopt emerging technologies and manage what follows from that adoption. GenAI is a clear example of a technology that has spread faster than the institutions using it have been able to adapt. Universities have responded with disclosure rules. These rules, however, apply to a technology whose use is hard to observe and whose limits remain unclear. The present study examines a behavioral consequence of this mismatch: when governance is unclear, users may adopt the technology while concealing that they have done so. This reflects a core idea in MOT: the value and risks of a technology depend on how people use it, not only on what it can do.

The study also speaks to the governance of emerging technology and the limits of control. Organizations typically manage technology use through policies, rules, and enforcement. Deterrence Theory represents this logic in its clearest form, as it assumes that the perceived certainty, severity, and swiftness of sanctions will discourage undesirable behavior. The findings of this study indicate that such formal mechanisms explain only part of the picture, and that attitudes and self-efficacy carry substantial weight. For technology managers, this suggests that rules alone do not produce transparent use, and that governance must also address how users view the behavior. Higher education serves as the empirical setting here. The underlying problem, namely how to govern a technology whose use cannot easily be verified, applies to organizations more broadly.

Finally, the approach taken in this thesis reflects the methodological orientation of the programme. Established behavioral theories are integrated into a single model, operationalized into measurable constructs, and tested quantitatively on survey data, which is the type of evidence-based analysis that MOT prepares its graduates to conduct.

1.6 Thesis Outline

The remainder of this thesis is organized into four chapters. Chapter 2 presents the theoretical framework, reviewing the literature on disclosure in GenAI assessment contexts and developing the hypotheses and the conceptual model. Chapter 3 describes the research methodology, covering the research design, the sample, the measurement instrument and its validation, and the data analysis strategy. Chapter 4 reports the results of the descriptive and regression analyses, followed by the exploratory mediation analysis, together with an overview of the hypothesis testing. Chapter 5 discusses these findings in relation to prior research, answers the central research question, and closes with the contributions, limitations, and directions for future research.

2. Literature Review and Hypothesis Development

This chapter outlines the theoretical framework of the study and the development of hypotheses.

2.1. *Theoretical Framework*

Explaining why students withhold information about their GenAI use requires a framework that captures two influences at once. The first is the governance environment in which the decision is made. Students face rules, detection efforts, and possible sanctions. Together, these external conditions define the potential risks and consequences associated with concealment. The second is the internal process by which a student evaluates these factors and decides how to act. A framework based on only one of these would leave the other unexplained. For this reason, the study does not rely on a single existing theory. It develops an integrated framework by combining two established theories that explain both aspects of the decision-making process. The two theories are Deterrence Theory and the Theory of Planned Behavior (TPB).

Deterrence Theory speaks to the governance side of the decision. It proposes that people are less likely to break a rule when they expect violations to be detected, sanctioned, and enforced quickly (Gibbs, 1979). These three perceptions are the certainty of detection, the severity of sanctions, and the swiftness of consequences. They apply directly to the GenAI assessment setting, where disclosure rules are backed by detection tools and disciplinary procedures (Luo, 2025; Zhang et al., 2025). The theory thus shows how the enforcement environment influences a student's reasoning. However, the theory primarily explains non-disclosure in terms of its anticipated costs and benefits. It offers limited insight into how students interpret and evaluate the behavior itself (Nguyen-Phuoc et al., 2024).

The Theory of Planned Behavior addresses what Deterrence Theory leaves open. It explains behavior as the result of an intention, which is formed by three components: a person's attitude toward the behavior, the subjective norms they perceive, and their perceived behavioral control (Ajzen, 1991). TPB has been used widely to study academic conduct such as cheating and plagiarism, and recent work applies it to GenAI disclosure (Gonsalves, 2025; Huang et al., 2025). It captures how students evaluate concealment, interpret their peers' expectations, and assess their own ability to disclose or withhold their GenAI use. Its limitation mirrors that of Deterrence Theory. On its own, TPB gives insufficient attention to the external enforcement environment that makes the decision matter (Gonsalves, 2025).

Combining the two theories is a deliberate choice. Each theory covers what the other misses: Deterrence Theory models the pressures students face, while TPB models how those pressures become behavior. This combination responds to recent calls to connect psychological explanations of GenAI disclosure with the enforcement context (Gonsalves, 2025; Nguyen-Phuoc et al., 2024). This integration forms the core of the framework.

The integrated model still treats students as alike, yet disclosure decisions clearly differ between students. Two individual-level differences are therefore added. The first is academic

self-efficacy, a student's confidence in completing assessed work without GenAI (Pajares, 1996). This links the framework to the purpose of assessment. A student who doubts they can meet the learning objectives independently of GenAI has more reason to use it, and in turn to hide that use (Zhang et al., 2024). The second is policy awareness, or how well students believe they understand the institutional rules on GenAI use and disclosure (McCabe et al., 2001). Together, these additions explain why students under the same rules may still disclose differently. The following sections develop each construct and derive the study's hypotheses.

2.2. Hypotheses Development

This section provides a review of relevant literature to elaborate on the relationships between the constructs within the conceptual model and to derive the study's hypotheses.

2.2.1. Disclosure in GenAI Assessment Contexts

Disclosure has long functioned as a core governance mechanism in higher education because it makes the origins of contributions to academic work visible and therefore assessable (Weaver, 2024). In academic assessment, disclosure practices are closely tied to maintaining integrity by clarifying what portion of a submission reflects the student's own achievement (Moorhouse et al., 2023). In this sense, disclosure is not only a communication standard but also an institutional tool for protecting assessment validity and academic standards (Moorhouse et al., 2023).

With the emergence of GenAI, institutions have started to extend this logic by requiring students to declare AI involvement in assessed work (Coates et al., 2025; Moorhouse et al., 2023). Moorhouse et al. (2023) show that many university guidelines treat disclosure as a transparency tool that clarifies how GenAI contributed to a student's work. Requiring a declaration, however, does not determine how it should be used. Guidelines typically present disclosure as a matter of transparency without specifying how acknowledged GenAI use should affect grading. As a result, teachers may be uncertain about how to respond once a declaration is made (den Hollander et al., 2025). Beyond this, the disclosure environment in higher education is rarely uniform (Moorhouse et al., 2023). At the institutional level, early reviews of the top universities show that GenAI guidelines are unevenly available and that differences exist in emphasis and specificity (Coates et al., 2025; Moorhouse et al., 2023). This unevenness creates variation in the interpretation of the requirements of disclosure and the expected level of detail (Gonsalves, 2025; Moorhouse et al., 2023). At the course level, instructors often differ in what they expect students to disclose, so that expectations can vary not only between programmes but between courses within the same programme (Adnin et al., 2025; Gonsalves, 2025). In other words, disclosure operates in a landscape that may be fragmented, locally interpreted, and inconsistently enforced (Gonsalves, 2025; Moorhouse et al., 2023).

This policy ambiguity is also visible in the Dutch higher education context. Survey evidence from Dutch universities suggests that both students and teachers already use AI tools, while simultaneously showing concerns about reliability and ethics of AI-assisted work (Bos et al., 2025). This highlights the need for clear norms, guidance and governance

structures (Bos et al., 2025). At the same time, research indicates that such governance is still fragmented (den Hollander et al., 2025). Decisions about acceptable GenAI use are often made at the faculty, program, or course level rather than through a unified institutional framework (den Hollander et al., 2025). As a result, policy practices often vary across contexts and may even lead to situations in which no clear policy exists at all (den Hollander et al., 2025).

This fragmentation is especially visible in assessment. In Dutch context, course coordinators already report uncertainty about detecting AI use in student work (den Hollander et al., 2025). This concern is consistent with broader research showing that AI detectors can produce false positives and perform poorly when asked to distinguish between human, AI-generated, and hybrid texts (Ardito, 2025; Deep et al., 2025).

The literature further suggests that disclosure does not appear to operate as a neutral act of transparency in this setting (Adnin et al., 2025; Gonsalves, 2025). Disclosing AI use can carry social and relational costs for the disclosing student (Schilke & Reimann, 2025). Schilke and Reimann (2025) show that disclosing AI use can reduce trust in the disclosing actor via lower perceived legitimacy, implying that transparency can carry reputational penalties. Disclosure can also be perceived as risky and self-incriminating when policy landscapes are ambiguous and enforcement is perceived as punitive (Gonsalves, 2025; Luo, 2025). This results in students intentionally and actively managing whether, when, and how to disclose GenAI use (Adnin et al., 2025). Concealment can therefore be understood as a strategic response that balances transparency against fears of negative consequences and distrust (Adnin et al., 2025; Qu et al., 2025).

These dynamics motivate a sharper conceptualization of the dependent variable in this study. Existing GenAI disclosure research commonly treats (non-)disclosure as a binary outcome, distinguishing between cases where agents disclose their GenAI use and cases where they do not disclose it (Adnin et al., 2025; Gonsalves, 2025; Luo, 2025; Schilke & Reimann, 2025; Yusuf et al., 2025). This binary approach may understate the complexity of disclosure behavior, as concealment often involves multiple strategies and selective self-presentation (Adnin et al., 2025; Gonsalves, 2025). Because disclosure can carry negative consequences, students may disclose “low-risk” GenAI use while withholding “high-risk” use, making partial disclosure a meaningful intermediate behavior that should be distinguished from complete non-disclosure (Barrett & Pack, 2023).

Building on this, the present study conceptualizes non-disclosure as two related but distinct behaviors. First, complete non-disclosure, where students provide no acknowledgement of GenAI involvement. Second, partial non-disclosure (under-disclosure), where students provide some acknowledgement but selectively omit material details about extent, type, or location of GenAI use. This distinction reflects the idea that students may seek to appear compliant while still managing exposure under perceived risk of negative impact (Adnin et al., 2025; Gonsalves, 2025).

Because disclosure is sensitive and difficult to observe directly, this study conceptualizes non-disclosure as a broader construct that includes both past non-disclosure behavior and non-disclosure intention (Ajzen, 1991; Nguyen & Goto, 2024). Measuring past behavior is important as it captures whether students have previously used concealment in academic contexts, rather than only expressing a willingness to do so (Adnin et al., 2025; Ajzen, 1991). At the same time, intention remains central in this study, since the Theory of Planned Behavior identifies intention as the most direct predictor of behavior (Ajzen, 1991). Ajzen (1991) notes that including past behavior in the prediction equation can help test the sufficiency of the theory. This makes it meaningful to examine prior non-disclosure together with future intention (Ajzen, 1991).

Including both past behavior and intention is also methodologically useful. Past behavior shows what students have already done in real situations. Intention shows whether they are willing to do the same in the future (Ajzen, 1991; Nguyen & Goto, 2024). This distinction is important because self-reports of past misconduct may be biased. Nguyen and Goto (2024) found that students may hide past AI-assisted misconduct when asked directly, reporting a significant difference between stated behavior and actual past behavior. In contrast, this biased reporting was not found for the measure of intended misconduct behavior (Nguyen & Goto, 2024). For this reason, intention is a useful additional indicator. It helps capture likely future behavior, especially when past behavior is socially sensitive and hard to measure without bias (Ajzen, 1991; Nguyen & Goto, 2024).

The distinction between full and partial non-disclosure raises two questions central to this study. The first issue is how common each form of non-disclosure is. Partial non-disclosure allows a student to appear compliant while limiting exposure. Acknowledging some GenAI use can signal good faith and make any later discovery seem less like deliberate deception. Under ambiguous rules and uncertain detection, this makes it a lower-risk strategy than withholding all information (Barrett & Pack, 2023). Partial non-disclosure is therefore expected to occur more often than full non-disclosure.

H1: Partial non-disclosure of GenAI use is more prevalent than full non-disclosure.

The second question is whether the two forms behave alike. If they are distinct behaviors rather than points on one scale, their drivers need not be the same. Full concealment is the higher-stakes, more clearly sanctionable act, so enforcement-based perceptions may weigh on it more heavily, while attitudes and social norms may shape both (Adnin et al., 2025; Barrett & Pack, 2023). Because the direction of these differences cannot be set in advance, this study predicts no specific pattern, testing instead whether predictors differ across all four outcomes.

H2: The predictors of non-disclosure differ across the four forms of non-disclosure, in particular between full and partial non-disclosure.

2.2.2. Deterrence Theory and Non-Disclosure

Deterrence theory (DT), stemming from the works of Jeremy Bentham and Cesare Beccaria, explains that individuals weigh perceived costs and benefits as a rational response to rule compliance (Lou et al., 2025). In simple terms, DT predicts that people are less inclined to break rules based on three perceptions: certainty, severity, and swiftness of sanctions. Applied to GenAI disclosure in assessment, DT is relevant because GenAI disclosure requirements create an enforcement context in which students may evaluate whether hiding AI use is “worth it,” given detection mechanisms and potential consequences.

Certainty

In a disclosure context, certainty refers to whether students believe undisclosed GenAI use is likely to be detected through AI detectors, instructor suspicion, or follow-up checks (Zhang et al., 2025). Several studies support the classic deterrence expectation that higher perceived detection probability reduces misconduct. Zhang et al. (2025) find that higher punishment certainty is associated with lower intention to engage in AI-based plagiarism. In a closely related GenAI policy setting, Lou et al. (2025) similarly show that perceived sanction certainty strengthens GenAI policy compliance intention, suggesting that certainty can function as a deterrent. Huang et al. (2025) add that perceptions of detection difficulty and the presence of AI detectors shape students’ GenAI cheating likelihood, reinforcing the idea that perceived detectability matters for behavior.

In the context of GenAI in assessment, certainty is difficult to establish because AI use cannot always be detected reliably. AI detection tools have increasingly been adopted by higher education institutions to identify potential misuse of GenAI in assessment (Deep et al., 2025; Hadra et al., 2026). However, a wide range of research suggests that these tools remain affected by false positives, inconsistent accuracy, and low transparency, making detection less certain than deterrence models might suggest (Ardito, 2025; Deep et al., 2025; Hadra et al., 2026). Their reliability further declines when assessing hybrid text, where AI output is paraphrased or edited (Deep et al., 2025; Hadra et al., 2026). This reduces how confidently undisclosed GenAI use can be identified and sanctioned (Perkins et al., 2024; Weber-Wulff et al., 2023). AI detection tools should therefore be understood as screening tools rather than proof of misconduct, meaning that sanction certainty still depends on instructor judgment and on the ability to establish whether and how GenAI was used in a submission (Ardito, 2025). Consequently, students may perceive sanction certainty as uneven and contestable. Here, suspicion may trigger investigation, but instructors often cannot conclusively determine whether, how much, or in what way GenAI was used in a submission.

These limitations in detection also help explain why certainty does not always lead to compliance. Adnin et al. (2025) highlight an adaptation mechanism where students strategically hide AI-mediated work to avoid detection. Using techniques such as paraphrasing, restructuring, or modifying outputs to avoid being flagged (Adnin et al., 2025). This suggests that deterrence pressure can shift behavior from complying to strategically concealing GenAI use rather than reducing use itself.

Severity

Severity captures how costly students believe discovery will be (Zhang et al., 2025). Across higher education, undisclosed GenAI use is commonly framed not as a minor procedural error but as a form of academic misconduct or plagiarism (An et al., 2025; Luo, 2024). In practice, this involves both formal and informal sanctions for students (Zhang et al., 2025). Formal sanctions commonly range from mark deductions to failure of the course, and in serious or repeated cases may result in suspension or even expulsion (Nguyen & Goto, 2024; Zhang et al., 2025). Informal sanctions can include social disapproval, negative judgements, stigma, shame and guilt (Gonsalves, 2025; Zhang et al., 2025).

Evidence often supports the deterrence prediction that harsher consequences reduce violations. Zhang et al. (2025) show a strong negative association between punishment severity and the intention to engage in AI-based plagiarism. Goyal et al. (2026) also find that severity of formal sanctions reduces students' intention to misuse GenAI for scholarly work, offering additional support that severity can deter GenAI misuse. In the context of content entrepreneurship, Lou et al. (2025) further report that perceived sanction severity strengthens GenAI policy compliance intention, showing that severity operates as a deterrence cue across governance domains.

Yet severity can also create incentives to conceal. Nguyen and Goto (2024) show that strict sanctions and stigma are linked to concealment and under-reporting. This indicates that “harsher penalties” can push behavior to conceal rather than increase honest reporting (Nguyen & Goto, 2024). Similarly, Qu et al. (2025) suggest that restrictive guidelines can reduce both compliance and truthfulness, suggesting that overly rigid or poorly framed policies may unintentionally backfire. Zhang et al. (2025) further complicate the picture by showing that informal sanctions (such as shame and guilt) can be more influential than formal sanctions, suggesting that “severity” may not be a dominant predictor in all settings.

Swiftness

Swiftness refers to the extent to which students expect consequences to follow quickly once undisclosed GenAI use is suspected or detected. Compared to certainty and severity, swiftness is less frequently measured in higher education. Nevertheless, swiftness has clear precedent in GenAI governance research. Lou et al. (2025) explicitly measure sanction swiftness and show that higher perceived swiftness strengthens GenAI policy compliance intention in the context of social content entrepreneurship.

Evidence from related academic integrity research further supports the relevance of swiftness. Dench and Joyce (2022) demonstrate that credible and immediate enforcement signals substantially reduce subsequent cheating, indicating that timely responses can shape student behavior. Nagin and Pogarsky (2004) provide theoretical reasoning for why timing matters, showing that delayed consequences carry less weight because present-oriented individuals discount future costs or overlook them altogether.

Taken together, these findings justify modeling swiftness of sanctions as a meaningful deterrence perception in higher educational contexts, where slow or ambiguous procedures may weaken deterrence perceptions.

In conclusion, Deterrence Theory predicts that stronger perceptions of detection risk and sanctions reduce non-disclosure. The evidence reviewed above suggests these same perceptions may also encourage concealment, a tension revisited in the discussion. Following the classic prediction, the hypotheses are stated as negative relationships.

H3: Perceived certainty of detection is negatively related to students' non-disclosure.

H4: Perceived severity of sanctions is negatively related to students' non-disclosure.

H5: Perceived swiftness of sanctions is negatively related to students' non-disclosure.

2.2.3. Theory of Planned Behavior and Non-Disclosure

The Theory of Planned Behavior (TPB) explains behavior as the outcome of intentions, which are formed through three main factors: attitudes, subjective norms, and perceived behavioral control (PBC) (Ajzen, 1991). TPB has been widely used to explain academic integrity behaviors such as cheating and plagiarism (Gonsalves, 2025; Huang et al., 2025). Recent GenAI research suggests that disclosure decisions need more attention, because disclosure is not only about compliance but also has social and psychological meaning (Gonsalves, 2025; Yusuf et al., 2025).

Attitudes

In this study, the relevant attitude concerns non-disclosure itself, that is, whether concealing GenAI use is seen as acceptable and low-risk or as wrong and hazardous. Because disclosing and concealing are two sides of one decision, a negative view of disclosure implies a more favorable view of non-disclosure (Adnin et al., 2025; Gonsalves, 2025). A dominant pattern in research suggests that attitudes toward disclosure are often negative as disclosure is not viewed as a neutral “transparency step,” but as an action that can create personal consequences for the disclosing student (Gonsalves, 2025; Luo, 2025). Students describe disclosure as potentially self-incriminating, particularly when rules are ambiguous and enforcement is inconsistent, which makes concealment appear safer than openness (Adnin et al., 2025; Gonsalves, 2025).

Importantly, this negative stance is not purely speculative. Causal evidence indicates that disclosure can carry real social costs for the disclosing agent. Schilke and Reimann (2025) show across experiments that AI disclosure reduces trust, largely because disclosure lowers perceived legitimacy. This implies that transparency can produce a reputational penalty even when it is socially desirable (Tan et al., 2024). In a higher education assessment context, Luo (2025) similarly highlights that opaque evaluation practices reduce students' trust in the responsible use of disclosure. Together, these findings support the idea that negative attitudes toward disclosure are grounded in evidenced expectations about trust, legitimacy, and

evaluative consequences, which in turn can foster non-disclosure (Luo, 2025; Schilke & Reimann, 2025; Tan et al., 2024).

At the same time, the literature also suggests a counter-mechanism in which negative internal feelings can lead to more disclosure rather than less. Yusuf et al. (2025) show that cognitive dissonance is positively associated with disclosure behavior. This indicates that when students experience internal conflict about concealing GenAI-use, disclosure can function as a way to restore self-consistency and reduce discomfort (Yusuf et al., 2025). This implies that attitudes toward disclosure are not shaped purely by perceived external costs, but also by self-evaluations. This might motivate disclosure even when reputational risks are present (Schilke & Reimann, 2025; Yusuf et al., 2025). Barrett and Pack (2023) further show that students generally regard the undisclosed use of GenAI as inappropriate across a range of writing-related tasks. Even though students were somewhat more accepting than teachers in some cases, undisclosed use consistently received the lowest appropriateness ratings among students themselves (Barrett & Pack, 2023). This suggests that, even among the actors performing the behavior, non-disclosure is often evaluated negatively.

Subjective norms

Subjective norms capture the perceived social expectations surrounding disclosure (Ajzen, 1991; Gonsalves, 2025). Here the relevant norm concerns non-disclosure, meaning the extent to which students perceive concealment as normal, expected, and socially safe. In the context of GenAI disclosure, dominant evidence suggests that norms can discourage disclosure by making non-disclosure appear normal and socially safer (Gonsalves, 2025). Students report that GenAI use is widespread and that non-declaration is often perceived as routine, which can normalize concealment and reduce perceived social pressure to disclose (Gonsalves, 2025; Nguyen & Goto, 2024; Yusuf et al., 2025). Similarly, students may fear stigmatization or negative judgment if they disclose, which shifts perceived norms away from openness and toward selective or full non-disclosure (Adnin et al., 2025; Luo, 2024; Nguyen & Goto, 2024).

Contrarily, evidence indicates that normative acceptance can reduce the social costs of disclosure and thereby weaken incentives for concealment (Schilke & Reimann, 2025). For instance, experimental evidence shows that the trust penalty of AI disclosure is smaller when AI use is perceived as more collectively accepted (Schilke & Reimann, 2025). Qualitative findings further suggest that in some settings disclosure becomes more acceptable when the supportive signals dominate. When instructors are transparent and AI-literate, students report higher comfort disclosing AI use, shifting perceived norms toward openness (Adnin et al., 2025; Luo, 2025). Taken together, the literature suggests that subjective norms may either promote non-disclosure through normalization and stigma, or support disclosure when social approval and institutional communication reduce reputational risk (Gonsalves, 2025; Luo, 2025).

Perceived behavioral control

Perceived behavioral control (PBC) reflects students' perceived ability to perform a behavior, shaped by internal and external factors (Ajzen, 1991). In this study, PBC refers to perceived control over non-disclosure, that is, how easily students feel they can withhold or limit information about GenAI use if they choose to (Adnin et al., 2025; Gonsalves, 2025). Much of the evidence suggests that unclear guidance raises this sense of control, because when rules are vague students feel both less able to disclose correctly and more able to conceal without consequence (Adnin et al., 2025; Gonsalves, 2025). Students report that vague or inconsistent instructions reduce perceived control and make non-disclosure an attractive strategy (Adnin et al., 2025; Gonsalves, 2025; Moorhouse et al., 2023). At the same time, students often perceive high control over concealment. By adopting strategies to hide GenAI use and managing perceived detection risk, perceived feasibility of non-disclosure is increased (Adnin et al., 2025).

In conclusion, the Theory of Planned Behavior suggests that students' intention to withhold GenAI use is shaped by their attitudes toward disclosure, perceived subjective norms, and perceived behavioral control (Ajzen, 1991). Therefore, the following hypotheses are formulated as follows:

H6: Attitude toward GenAI non-disclosure is positively related to students' non-disclosure.

H7: Subjective norms toward non-disclosure are positively related to students' non-disclosure.

H8: Perceived behavioral control over non-disclosure is positively related to students' non-disclosure.

2.2.4. Integrating Theory

Deterrence Theory and the Theory of Planned Behavior offer complementary explanations for students' GenAI non-disclosure. Deterrence Theory focuses on how students respond to governance pressure, conceptualized as perceived certainty of sanctions, severity of sanctions, and swiftness of consequences (Lou et al., 2025). The Theory of Planned Behavior explains non-disclosure as an intention-driven behavior shaped by attitudes, subjective norms, and perceived behavioral control (Ajzen, 1991). Studying each theory in isolation risks missing key parts of the disclosure decision process (Gonsalves, 2025; Nguyen-Phuoc et al., 2024; Yusuf et al., 2025).

Deterrence Theory alone can model rule-governance cues, but it does not fully capture how students interpret those cues and translate them into intentions. Evidence suggests that students not only "calculate" sanctions but also attach meaning to disclosure (Adnin et al., 2025; Gonsalves, 2025; Zhang et al., 2025). The Theory of Planned Behavior, by contrast, is strong on intention formation but gives too little weight to the external enforcement landscape that defines what is at stake (Gonsalves, 2025). In GenAI assessment contexts, policies, detection systems, and disciplinary procedures are background conditions that shape students' rationale when deciding to disclose (Luo, 2024; Moorhouse et al., 2023).

Recent literature explicitly supports the need for this integration. Gonsalves (2025) argues that TPB should be expanded or combined with other approaches to capture the complexity of GenAI disclosure compliance. Yusuf et al. (2025) similarly stress that disclosure behavior cannot be explained by psychological mechanisms alone. This calls for the need for greater attention to external influences such as institutional conditions and social pressures. Outside the higher education context, Nguyen-Phuoc et al. (2024) propose that deterrence models should be expanded by integrating factors related to social norms and individual perceptions. They thereby argue that attitudes should be positioned as a key channel through which perceptions of sanctions can influence the intentions to comply (Nguyen-Phuoc et al., 2024).

This raises the possibility that deterrence perceptions shape non-disclosure not only directly but also indirectly, by first shaping the psycho-social evaluations in the Theory of Planned Behavior. There is no strong theoretical basis for assuming a single channel, so more than one is considered. The most direct route runs through attitude. Deterrence Theory describes a weighing of the expected costs and benefits of an act, which maps closely onto the evaluative component of the Theory of Planned Behavior (Ajzen, 1991). A second plausible route runs through subjective norms. Sanctions not only impose costs but also signal what an institution considers unacceptable, which may influence what students believe their peers find acceptable. Both routes concern how the behavior is evaluated rather than whether it can be performed, which places them close together conceptually.

Perceived behavioral control is not treated as a channel. The items focus on students' perceived autonomy in reporting. They measure how much control students feel they have over what they disclose, not their perceptions of detection or enforcement. Sanction perceptions have no clear route to that sense of autonomy, because sanctions govern what follows an act rather than whether the act is within the student's control. Perceived behavioral control therefore remains a direct predictor of non-disclosure (H8), while its role as a channel is set aside.

Because these indirect routes are plausible but not firmly established, and because the design is cross-sectional, they are not stated as formal hypotheses. They are instead examined as an exploratory extension of the model. As all constructs were measured at a single point in time, this analysis can show only that the observed associations are consistent with an indirect pathway, not that such a pathway operates over time (Bryman, 2016).

2.2.5. Individual-Level Predictors of Non-Disclosure

The integrated model of Deterrence Theory and the Theory of Planned Behavior explains how governance signals and socio-cognitive evaluations shape students' non-disclosure decisions (Nguyen-Phuoc et al., 2024). However, this model may not fully capture why some students feel more compelled than others to conceal their GenAI use. To address this, the

present study extends the model with two additional individual-level factors: academic self-efficacy and policy awareness. These constructs complement the model of Deterrence theory and the Theory of Planned Behavior by accounting for differences in students' perceived capability and perceived understanding of institutional rules.

Academic self-efficacy

The first construct is academic self-efficacy, defined here as students' confidence in their ability to complete academic work without GenAI support (Pajares, 1996). This nuance is important. The construct does not refer to confidence in disclosing AI use, but to confidence in performing academic tasks independently (Pajares, 1996). Academic self-efficacy is task-specific: it reflects whether students believe they can complete the assessed work without relying on external AI assistance (Pajares, 1996).

This makes self-efficacy theoretically relevant for the present study. The integrated model fails to represent students' feelings regarding their capability to complete tasks independently. Assessments are designed to show whether students have mastered certain skills. Students who cannot perform those skills without GenAI are more likely to use it. When this use is not disclosed, it becomes difficult to determine whether the skills have truly been mastered. This perception of capability also helps explain why disclosure decisions differ across students. Especially as some see GenAI use as optional while others experience it as a necessity (Adnin et al., 2025). Students who believe they can complete assignments independently may have less reason to rely on GenAI and therefore less reason to conceal its use. In contrast, students with lower academic self-efficacy often turn to GenAI as a way to cope when tasks feel difficult or stressful (Zhang et al., 2024). For these students, disclosure can feel more uncomfortable because it draws attention to their dependence on AI, which they may see as a sign that they cannot complete the work on their own (Adnin et al., 2025; Gonsalves, 2025). Academic self-efficacy therefore adds an individual-level explanation for why students facing similar rules may still differ in their intention to conceal GenAI use.

Literature also provides support for the relevance of academic self-efficacy. Using student reflections, Adnin et al. (2025) show that task-specific confidence varies considerably amongst students. Some students report the feeling of being unable to complete assignments without GenAI support, whereas others describe being fully capable of performing the same academic tasks independently (Adnin et al., 2025). In a meta-analysis, Krou et al. (2021) found that self-efficacy is negatively associated with academic dishonesty. This suggests that students who feel more capable of succeeding academically are generally less likely to engage in academic misconduct. This aligns with Leaton Gray et al. (2025), who argue that efficacy beliefs can help students resist unethical shortcuts in AI-enabled academic settings. Taken together, these studies support the broader expectation that stronger efficacy beliefs reduce students' likelihood of using dishonest or self-serving strategies when facing academic pressure.

Based on this reasoning, the following hypothesis is proposed:

H9: Academic self-efficacy to complete assessed academic tasks without GenAI assistance is negatively related to students' non-disclosure.

Academic self-efficacy may also reach non-disclosure indirectly, by shaping how students evaluate the behavior rather than by acting on it alone. Efficacy beliefs influence not only what students do but also how they evaluate the options in front of them (Pajares, 1996). A student who feels able to complete assessed work without GenAI has little need to rely on it and little to conceal, which may lead them to view non-disclosure less favorably. Students who use GenAI to manage their academic workload may be more likely to view concealment as acceptable (Adnin et al., 2025; Gonsalves, 2025). This positions attitude as the most plausible channel through which self-efficacy operates (Ajzen, 1991). As with the deterrence perceptions and policy awareness, this indirect route cannot be tested conclusively with cross-sectional data. It is therefore not stated as a formal hypothesis, but is examined in the exploratory analysis of indirect effects.

Policy awareness

The second additional construct is policy awareness. Policy awareness refers to the extent to which students believe they understand institutional rules about AI use and disclosure (Samonte & Acedera, 2025). It captures their subjective assessment of how familiar and confident they feel with those rules, rather than their objective knowledge of the policy itself. This is important, as clear, accessible, and well-communicated policies are considered an important driver of student behavior (McCabe et al., 2001). In the context of generative AI, this is especially relevant because policies are still developing and may differ across institutions, courses, and assignments (Adnin et al., 2025; Moorhouse et al., 2023). Therefore, policy awareness captures students' self-reported understanding of when and how AI use should be disclosed.

The term awareness follows the compliance literature, in which policy awareness is the established label for an individual's grasp of the rules governing their behavior (Bulgurcu et al., 2010; D'Arcy et al., 2009; Lund et al., 2025). It should be noted, however, that the construct is measured as a self-assessment rather than as an objective test of policy knowledge, which is consistent with the scope of this study. This distinction matters, since students may feel confident about rules they have in fact misunderstood.

Awareness of institutional rules has been identified as an important antecedent of compliance behavior in several domains. Understanding the rules allows individuals to interpret them correctly and to judge which actions are acceptable and which may constitute a violation (Ishwardat et al., 2024). This mechanism is central to the information security policy literature, where awareness of organizational rules is treated as a key determinant of compliance. Bulgurcu et al. (2010) show that security policy awareness shapes employees' beliefs about the consequences of compliance and, through these beliefs, their attitude and intention to comply. Comparable patterns appear in deterrence-based research, where familiarity with rules and sanctions is associated with lower levels of policy violation (D'Arcy

et al., 2009). In educational contexts, students who report a clearer understanding of academic integrity rules are generally better able to recognize which practices are prohibited (Gullifer & Tyson, 2014). Taken together, this suggests that students with higher policy awareness may be less inclined to leave their use of AI undisclosed.

However, having knowledge of rules does not automatically lead to compliant behavior. Lund et al. (2025) found that policy awareness did not significantly predict students' perceptions of the ethicality of AI use or their AI-related behavior, suggesting that knowledge alone may be insufficient to influence decision-making. This aligns with broader findings from the Theory of Planned Behavior, which emphasize that behavior is driven by beliefs and attitudes rather than knowledge alone (Ajzen, 1991). In this perspective, knowledge functions as an informational basis that shapes beliefs about consequences, which in turn influence attitudes and behavioral intentions.

Further, the effectiveness of policy awareness may be limited because students often face an ambiguous and inconsistent policy landscape (An et al., 2025). As GenAI guidelines are still evolving and often differ significantly between institutions and courses, it may be difficult for students to form a clear and accurate understanding of disclosure expectations (Adnin et al., 2025; Moorhouse et al., 2023). Therefore, when formal guidance is unclear, students may rely more on personal interpretations and social norms when deciding whether to disclose AI use (Moorhouse et al., 2023).

From a theoretical perspective, policy awareness can be linked to both deterrence theory and the Theory of Planned Behavior. Deterrence theory suggests that individuals are less likely to violate rules when they perceive a high likelihood of detection and punishment (Gibbs, 1979). These perceptions depend partly on whether individuals understand which behaviors are prohibited and what consequences may follow (Gibbs, 1979). If students are unsure when AI use must be disclosed, they may be unable to accurately assess the risks involved.

Taken together, this reasoning suggests that students who better understand the rules are more able to judge what must be disclosed and what risks non-disclosure carries, and are therefore less likely to leave their GenAI use undisclosed.

Based on this reasoning, the following hypothesis is proposed:

H10: Policy awareness is negatively related to students' non-disclosure.

Policy awareness may also operate indirectly, by shaping the evaluations described in the Theory of Planned Behavior rather than acting on non-disclosure on its own. Because this indirect route cannot be tested conclusively with cross-sectional data, it is not stated as a formal hypothesis but is included in the exploratory analysis of indirect effects described earlier.

2.2.6. Conceptual Model

The literature review and the hypotheses developed above are brought together in the conceptual model shown in Figure 2. The model explains students' non-disclosure of GenAI use through three groups of predictors. The first comes from Deterrence Theory and captures the enforcement environment: perceived certainty of detection, severity of sanctions, and swiftness of consequences (H3 to H5). The second comes from the Theory of Planned Behavior and captures how students evaluate the behavior: attitude toward non-disclosure, subjective norms, and perceived behavioral control (H6 to H8). The third consists of the two individual-level factors, academic self-efficacy and policy awareness (H9 and H10). All eight predictors are modelled as direct effects on non-disclosure.

On the outcome side, non-disclosure is not treated as a single act. In line with the conceptualization developed in Section 2.2.1, it is distinguished along two dimensions: whether concealment is full or partial, and whether it concerns past behavior or intention. This yields four outcomes: past full non-disclosure, past partial non-disclosure, full non-disclosure intention, and partial non-disclosure intention. The relative prevalence of these forms is examined first (H1). Each predictor is then tested against all four outcomes, which makes it possible to assess whether the predictors differ across the forms of non-disclosure (H2).

Beyond these direct effects, the model allows for the possibility that the deterrence perceptions, policy awareness, and academic self-efficacy reach non-disclosure indirectly, by first shaping the evaluations in the Theory of Planned Behavior. Because the cross-sectional design cannot establish such a pathway, this possibility is not part of the main model. It is examined separately as an exploratory analysis, reported in Chapter 4.




Figure 2 – Enhanced Conceptual Model

3. Methodology

This chapter describes how the study was designed and carried out. It first presents the research design (Section 3.1) and the sample, including the population, the sampling strategy, and the respondents (Section 3.2). It then sets out the measurement instrument and how the measures were validated (Section 3.3), followed by the operational conceptual model (Section 3.4) and the data analysis strategy used to prepare the data and test that model (Section 3.5).

3.1. Research Design

This study employed a quantitative, cross-sectional survey design to investigate students' non-disclosure of GenAI use in assessed academic work. The approach follows a deductive logic, in which hypotheses were derived from Deterrence Theory and the Theory of Planned Behavior and then tested empirically against survey data. A survey suited this aim well, because the central constructs are perceptual and latent, such as perceived certainty of detection, perceived severity of consequences, attitudes, subjective norms, and perceived behavioral control (Boateng et al., 2018). These beliefs cannot be observed directly and are therefore best captured through self-report measures (Boateng et al., 2018).

A cross-sectional design was chosen for both practical and substantive reasons. It measures all variables at a single point in time, which fits the limited timeframe of the project and allows a broad overview across different institutions, fields of study, and study levels (Capili, 2021). This breadth is valuable given how little empirical work has examined student non-disclosure of GenAI use, particularly in the Dutch higher education context (Bos et al., 2025; den Hollander et al., 2025; Gonsalves, 2025). The design does carry a clear limitation, as data gathered at one moment cannot establish temporal order or support strong causal claims (Bryman, 2016). Nonetheless it offers an appropriate first mapping of the relationships proposed in the conceptual model.

The topic is sensitive, since students may associate non-disclosure with academic misconduct (Gonsalves, 2025; Nguyen & Goto, 2024). An anonymous, self-administered questionnaire was therefore preferred over interviewer-based methods, as respondents tend to answer more openly when no interviewer is present (Gnambs & Kaspar, 2015). Self-administration also reduces social desirability bias and underreporting, which improves the accuracy of responses on sensitive behavior (Gnambs & Kaspar, 2015; Tourangeau & Yan, 2007). Data were collected through an online questionnaire administered in Qualtrics.

Throughout the research process, ethical considerations were taken into account. The study received formal approval from the Human Research Ethics Committee of Delft University of Technology on 8 May 2026. Participation was voluntary, and respondents were informed about the purpose of the study before taking part. Informed consent was obtained before the questionnaire began. Responses were collected anonymously and treated confidentially, so that individual students could not be identified from the data.

3.2. Sample

The target population consisted of higher education students completing assessed academic work in settings where GenAI use may be regulated through disclosure requirements. The study was conducted in Dutch higher education. In this setting, generative AI has become part of everyday academic practice, while institutional approaches to governance and support are still taking shape (Bos et al., 2025; den Hollander et al., 2025). Data were collected between June 3 and August 6, 2026 through an online questionnaire administered using Qualtrics.

Recruitment relied on non-probability sampling, combining convenience and snowball sampling. This choice suited the limited timeframe of the project and helped maximize reach and sample size, although it reduces the representativeness of the wider student population (Golzar et al., 2022; Ting et al., 2025). Convenience recruitment took place through university channels and student networks, which gave direct access to the relevant population. The questionnaire was distributed across several social media platforms, including LinkedIn, WhatsApp, Instagram, and Facebook. Recruitment also targeted students from multiple universities to avoid limiting the sample to a single institution. Snowball sampling complemented this approach. The final page of the questionnaire deliberately invited respondents to share the survey with peers who met the inclusion criteria, a technique that can raise participation when a topic is sensitive or when parts of the population are harder to reach (Ting et al., 2025).

The required sample size was established through an a priori power analysis in G*Power 3.1 (Faul et al., 2007). Assuming a medium effect size ($f^2 = .15$; Cohen, 1988), an alpha of .05, a power of .95, and eight predictors, the analysis indicated a minimum of 160 respondents. The full calculation is reported in Appendix 4.

The questionnaire opened with an informed consent statement, after which respondents first answered detailed questions about their GenAI use and only then the questions on disclosure. This ordering was intentional. GenAI use is not a single act, since students rely on these tools for a range of different tasks (Freeman, 2025). Measuring these specific uses first helped to break a broad and sensitive topic into concrete components, which improves the clarity and consistency of the later disclosure measures (Peytcheva & Yan, 2025).

The questionnaire was offered in two languages, English and Dutch, and respondents selected their preferred version at the start of the survey. It contained 47 items measuring twelve constructs: past full non-disclosure (PFD), past partial non-disclosure (PPD), full non-disclosure intention (FDI), partial non-disclosure intention (PDI), certainty (CER), severity (SER), swiftness (SOS), attitude (ATT), subjective norms (SUN), perceived behavioral control (PBC), policy awareness (PA), and academic self-efficacy (ASE). Five demographic questions were included as well. The full set of items is provided in Appendix 1 for the English version and in Appendix 3 for the Dutch version. All responses were stored securely and were accessible only to the research team.

The survey attracted 728 responses in total. Of these, 139 were incomplete and were discarded before the data were imported into SPSS, which left 589 complete responses. The

dataset was then cleaned in six sequential steps, and 431 responses were retained for analysis, corresponding to a retention rate of 73%.

The largest single reduction came from the response-time criterion, which removed 116 respondents who completed the questionnaire faster than the minimum threshold (198 seconds for the English version and 160 seconds for the Dutch version) or slower than the maximum threshold of one hour (Brysbaert et al., 2021; Rayner et al., 2016). The remaining steps each removed only a small number of cases. 23 respondents selected "prefer not to say" on one or more past-behavior items, which made their responses unusable for those scales, and one respondent showed zero intra-individual response variability. Four respondents gave logically impossible answers, either an implausible age or a contradictory pattern on the GenAI use question, and eight were removed as outliers after individual inspection. A final six responses were discarded because the open-ended answers on field or level of study were illegible or non-compliant. The full criteria, the underlying reasoning, and the cleaning logbook are documented in Appendix 5.

Despite these exclusions, the final sample of 431 respondents comfortably exceeds the minimum of 160 respondents indicated by the a priori power analysis (Appendix 4).

3.2.1. Demographic Information

This subsection describes the demographic composition of the final sample ($n = 431$). Table 1 (p. 25) reports the demographic characteristics, Table 2 (p. 25) summarizes the GenAI use cases that respondents reported, and Table 3 (p. 26) shows the number of use cases selected per respondent.

Respondents had a mean age of 24.16 years ($SD = 4.59$). The sample contained somewhat more women (60.3%) than men (36.7%), with the remainder identifying as other (0.9%) or preferring not to say (2.1%). In terms of field of study, the largest groups were Business, Finance, and Law (36.9%) and STEM (30.2%), followed by Social Sciences and Education (18.3%), while Arts and Humanities (7.9%) and Medicine and Health (6.7%) were smaller. The sample was concentrated at the master's level (55.9%) and among third-year bachelor students (23.2%), with smaller shares of first-year (5.8%) and second-year (7.4%) bachelor students, PhD candidates (3.0%), and other students (4.6%). Regarding survey language, 38.7% of respondents chose to complete the questionnaire in Dutch, while 61.3% selected the English version.

The gender balance of the sample is broadly in line with the national picture, though it leans slightly more female. In Dutch higher education, women make up a small majority of enrolments, at around 53 to 54 percent (CBS, 2023). At 60.3%, women are therefore modestly over-represented in this sample. The distribution across study levels is less representative, as master's students form a clear majority, whereas bachelor students are more numerous in the wider student population (CBS, 2026; OECD, 2025). This skew most likely reflects the recruitment channels used and should be kept in mind when interpreting the results (Golzar et al., 2022; Ting et al., 2025).

Table 1 – Descriptive statistics regarding demographic characteristics of the sample

Variable	Freq (%)	$\bar{x}$	S
1. Age		24.16	4.59
2. Gender			
Male	36.7		
Female	60.3		
Other	.9		
Prefer not to say	2.1		
3. Field of study			
Arts & Humanities	7.9		
Business, Finance, & Law	36.9		
Medicine & Health	6.7		
STEM	30.2		
Social Sciences & Education	18.3		
4. Level of study			
First-year bachelor student	5.8		
Second-year bachelor student	7.4		
Third-year bachelor student	23.2		
Master's student	55.9		
PhD	3.0		
Other	4.6		
5. Selected language			
English	61.3		
Dutch	38.7		

Note. $n = 431$, Freq (%) = frequency in percentage of the total respondents, $\bar{x}$ = mean, S = standard deviation.

Respondents were also asked which types of GenAI use they had engaged in over the past academic year. Because multiple answers were allowed, the responses are summarized as a multiple-response frequency table (Table 2). The most frequently reported uses were explaining concepts (79.8% of respondents), suggesting ideas for research (74.5%), and checking grammar, spelling, or flow (71.2%). More direct forms of use were less common. Including AI-generated text after editing it with a digital tool was reported by 20.2% of respondents, and including such text directly in assessed work by 18.3%. Almost all respondents reported at least one form of use, as only 1.4% selected "none of the above."

Table 2 – Reported GenAI use cases (multiple responses allowed)

AI use case	Responses		Percent of total
	N	Percent	
Used AI to suggest ideas for research	321	13.5%	74.5%
Used AI to explain concepts to me	344	14.4%	79.8%
Used AI to structure my thoughts	253	10.6%	58.7%
Used GenAI to search for sources, literature, or references	291	12.2%	67.5%

Note. The table is continued on the next page.

Table 2 Continued

AI use case	Responses		Percent of total
	<i>N</i>	<i>Percent</i>	
Used AI to summarize a relevant article	297	12.5%	68.9%
Used AI to generate visual aids, audio or other media	146	6.1%	33.9%
Used AI to check the grammar, spelling, or flow of my own text	307	12.9%	71.2%
Included AI-generated text in assessment after editing it myself	220	9.2%	51.0%
Included AI-generated text in assessment after editing it with a digital tool	87	3.7%	20.2%
Included AI-generated text directly in a piece of assessed work	79	3.3%	18.3%
Other	32	1.3%	7.4%
None of the above	6	0.3%	1.4%
<i>Total</i>	<i>2383</i>	<i>100.0%</i>	<i>552.9%</i>

Note. $n = 431$. Multiple responses were allowed, so percentages of cases sum to more than 100%. N = number of times the option was selected; % of responses = share of all 2383 selections; % of cases = share of the 431 respondents who selected the option.

Finally, the number of distinct use cases per respondent was examined to indicate the breadth of GenAI use (Table 3). Most respondents reported using GenAI for several purposes rather than just one. The distribution peaks at five and six use cases (16.9% each), and relatively few respondents reported either a single use case (4.9%) or a very high number.

Table 3 – Distribution of the number of reported GenAI use cases per respondent

Number of cases selected	<i>N</i>	<i>%</i>
1	21	4.9
2	30	7.0
3	37	8.6
4	53	12.3
5	73	16.9
6	73	16.9
7	52	12.1
8	43	10.0
9	24	5.6
10	23	5.3
11	2	0.5
<i>Total</i>	<i>431</i>	<i>100.0</i>

Note. $n = 431$. The maximum possible number of use cases is 11. Contradictory combinations, such as selecting "none of the above" together with one or more specific use cases, were removed during data cleaning (see Appendix 5).

3.3. Measurement and Survey

This section describes how the constructs in the conceptual model were measured. It first explains how the survey items were developed and operationalized (Section 3.3.1), then reports the exploratory factor analysis used to examine the underlying structure of the items (Section 3.3.2), and finally presents the reliability analysis of the resulting scales (Section 3.3.3). The full instrument is provided in Appendix 1 for the English version and Appendix 3 for the Dutch version.

3.3.1. Measures

The questionnaire measured twelve constructs: four non-disclosure outcomes (past full non-disclosure, past partial non-disclosure, full non-disclosure intention, and partial non-disclosure intention), three deterrence perceptions (certainty, severity, and swiftness), three Theory of Planned Behavior constructs (attitude, subjective norms, and perceived behavioral control), and two individual-level predictors (policy awareness and academic self-efficacy). Wherever possible, items were adapted from established, peer-reviewed instruments and reworded to fit the context of GenAI use and non-disclosure in assessed academic work. Reusing validated measures supports content validity, while adapting their wording is appropriate when an instrument is transferred to a new setting (Boateng et al., 2018).

The four non-disclosure constructs were an exception, as they were developed specifically for this study. Existing GenAI disclosure measures do not distinguish between full and partial non-disclosure, so parallel item sets were written for each form. The intention items were phrased as behavior-specific, self-directed, and future-oriented statements, in line with Theory of Planned Behavior guidance (Ajzen, 2006), whereas the past-behavior items followed non-disclosure research suggesting that withholding is best captured through specific types of omission rather than a single general question (Yusuf et al., 2025). To support accurate recall, the behavior items were also bounded to a defined timeframe (Grigoryeva & Robbins, 2024).

All constructs were measured on five-point scales, with anchors adapted to the nature of each construct. Past-behavior items used a frequency scale (1 = never, 5 = always), intention items used a likelihood scale (1 = very unlikely, 5 = very likely), the deterrence, subjective norms, perceived behavioral control, policy awareness, and academic self-efficacy items used an agreement scale (1 = strongly disagree, 5 = strongly agree), and attitude was measured with a five-point semantic differential. Applying a consistent five-point format across the questionnaire keeps the response process simple and reduces respondent burden in a relatively long instrument (Kusmaryono et al., 2022). A full justification of each construct, including the reasoning behind every adaptation, is provided in Appendix 2.

Before the survey was distributed, the instrument was refined through a two-stage pilot test. In the first stage, the draft questionnaire was reviewed by two experts. This review focused on the conceptualization of the constructs and the wording of the items, and it served mainly to strengthen content validity and the overall design quality of the survey (Boateng et al., 2018; Elangovan & Sundaravel, 2021). In the second stage, cognitive interviews were conducted with eight participants across both the Dutch and English versions of the

questionnaire. By observing respondents as they completed the survey, the researcher could assess how individual items were interpreted and identify points of hesitation or ambiguity. Such pretesting is widely applied, as respondent feedback helps surface issues of wording, comprehension, and clarity prior to full data collection (Hashim et al., 2022). Minor issues relating to wording and spelling that surfaced during the sessions were resolved before launch. Overall, the pilot test confirmed that the items were understood as intended and that the instrument was clear and ready for the main data collection.

3.3.2. Exploratory Factor Analysis

An exploratory factor analysis (EFA) was conducted to examine the underlying structure of the survey items and to assess the construct validity of the measures. It served three related purposes: to identify the latent constructs underlying the items, to examine whether the items grouped into the expected structure, and to evaluate whether each construct was measured by a coherent set of items (Fabrigar et al., 1999; Worthington & Whittaker, 2006). In this way, the EFA provides evidence of construct validity, since it shows whether items intended to capture the same construct load together and remain distinct from those measuring other constructs (Worthington & Whittaker, 2006).

An exploratory rather than a confirmatory approach was appropriate here. Confirmatory factor analysis tests a factor structure that is already well established, whereas EFA is suited to situations in which the structure is not yet fixed, such as when instruments are newly developed or adapted to a new context (Boateng et al., 2018; Worthington & Whittaker, 2006). Given the adapted and newly developed items in this study, the factor structure could not be assumed in advance, which made EFA the more suitable choice.

The analysis was carried out as two separate models. This theory-driven decision reflected the fact that the constructs represented distinct conceptual domains. The first model combined the perception and intention items, because these belong to the same theoretical framework derived from the Theory of Planned Behavior and Deterrence Theory. Constructs within a shared framework are expected to be related, so moderate correlations among these items are theoretically anticipated rather than a sign of poor discrimination (Ajzen, 1991). The second model contained the past behavior items, which reflect a reported behavioral outcome rather than a psychological or motivational construct (Ajzen, 1991). Analyzing the two sets separately keeps conceptually distinct domains apart and allows the structure within each to be assessed on its own terms.

Both sets of results are reported transparently in the subsections that follow. Notably, the distinction between full and partial non-disclosure separated more clearly than that between intention and past behavior. The following subsections present each model in turn, describing the settings applied, the strength of the solution, and the items that were removed, with the supporting output reported in Appendix 7.

Model 1: Intentions and predictors

Before interpreting the factor solution, the suitability of the data for factor analysis was assessed. The Kaiser-Meyer-Olkin (KMO) measure of sampling adequacy was 0.866, which indicates a good level of common variance among the items and is well above the accepted minimum of 0.50 (Hair et al., 2014). Bartlett's test of sphericity was significant, $\chi^2(741) = 9906.45$, $p < .001$, which confirms that the correlation matrix differed from an identity matrix (Shrestha, 2021). Together, these results show that the data were suitable for factor analysis.

Principal component analysis was used as the extraction method. Rather than relying on the eigenvalue-greater-than-one criterion, which can over- or under-extract components, the number of components was fixed at ten to match the ten constructs specified in the theoretical model (Costello & Osborne, 2005; Fabrigar et al., 1999). Because these constructs were expected to correlate, an oblique rotation was appropriate, and Direct Oblimin rotation with Kaiser normalization was applied (Fabrigar et al., 1999). The rotation converged after 12 iterations.

The communalities of the retained items ranged from 0.489 to 0.838. With the single exception of CER_4, which fell marginally below the 0.50 guideline, these values are acceptable and indicate that each item shared sufficient variance with the extracted components (Shrestha, 2021). A minimum loading threshold of 0.30 was applied when assigning items to components. This threshold is appropriate for the sample size in this study, since for samples of this size, loadings of approximately 0.30 and above can be considered meaningful (Stevens, 2002). Where an item loaded on more than one component, it was assigned to the component on which it loaded most strongly, provided the difference between the primary and secondary loading exceeded 0.30.

In the initial solution, the severity and swiftness items did not form clean, separate components, and two items were removed to resolve this. SOS_1 loaded most strongly on the severity component rather than on the two remaining timing items, and SER_5, which reflects the social consequences of discovery rather than the formal severity of a sanction, was the severity indicator whose removal produced the clearest separation. Both removals are documented in full in Appendix 7. One further item, FDI_2, showed a secondary loading but loaded most strongly on its intended component, with an absolute difference clearly exceeding 0.30, so it was retained.

The ten retained components together explained 72.94% of the variance, which is well above the commonly recommended minimum for social science research (Hair et al., 2014). The full output is reported in Appendix 7.

The rotated solution produced a clean structure in which each component corresponded to one intended construct. Table 4 (p. 30) summarizes the components, their associated items, and the range of loadings. Several components loaded negatively, which is an artifact of the direction of the oblique rotation rather than a substantive result, so the loadings are reported here in absolute terms.

Table 4 – Component structure of Model 1

Component	Construct	Items	Loading range
1	Partial non-disclosure intention	PDI_1–PDI_6	.696 to .846
2	Severity of sanctions	SER_1–SER_4	.738 to .870
3	Policy awareness	PA_1–PA_4	.821 to .901
4	Academic self-efficacy	ASE_1–ASE_4	.645 to .896
5	Certainty of sanctions	CER_1–CER_4	.586 to .874
6	Subjective norms	SUN_1–SUN_4	.777 to .848
7	Perceived behavioral control	PBC_1–PBC_3	.695 to .839
8	Attitude	ATT_1–ATT_4	.761 to .811
9	Full non-disclosure intention	FDI_1–FDI_2	.668 to .720
10	Swiftiness of sanctions	SOS_2–SOS_3	.757 to .762

Note. Loadings are reported as absolute values; FDI = Full Non-Disclosure Intention, PDI = Partial Non-Disclosure Intention, CER = Certainty, SER = Severity, SOS = swiftiness, ATT = Attitude, SUN = Subjective Norms, PBC = Perceived Behavioral Control, PA = Policy awareness, ASE = Academic Self-Efficacy.

As Table 4 shows, the items grouped onto their intended constructs, and each construct emerged as a distinct component. This pattern is meaningful in two ways. Items measuring the same construct loaded together, which supports convergent validity, and items measuring different constructs separated into distinct components, which supports discriminant validity. The outcome constructs behaved as intended as well, since full and partial non-disclosure intention formed two separate components, 9 and 1 respectively. Policy awareness emerged as a distinct component, separate from attitude and perceived behavioral control, supporting its treatment as an independent construct.

Model 2: Past behavior

The suitability of the data was assessed in the same way as for the first model. The Kaiser-Meyer-Olkin measure of sampling adequacy was 0.887, a value considered acceptable for factor analysis (Hair et al., 2014). Bartlett's test of sphericity was significant, $\chi^2(28) = 1659.24, p < .001$, which confirms that the correlation matrix differed from an identity matrix (Shrestha, 2021). The data were therefore suitable for factor analysis.

The same extraction and rotation settings were applied as in the first model. Principal component analysis was used, and the number of components was fixed at two to match the two past behavior constructs (Costello & Osborne, 2005; Fabrigar et al., 1999). Direct Oblimin rotation with Kaiser normalization was applied, since the two constructs were expected to correlate, and the rotation converged after four iterations. The same 0.30 loading threshold and cross-loading criterion were used as before (Stevens, 2002).

The final solution retained five items across two components, after the removal of three past partial non-disclosure items, for which the full reasoning is documented in Appendix 7. The retained items showed communalities ranging from 0.610 to 0.820, which are all acceptable (Shrestha, 2021). Component 1 corresponded to past partial non-disclosure, with loadings between 0.755 and 0.901, and Component 2 corresponded to past full non-

disclosure, with loadings of -0.898 and -0.901. Together, the two components explained 73.78% of the variance, which is well above the commonly recommended minimum for social science research (Hair et al., 2014).

The two behavior constructs emerged as distinct components. Items measuring the same construct loaded together, and the two forms of past non-disclosure separated cleanly, which supports the convergent and discriminant validity of these measures. As in the first model, full and partial non-disclosure formed two separate components rather than collapsing into a single behavioral factor.

Taken together, the two factor analyses show that the final measurement model differs only modestly from the originally proposed set of items. In Model 1, SER_5 and SOS_1 were removed due to cross-loading, which reduced severity from five items to four and swiftness from three items to two. In Model 2, three items measuring past partial non-disclosure (PPD_2, PPD_3, PPD_4) were removed, leaving this construct with three items instead of six. Full non-disclosure intention and past full non-disclosure were not affected by these reductions, as both had already been designed as two-item scales. Together with swiftness, they are therefore the only constructs measured with two items in the final instrument.

3.3.3. Reliability Analysis

Following the exploratory factor analysis, a reliability analysis was conducted to assess the internal consistency of the retained scales. For constructs measured by three or more items, Cronbach's alpha was calculated. A value of .70 or higher was treated as the preferred level of internal consistency, while a value above .60 was regarded as acceptable (Nunnally, 1978). Cronbach's alpha is not a suitable measure for two-item scales, so for these constructs the Pearson correlation between the two items was reported instead (Eisinga et al., 2013). This applied to past full non-disclosure, full non-disclosure intention, and swiftness, each of which was retained as a two-item scale after the factor analysis. The reliability and item-total statistics for all scales are reported in Table 5 (p. 32).

The multi-item scales showed acceptable to strong internal consistency, with Cronbach's alpha ranging from .731 to .910. The highest value was found for partial non-disclosure intention ($\alpha = .910$), followed by policy awareness ($\alpha = .882$) and subjective norms ($\alpha = .872$). Perceived behavioral control produced the lowest value ($\alpha = .731$), which still exceeds the preferred .70 threshold, so every multi-item scale met or surpassed the preferred level of internal consistency (Nunnally, 1978). The corrected item-total correlations supported these results, as all items correlated with their scale above the .30 threshold, with values ranging from .421 to .831 (Field, 2018). Removing individual items would not have meaningfully improved the reliability of any scale, so all retained items were kept.

The three two-item scales also showed satisfactory internal coherence. In each case, the two items were positively and significantly correlated ($p < .001$), with values of $r = .786$ for swiftness, $r = .751$ for full non-disclosure intention, and $r = .651$ for past full non-disclosure. These relatively strong inter-item correlations provide evidence that the paired items captured a common underlying construct and support treating each pair as a combined two-item measure (Eisinga et al., 2013).

Taken together, the reliability analysis indicates that the measurement scales were internally consistent. All multi-item scales exceeded the preferred .70 threshold, and the two-item scales showed adequate inter-item correlations, which supports the reliability of the instrument used in the subsequent analyses.

Table 5 – Reliability and item-total statistics for the measurement scales

Factor Name	Items	Cronbach's α / Pearson's r	Total Correlation
Past Full Non-Disclosure (PFD)	PFD_1 - I did not disclose any of my GenAI use when I was expected to. PFD_2 - I deliberately left out any mention of having used GenAI in my work.	$r = .651^{***}$	–
Past Partial Non-Disclosure (PPD)	PPD_1 - I disclosed some of my GenAI use but not all of it. PPD_5 - I only reported minor uses of GenAI but left out more substantial contributions. PPD_6 - I described my GenAI use in general terms rather than in specific detail.	$\alpha = .753$	.541 .635 .572
Full Non-Disclosure Intention (FDI)	FDI_1 - I will keep all of my GenAI use undisclosed when disclosure is expected. FDI_2 - I will deliberately leave out any mention of having used GenAI in my work.	$r = .751^{***}$	–
Partial Non-Disclosure Intention (PDI)	PDI_1 - I will disclose some of my GenAI use but not all of it. PDI_2 - I will leave out details about how I used GenAI. PDI_3 - I will make my GenAI use seem less extensive than it actually was. PDI_4 - I will give limited information about how much my final work depended on GenAI output. PDI_5 - I will only report minor uses of GenAI but leave out more substantial contributions. PDI_6 - I will describe my GenAI use in general terms rather than in specific detail.	$\alpha = .910$	.684 .748 .831 .797 .756 .677
Certainty (CER)	CER_1 - It is likely that undisclosed GenAI use will be detected. CER_2 - My examiner would detect the extent of my GenAI use. CER_3 - AI detection tools would detect the extent of my GenAI use. CER_4 - An oral defense or follow-up interview would reveal the extent of my GenAI use.	$\alpha = .766$	.571 .677 .620 .421
Severity (SER)	SER_1 - If this was discovered, the formal consequences would be serious. SER_2 - If this was discovered, my academic progress would be hindered. SER_3 - If this was discovered, I would have to go through a difficult disciplinary process. SER_4 - If this was discovered, my reputation will be harmed.	$\alpha = .860$	.718 .794 .741 .586

Note. The table is continued on the next page.

Table 5 Continued

Factor Name	Items	Cronbach's α / Pearson's r	Total Correlation
Swiftiness (SOS)	SOS_2 - If this was discovered, a final decision would be reached in a very short time. SOS_3 - If this was discovered, it would not take long for me to receive a sanction.	$r = .786^{***}$	– –
Attitude (ATT)	ATT_1 - Bad : Good ATT_2 - Foolish : Wise ATT_3 - Useless : Useful ATT_4 - Harmful : Beneficial	$\alpha = .864$	.691 .728 .675 .757
Subjective Norms (SUN)	SUN_1 - Most people who are important to me would find this acceptable. SUN_2 - Most people who are important to me would understand this behavior. SUN_3 - My peers would find this acceptable. SUN_4 - My peers would understand this behavior.	$\alpha = .872$	.685 .733 .755 .734
Perceived Behavioral Control (PBC)	PBC_1 - I can easily report my GenAI use less extensively than required. PBC_2 - I am confident I can control what I report about my GenAI use. PBC_3 - It is entirely up to me how much I disclose about my GenAI use.	$\alpha = .731$	.523 .659 .494
Policy awareness (PA)	PA_1 - I understand the GenAI rules that apply to my courses and assessments. PA_2 - I understand when GenAI use is allowed in assessed work. PA_3 - I know what types of GenAI use need to be reported in assessed work. PA_4 - It is clear to me how much detail I should provide about my GenAI use.	$\alpha = .882$	.745 .753 .791 .699
Academic Self-Efficacy (ASE)	ASE_1 - I can solve difficult academic tasks without using GenAI tools if I try hard enough. ASE_2 - I know I can stick to my study goals and achieve them without relying on GenAI tools. ASE_3 - I am confident I can complete difficult assignments without GenAI. ASE_4 - If my peers can complete study tasks without GenAI, so can I.	$\alpha = .845$	.673 .721 .768 .567

Note. $n = 431$; $p < .05 = *$, $p < .01 = **$, $p < .001 = ***$; α = Cronbach's Alpha, r = Pearson correlation coefficient.

3.4. Operational Conceptual Model

Following the exploratory factor analysis (EFA) and reliability analysis, the constructs PFD, PPD, FDI, PDI, CER, SER, SOS, ATT, SUN, PBC, PA, and ASE were found to demonstrate satisfactory levels of validity and reliability. Based on these results, an operational conceptual model was developed. The model outlines the proposed hypotheses, the number of measurement items associated with each construct, and the original sources from which these items were adapted (see Figure 3)




Figure 3 – Operational conceptual model

3.5. Data Analysis Approach

This section describes how the collected data were analyzed. All analyses were conducted in SPSS. Before the main analysis, several preliminary checks were carried out to ensure the quality and integrity of the dataset.

The data preparation involved four steps. First, a missing data analysis was conducted to identify missing values across the dataset. Second, outlier detection was performed to identify extreme values. Third, Insufficient Effort Responding (IER) was carried out to detect careless or inattentive responding, using post-hoc indicators based on response time and intra-individual response variability (Huang et al., 2015). Fourth, responses containing logically impossible answers were removed. The reasoning and results for each step are described in detail in Appendix 5.

After data preparation, the measurement instrument was examined. An exploratory factor analysis was used to assess the underlying factor structure and construct validity of the scales. Subsequently, a reliability analysis was performed to evaluate their internal consistency. Both are reported in Section 3.3.

The conceptual model was tested using multiple hierarchical regression (Field, 2018). Because the model contains four dependent variables, each non-disclosure outcome was analyzed in a separate regression, namely past full non-disclosure, past partial non-disclosure, full non-disclosure intention, and partial non-disclosure intention. In each regression, predictors were entered in three blocks that follow the theoretical structure of the model. The first block contained the deterrence perceptions, namely the certainty, severity, and swiftness of sanctions. The second block added the Theory of Planned Behavior constructs, namely attitude, subjective norms, and perceived behavioral control. The third block added the two individual-level predictors, academic self-efficacy and policy awareness. This ordering moves from the external enforcement environment, to how students evaluate the act, and finally to individual characteristics. Entering deterrence first also provides a conservative test of the Theory of Planned Behavior, since its constructs are credited only with the variance they explain beyond deterrence.

This block structure shows how much each theoretical perspective contributes to the explanation of non-disclosure. For every block, the explained variance (R^2), the change in explained variance (ΔR^2), and the associated F-change test are reported (Field, 2018). The standardized coefficients are interpreted at the final block, where all predictors are entered together. One hierarchical table is presented in the main text for each of the four outcomes. To make the comparison between the two theories fully transparent, three further models were estimated per outcome (a deterrence-only model, a Theory of Planned Behavior-only model, and a combined model), which are reported in Appendix 13. The assumptions underlying linear regression were assessed before interpretation and are described in Section 3.5.1 and reported in Appendices 8 to 11.

Finally, an exploratory mediation analysis examined whether the influence of the deterrence perceptions, policy awareness and academic self-efficacy on non-disclosure operated indirectly through the way students evaluate the behavior. In line with the reasoning developed in Chapter 2, this analysis was treated as exploratory rather than as a formal test of hypotheses. Its purpose was to identify possible indirect pathways rather than to establish their existence. Two Theory of Planned Behavior constructs were considered as channels, attitude and subjective norms. Perceived behavioral control was not examined as a mediator. Its items capture perceived capability to report GenAI use, rather than an evaluative or normative channel through which deterrence would plausibly operate. For transparency, the corresponding output is nonetheless reported in Appendix 14.

The indirect effects were estimated using the PROCESS macro Model 4 for SPSS (Hayes, 2013). Model 4 was selected because it separates each predictor's total effect into a direct effect on the outcome and an indirect effect that runs through the proposed mediator. Bias-corrected bootstrapped confidence intervals were calculated based on 5,000 resamples. An indirect effect was regarded as meaningful when the 95% bootstrapped confidence interval did not include zero. Because the design is cross-sectional, these results describe patterns of association that are consistent with mediation, but they cannot establish the temporal order that mediation implies (Bryman, 2016). Figure 4 (p. 36) illustrates the mediation model and the relationships that were tested.




Figure 4 – Hayes' (2013) PROCESS macro Model 4 for mediation

3.5.1. Regression Assumptions

The conceptual model was tested using hierarchical regression (Section 3.5). Before interpreting the models, the four assumptions underlying linear regression were checked for each outcome: linearity and normality of residuals, homoscedasticity, absence of multicollinearity, and independence of observations (Field, 2018). The supporting output is reported in Appendices 8 to 11.

Linearity and normality were assessed with P-P plots and histograms of the residuals, supported by skewness and kurtosis values and by formal tests (Appendix 8). The histograms were reasonably bell-shaped and the P-P plots tracked the diagonal closely, and no skewness or kurtosis value exceeded the ± 1.0 range (George & Mallery, 2019). The Kolmogorov-Smirnov and Shapiro-Wilk tests were significant for most models, but these tests flag trivial deviations in large samples (Field, 2018). Given the visual evidence, the modest skewness and kurtosis values, and a sample of 431, the residuals were considered sufficiently normal (Hair et al., 2014).

Homoscedasticity was examined using scatterplots of the standardized residuals against the standardized predicted values (Appendix 9). The residuals were spread evenly across the predicted values with no systematic pattern, so the assumption of constant variance was met (Field, 2018).

Multicollinearity was assessed using tolerance and the variance inflation factor (Appendix 10). Because the four models share the same eight predictors, their statistics are identical: tolerance ranged from 0.617 to 0.929 and VIF from 1.076 to 1.620, comfortably within accepted limits (Field, 2018). Independence was assessed using the Durbin-Watson statistic, which ranged from 1.823 to 2.065 across the four models, close to 2.0 and within the acceptable range (Appendix 11) (Field, 2018).

Overall, the assumptions were satisfied across all models. The residuals were approximately normal, showed constant variance, and were independent, and the predictors were free of problematic multicollinearity, so the models were considered suitable for interpretation.

4. Results

This chapter presents the results of the study. It first reports the descriptive statistics and correlations (Section 4.1), the prevalence of full and partial non-disclosure (Section 4.2), the regression analyses of the direct effects (Section 4.3), and an exploratory analysis of the indirect effects (Section 4.4). It then provides an overview of the hypothesis testing (Section 4.5) and closes with a set of robustness and supplementary analyses (Section 4.6).

4.1. Descriptive Statistics and Correlations

Before the conceptual model was tested, the means, standard deviations, and bivariate correlations of all study variables were examined. This preliminary step served two purposes. It offered a first description of how students scored on each construct, and it allowed the intercorrelations among the predictors to be screened for signs of multicollinearity before the regression models were estimated (Field, 2018).

The strength of the correlations was interpreted using common benchmarks, where coefficients up to .30 are treated as weak, values between .30 and .70 as moderate, and values above .70 as strong (Field, 2018). The full set of descriptive statistics and correlations is presented in Table 6 (p. 39).

Across the four non-disclosure outcomes, mean scores clustered around the midpoint of the five-point scale. Partial non-disclosure intention was the most strongly endorsed outcome ($M = 3.15$, $SD = 1.07$), whereas the two full non-disclosure measures scored lowest (past full non-disclosure $M = 2.51$, $SD = 1.21$; full non-disclosure intention $M = 2.52$, $SD = 1.24$). The partial forms were therefore rated somewhat higher than the full forms, which fits the idea that students lean more towards withholding detail than towards concealing GenAI use entirely. Among the predictors, academic self-efficacy had the highest mean ($M = 4.05$, $SD = 0.78$), followed by severity ($M = 3.67$, $SD = 0.94$), perceived behavioral control ($M = 3.64$, $SD = 0.82$), and subjective norms ($M = 3.58$, $SD = 0.91$). Certainty ($M = 2.88$, $SD = 0.94$) and attitude ($M = 2.74$, $SD = 1.04$) scored below the midpoint, which indicates that, on average, students neither viewed detection as highly likely nor evaluated non-disclosure favorably.

Turning to the correlations, the four non-disclosure outcomes were positively and significantly related to one another. The strongest associations appeared within each type of behavior across time. Past full non-disclosure correlated strongly with full non-disclosure intention ($r = .655$, $p < .001$), and past partial non-disclosure correlated strongly with partial non-disclosure intention ($r = .684$, $p < .001$). Correlations across the full and partial forms were weaker, for instance between past partial non-disclosure and full non-disclosure intention ($r = .358$, $p < .001$). This pattern supports the conceptual separation of full and partial non-disclosure, since the two forms did not collapse into a single behavior.

Among the Theory of Planned Behavior predictors, attitude showed the most consistent and substantial links with non-disclosure. Attitude correlated moderately with all four outcomes, most strongly with partial non-disclosure intention ($r = .555$, $p < .001$). Subjective norms followed a similar but weaker pattern, for example with partial non-disclosure intention

($r = .486$, $p < .001$), as did perceived behavioral control with the same outcome ($r = .370$, $p < .001$). The positive correlations indicate that students with more favorable attitudes toward non-disclosure, stronger perceptions of its social acceptability, and greater perceived control over reporting behavior reported higher levels of non-disclosure.

The deterrence perceptions were more weakly related to the outcomes. Certainty of detection showed small negative correlations with non-disclosure, which reached significance for past full non-disclosure ($r = -.177$, $p < .001$) and full non-disclosure intention ($r = -.187$, $p < .05$) but not for the two partial forms. This is in line with the deterrence expectation that a higher perceived chance of detection accompanies less concealment. Severity showed weak but consistent negative correlations, and these were significant for all four outcomes, for example with full non-disclosure intention ($r = -.178$, $p < .001$). Swiftness was the least related of the three. Its correlations were close to zero, and only its association with partial non-disclosure intention reached significance, and then only weakly ($r = -.103$, $p < .05$), while the other three were non-significant.

The two individual-level predictors behaved differently from one another. Academic self-efficacy correlated negatively and significantly with all four outcomes, with coefficients ranging from $r = -.181$ to $r = -.210$ (all $p < .001$). This suggests that students who felt more capable of completing academic work without GenAI also reported less non-disclosure. Policy awareness showed weaker and less consistent negative correlations. It reached significance for past partial non-disclosure ($r = -.157$, $p < .01$) and partial non-disclosure intention ($r = -.175$, $p < .001$), but not for past full non-disclosure ($r = -.089$, $p > .05$) or full non-disclosure intention ($r = -.090$, $p > .05$).

Finally, the correlations among the predictors were inspected for multicollinearity. The strongest association was that between severity and swiftness ($r = .574$, $p < .001$), which only reflects a moderate relationship. All remaining predictor correlations were weaker, and none approached the .80 threshold. Together with the tolerance and VIF values reported earlier in Section 3.5.1, this indicates that multicollinearity was not a concern (Field, 2018).

Table 6 – Descriptive statistics and intercorrelations of study variables

Variable	Mean	SD	1	2	3	4	5	6	7	8	9	10	11	12
1. PFD	2.51	1.21	–											
2. PPD	2.81	1.11	.469***	–										
3. FDI	2.52	1.24	.655***	.358***	–									
4. PDI	3.15	1.07	.525***	.684***	.582***	–								
5. CER	2.88	0.94	-.177***	-.074	-.187*	-.123	–							
6. SER	3.67	0.94	-.134**	-.160***	-.178***	-.164***	.213***	–						
7. SOS	3.24	1.08	-.082	-.080	-.027	-.103*	.202***	.574***	–					
8. ATT	2.74	1.04	.527***	.462***	.496***	.555***	-.245***	-.332***	-.275***	–				
9. SUN	3.58	0.91	.381***	.352***	.398***	.486***	-.203***	-.312***	-.270***	.530***	–			
10. PBC	3.64	0.82	.263***	.284***	.247***	.370***	-.210***	-.164***	-.141**	.366***	.382***	–		
11. PA	3.55	0.99	-.089	-.157**	-.090	-.175***	-.014	.069	.141**	-.076	-.109*	-.030	–	
12. ASE	4.05	0.78	-.191***	-.210***	-.181***	-.206***	-.030	.032	.027	-.216***	-.107*	.033	.214***	–

Note. $n = 431$, $p < .05 = *$, $p < .01 = **$, $p < .001 = ***$; M = mean, SD = standard deviation; PFD = Past Full Non-Disclosure, PPD = Past Partial Non-Disclosure, FDI = Full Non-Disclosure Intention, PDI = Partial Non-Disclosure Intention, CER = Certainty, SER = Severity, SOS = Swiftness, ATT = Attitude, SUN = Subjective Norms, PBC = Perceived Behavioral Control, PA = Policy awareness, ASE = Academic Self-Efficacy.

4.2. *Prevalence of Full and Partial Non-Disclosure*

The first hypothesis concerned how common each form of non-disclosure was. H1 proposed that partial non-disclosure is more prevalent than full non-disclosure. Since each respondent reported on both forms, the two were compared using paired-samples t-tests (Field, 2018). The comparison was carried out separately for past behavior and for intention, so that the contrast between full and partial non-disclosure was examined while holding the time dimension constant. Given the sample size of 431, the test is robust to departures from normality in the difference scores, so the parametric test was retained (Field, 2018). A Wilcoxon signed-rank test was run alongside each comparison as a robustness check. Cohen's d is reported for each test, because the large sample means that even small differences can reach significance, so effect size is needed to judge their practical weight.

For past behavior, partial non-disclosure ($M = 2.81$, $SD = 1.11$) was reported more often than full non-disclosure ($M = 2.51$, $SD = 1.21$). This difference was significant, $t(430) = 5.14$, $p < .001$, and the effect was small ($d = 0.25$, 95% CI [0.15, 0.34]).

The same pattern was stronger for intention. Partial non-disclosure intention ($M = 3.15$, $SD = 1.07$) clearly exceeded full non-disclosure intention ($M = 2.52$, $SD = 1.24$), $t(430) = 12.19$, $p < .001$, with a medium effect ($d = 0.59$, 95% CI [0.49, 0.69]).

A Wilcoxon signed-rank test confirmed both comparisons. Partial non-disclosure ranked significantly higher than full non-disclosure for past behavior ($Z = 6.56$, $p < .001$) and for intention ($Z = 11.03$, $p < .001$), which matches the results of the paired-samples t-tests.

H1 was therefore supported for both past behavior and intention. Students were more inclined to withhold some details about their GenAI use than to conceal it entirely. This pattern was evident both in what they reported having done and in what they intended to do. This is consistent with the idea that partial non-disclosure is a lower-risk strategy, since acknowledging some use can signal good faith while still limiting exposure (Barrett & Pack, 2023).

4.3. Hierarchical Regression Analysis

This section reports the hierarchical regression analyses used to test the direct effects in the conceptual model. Each of the four non-disclosure outcomes was regressed on the eight predictors, entered in three blocks that follow the theoretical structure of the model: the deterrence perceptions first (Model 1), the Theory of Planned Behavior constructs second (Model 2), and the two individual-level predictors last (Model 3). This ordering makes it possible to assess how much each theoretical perspective contributes beyond the one entered previously. Standardized coefficients are interpreted using Model 3, in which all predictors are included simultaneously. The results for the four outcomes are reported in Table 7 to Table 10 (p. 44 and 45)

All four final models were statistically significant ($p < .001$) and explained a moderate to substantial share of the variance in non-disclosure (Cohen, 1988; Plonsky & Ghanbar, 2018). The explained variance was highest for partial non-disclosure intention (adjusted $R^2 = .396$) and lowest for past partial non-disclosure (adjusted $R^2 = .258$), with the two remaining outcomes in between (past full non-disclosure adjusted $R^2 = .300$; full non-disclosure intention adjusted $R^2 = .297$). According to commonly used benchmarks, these values reflect moderate to large effects and compare favorably with results reported in related studies of GenAI-related academic misconduct (Plonsky & Ghanbar, 2018).

The block structure revealed a consistent pattern across all four outcomes. The deterrence perceptions were entered first and explained only a small share of the variance, ranging from 2.8% for past partial non-disclosure to 6.6% for full non-disclosure intention. Adding the Theory of Planned Behavior constructs produced a large and highly significant increase in explained variance in every model (ΔR^2 between .220 and .350, all $p < .001$). The individual-level predictors were entered last. They added a small amount of explained variance. This increase was significant for the two partial outcomes (past partial $\Delta R^2 = .024$, $p < .001$; partial intention $\Delta R^2 = .022$, $p < .001$) but not for the two full outcomes (both $\Delta R^2 = .009$, $p > .05$). The psycho-social constructs accounted for the largest share of the explained variance. They added approximately four to nine times as much explanatory power as the deterrence block. This is the first indication that non-disclosure is shaped more by how students evaluate the behavior than by the enforcement environment.

To make the contrast between the two frameworks explicit, each was also estimated on its own and then combined, for all four outcomes (Appendix 13). On their own, the deterrence perceptions explained little variance in non-disclosure (R^2 between .028 and .066), whereas the Theory of Planned Behavior constructs explained far more (R^2 between .239 and .376). Combining the two added only a marginal gain (ΔR^2 between .008 and .029). In the combined model, certainty and severity lost significance for every outcome, and only swiftness remained, again in relation to full non-disclosure intention alone. Combining the theories therefore clarified the structure of the effects rather than adding explanatory power.

Deterrence perceptions

The three deterrence perceptions were weak predictors once the other constructs were controlled. Certainty of detection had no significant effect on any of the four outcomes in the final models (all $p > .05$), so H3 was not supported. Notably, certainty was a significant negative predictor in Model 1 for the two full outcomes. However, this effect disappeared after the Theory of Planned Behavior constructs were added. This might indicate that certainty influenced non-disclosure indirectly through students' evaluations of the behavior rather than through a direct effect. Severity of sanctions likewise showed no significant direct effect in any final model (all $p > .05$), so H4 was not supported. Swiftiness produced the most notable result. Contrary to expectations, the relationship was positive rather than negative. Swiftiness was significantly related to full non-disclosure intention ($\beta = .204, p < .001$) and, more weakly, to partial non-disclosure intention ($\beta = .093, p = .046$). It also approached significance for past partial non-disclosure ($\beta = .098, p = .058$). Because these effects ran opposite to the hypothesized direction, H5 was not supported. This counterintuitive pattern indicates that a faster expected sanction did not reduce non-disclosure. Instead, it was associated with a greater tendency to withhold disclosure, most notably in relation to intentions to conceal GenAI use entirely.

Theory of Planned Behavior

The psycho-social constructs were the strongest and most consistent predictors in the study. Attitude toward non-disclosure emerged as the strongest predictor in every model. The effect was consistently positive and highly significant, ranging from $\beta = .346$ ($p < .001$) for past partial non-disclosure to $\beta = .431$ ($p < .001$) for past full non-disclosure. H6 was therefore supported across all four outcomes. Subjective norms were also positively and significantly related to non-disclosure in every model (β between .120 and .236, all $p \leq .021$), so H7 was supported. Perceived behavioral control was more selective. It was positively and significantly related to past partial non-disclosure ($\beta = .129, p = .006$) and partial non-disclosure intention ($\beta = .160, p < .001$). However, it was not significantly related to either of the full non-disclosure outcomes (both $p > .05$). H8 was therefore supported for the partial forms of non-disclosure but not for the full forms.

Individual-level predictors

The two individual-level predictors showed the same full-versus-partial distinction as perceived behavioral control. Academic self-efficacy was negatively and significantly related to the two partial outcomes. These were past partial non-disclosure ($\beta = -.105, p = .017$) and partial non-disclosure intention ($\beta = -.081, p = .041$). Similar negative relationships were observed for the two full outcomes, but these did not reach statistical significance (past full non-disclosure $p = .055$; full non-disclosure intention $p = .071$). H9 was therefore supported for the partial forms and only marginal for the full forms. Policy awareness followed the same pattern. It was negatively and significantly related to past partial non-disclosure ($\beta = -.101, p = .019$) and partial non-disclosure intention ($\beta = -.112, p = .004$). However, it had no significant effect on either of the full non-disclosure outcomes (both $p > .05$). H10 was therefore supported for the partial forms of non-disclosure only.

Taken together, the four models point to a consistent core of predictors alongside a set of effects that were specific to particular forms of non-disclosure. Attitude and subjective norms predicted all four outcomes and were the two strongest predictors throughout. In contrast, perceived behavioral control, academic self-efficacy, and policy awareness were significant only for the partial forms of non-disclosure, not for the full forms. The deterrence perceptions were weak throughout, with swiftness the only one to reach significance and then in the direction opposite to that predicted. A notable pattern was that the Theory of Planned Behavior constructs outperformed deterrence perceptions. Their contribution to explained variance was four to nine times greater across the four models.

Table 7 – Hierarchical regression results for past full non-disclosure

Explanatory variables	Past full non-disclosure								
	Model 1		Model 2		Model 3		95% CI	β	p
	B	t	B	t	B	t			
<i>Deterrence</i>									
Certainty	−.201	−3.21**	−.064	−1.16	−.074	−1.35	[−.181, .034]	−.058	.179
Severity	−.137	−1.83	.045	0.69	.041	0.62	[−.088, .169]	.032	.534
Swiftness	.012	0.19	.088	1.57	.094	1.66	[−.018, .205]	.083	.099
<i>Theory of Planned Behavior</i>									
Attitude			.528	9.04***	.501	8.39***	[.383, .618]	.431	<.001
Subjective norms			.190	2.87**	.182	2.75**	[.052, .312]	.137	.006
Perceived behavioral control			.072	1.08	.088	1.32	[−.043, .218]	.059	.189
<i>Individual-level</i>									
Academic self-efficacy					−.127	−1.93	[−.256, .002]	−.082	.055
Policy awareness					−.045	−0.88	[−.145, .055]	−.037	.378
R ²			.041		.304				.313
Adjusted R ²			.034		.295				.300
ΔR^2			.041		.263				.009
F-statistic		6.09***		30.92***					24.06***
df (df1, df2)		(3, 427)		(6, 424)					(8, 422)

Note. n = 431. B = unstandardized coefficient; t = t-value; β = standardized coefficient; CI = 95% confidence interval for B; p = significance of the predictor in the full model. CI, β , and p are reported for the full model (Model 3). Empty cells indicate predictors not yet entered at that step. p < .05 = *, p < .01 = **, p < .001 = ***.

Table 8 – Hierarchical regression results for past partial non-disclosure

	Past partial non-disclosure								
	Model 1		Model 2		Model 3				
Explanatory variables	B	t	B	t	B	t	95% CI	β	p
<i>Deterrence</i>									
Certainty	−.051	−0.88	.071	1.35	.055	1.07	[−.046, .157]	.047	.285
Severity	−.193	−2.80**	−.043	−0.69	−.051	−0.82	[−.172, .070]	−.043	.411
Swiftness	.024	0.40	.087	1.63	.101	1.90	[−.003, .206]	.098	.058
<i>Theory of Planned Behavior</i>									
Attitude			.400	7.20***	.368	6.54***	[.257, .478]	.346	<.001
Subjective norms			.160	2.54*	.145	2.32*	[.022, .268]	.120	.021
Perceived behavioral control			.154	2.43*	.173	2.76**	[.050, .297]	.129	.006
<i>Individual-level</i>									
Academic self-efficacy					−.148	−2.39*	[−.270, −.026]	−.105	.017
Policy awareness					−.113	−2.35*	[−.207, −.019]	−.101	.019
R ²			.028		.247				.272
Adjusted R ²			.021		.237				.258
ΔR ²			.028		.220				.024
F-statistic		4.05**		23.24***					19.70***
df (df1, df2)		(3, 427)		(6, 424)					(8, 422)

Note. n = 431. B = unstandardized coefficient; t = t-value; β = standardized coefficient; CI = 95% confidence interval for B; p = significance of the predictor in the full model. CI, β , and p are reported for the full model (Model 3). Empty cells indicate predictors not yet entered at that step. p < .05 = *, p < .01 = **, p < .001 = ***.

Table 9 – Hierarchical regression results for full non-disclosure intention

	Full non-disclosure intention								
	Model 1		Model 2		Model 3				
Explanatory variables	B	t	B	t	B	t	95% CI	β	p
<i>Deterrence</i>									
Certainty	−.220	−3.49***	−.090	−1.60	−.100	−1.79	[−.211, .010]	−.076	.075
Severity	−.286	−3.78***	−.106	−1.58	−.111	−1.66	[−.243, .020]	−.085	.097
Swiftness	.151	2.29*	.227	3.93***	.234	4.04***	[.120, .348]	.204	<.001
<i>Theory of Planned Behavior</i>									
Attitude			.468	7.81***	.441	7.22***	[.321, .562]	.371	<.001
Subjective norms			.260	3.83***	.251	3.69***	[.117, .384]	.185	<.001
Perceived behavioral control			.046	0.67	.061	0.90	[−.073, .195]	.041	.370
<i>Individual-level</i>									
Academic self-efficacy					−.122	−1.81	[−.255, .010]	−.077	.071
Policy awareness					−.060	−1.14	[−.162, .043]	−.048	.254
R ²			.066		.301				.310
Adjusted R ²			.060		.291				.297
ΔR ²			.066		.234				.009
F-statistic		10.13***		30.37***					23.68***
df (df1, df2)		(3, 427)		(6, 424)					(8, 422)

Note. n = 431. B = unstandardized coefficient; t = t-value; β = standardized coefficient; CI = 95% confidence interval for B; p = significance of the predictor in the full model. CI, β , and p are reported for the full model (Model 3). Empty cells indicate predictors not yet entered at that step. p < .05 = *, p < .01 = **, p < .001 = ***.

Table 10 – Hierarchical regression results for partial non-disclosure intention

	Partial non-disclosure intention								
	Model 1		Model 2		Model 3				
Explanatory variables	B	t	B	t	B	t	95% CI	β	p
<i>Deterrence</i>									
Certainty	−.103	−1.87	.044	0.97	.030	0.68	[−.058, .118]	.027	.500
Severity	−.162	−2.45*	.022	0.41	.015	0.29	[−.090, .121]	.014	.774
Swiftness	−.002	−0.04	.077	1.65	.092	2.00*	[.002, .183]	.093	.046
<i>Theory of Planned Behavior</i>									
Attitude			.416	8.61***	.391	8.01***	[.295, .487]	.382	<.001
Subjective norms			.290	5.29***	.276	5.09***	[.169, .382]	.236	<.001
Perceived behavioral control			.192	3.50***	.207	3.81***	[.100, .315]	.160	<.001
<i>Individual-level</i>									
Academic self-efficacy					−.111	−2.05*	[−.216, −.005]	−.081	.041
Policy awareness					−.120	−2.89**	[−.202, −.038]	−.112	.004
R ²			.035		.385				.407
Adjusted R ²			.028		.377				.396
ΔR ²			.035		.350				.022
F-statistic		5.16**		44.30***					36.25***
df (df1, df2)		(3, 427)		(6, 424)					(8, 422)

Note. n = 431. B = unstandardized coefficient; t = t-value; β = standardized coefficient; CI = 95% confidence interval for B; p = significance of the predictor in the full model. CI, β , and p are reported for the full model (Model 3). Empty cells indicate predictors not yet entered at that step. p < .05 = *, p < .01 = **, p < .001 = ***.

4.3.1. Comparing Predictors Across Non-Disclosure Forms

The four regressions above suggested that some predictors operated differently across the forms of non-disclosure. Hypothesis 2 proposed that these predictors do indeed differ across the four outcomes, in particular between full and partial non-disclosure. Testing this directly requires more than comparing whether a predictor was significant in one model but not another. A predictor may be significant in one model and not in another, even when the two coefficients do not differ significantly from each other. Because the four outcomes were measured on the same respondents, their coefficients are also statistically dependent, which an ordinary comparison would ignore. Because the four outcomes were measured for the same respondents, their coefficients were not independent. To account for this, the four equations were estimated jointly using seemingly unrelated regression, and Wald tests were used to compare predictor coefficients across outcomes (Zellner, 1962). The full set of tests is reported in Appendix 12.

At the overall level, the results supported Hypothesis 2. A global test rejected the null hypothesis that all predictors had identical coefficients across the four outcomes ($\chi^2(24) = 39.46, p = .024$). This indicates that the predictors of non-disclosure operated differently across the four forms. This difference came mainly from the deterrence perceptions. Looking at the predictors one by one, the effects of severity and swiftness varied across outcomes, while attitude, subjective norms, perceived behavioral control, academic self-efficacy, and policy awareness did not. In short, the psycho-social and individual-level predictors worked in a similar way across all four forms, whereas the deterrence perceptions did not.

The differences between the forms of non-disclosure were more nuanced than a simple distinction between full and partial non-disclosure. Overall, most predictors showed similar effects for both forms. This suggests that full and partial non-disclosure were not driven by fundamentally different predictors. Instead, the main difference concerned the deterrence perceptions. After correcting for multiple testing, only one effect remained significant. The relationship between severity and non-disclosure differed between past behavior and intention (Holm-adjusted $p = .025$). Severity was more strongly associated with full non-disclosure in past behavior, but with partial non-disclosure in intention. Swiftness showed a similar pattern, but this effect was no longer significant after correction for multiple testing. All other differences were significant only before correction and should therefore be interpreted with caution.

Hypothesis 2 was supported, although the differences between the forms of non-disclosure were limited. Predictor effects varied across the four outcomes, but these differences were mainly found in the deterrence perceptions. In contrast, students' evaluations of the behavior and their individual characteristics showed similar effects across all forms of non-disclosure. Full and partial non-disclosure were therefore not driven by different sets of predictors. Instead, the main difference was that deterrence perceptions, particularly severity, related to the two forms differently for past behavior and future intentions.

4.4. Exploratory Mediation Analysis

The regressions above tested how each predictor related to non-disclosure directly. This section asks whether those effects also ran through the psycho-social constructs of the Theory of Planned Behavior. Because all constructs were measured at one point in time, the question cannot be answered causally. The analysis is therefore exploratory, showing whether the associations are consistent with an indirect pathway rather than proving one (Bryman, 2016). Indirect effects were estimated with parallel multiple mediation, using attitude, subjective norms, and perceived behavioral control as simultaneous channels (Hayes, 2013). The full models are reported in Appendix 14.

Attitude functioned as the main channel. It carried the effects of certainty, severity, and academic self-efficacy onto all four outcomes. The strongest of these paths ran through attitude for certainty and severity, so their link to non-disclosure was fully indirect. This is most striking for severity. It showed no significant direct effect in the regressions, yet it was clearly related to non-disclosure once attitude was modelled as a channel. Harsher expected sanctions were tied to a less favorable view of non-disclosure, which was in turn tied to less concealment.

Subjective norms formed a second, narrower channel. They carried certainty and severity, but not academic self-efficacy. Perceived behavioral control played almost no role, transmitting only a weak effect for certainty on the partial outcomes. Swiftness and policy awareness did not work through any channel at all. Where they mattered, their effect was direct, as reported in Section 4.2.

Taken together, the pattern supports the study's integration logic (Nguyen-Phuoc et al., 2024). Deterrence perceptions that carry evaluative or cost content, namely certainty and severity, reached non-disclosure partly through the way students judged the behavior. Perceptions about the timing of enforcement, or about knowledge of the rules, did not. This suggests that combining the two theories added explanatory value that a model of direct effects alone would miss.

4.5. Overview of Hypothesis Testing

This section summarizes the outcomes of the ten hypotheses tested in Chapter 4. Table 11 presents the full overview. Most predictors were tested against all four non-disclosure outcomes, so the table reports the result for each outcome separately rather than as a single overall verdict. The exploratory indirect-effects analysis in Appendix 14 carried no hypotheses and is therefore not included here.

The first two hypotheses concerned the structure of non-disclosure itself. H1 held that partial non-disclosure is more prevalent than full non-disclosure. It was supported for both past behavior and intention. H2 held that the predictors differ across the four outcomes. It was also supported. The global test was significant, and the differences were concentrated in the deterrence perceptions, most clearly for severity.

The remaining hypotheses concerned the direct effects of the predictors. Here the results showed a clear divide between the two theoretical perspectives. The deterrence perceptions were weak. Certainty (H3) and severity (H4) had no significant direct effect on any outcome. Swiftness (H5) was significant for the two intention outcomes, but in the direction opposite to the one predicted, so it too was not supported.

The Theory of Planned Behavior constructs performed far more consistently. Attitude (H6) and subjective norms (H7) were positively related to all four outcomes, so both were supported. Perceived behavioral control (H8) was positive and significant for the two partial outcomes only.

The two individual-level predictors showed the same full-versus-partial split. Academic self-efficacy (H9) and policy awareness (H10) were both negatively related to the two partial outcomes, but not to the two full outcomes. A consistent pattern therefore runs through the results. Several predictors were specific to the partial forms of non-disclosure, whereas attitude and subjective norms held across every outcome.

Table 11 – Hypothesis overview

Hypotheses	Results
H1: Partial non-disclosure of GenAI use is more prevalent than full non-disclosure.	Supported
H2: The predictors of non-disclosure differ across the four forms of non-disclosure, in particular between full and partial non-disclosure	Supported
H3: Perceived certainty of detection is negatively related to students' non-disclosure.	Not supported
H4: Perceived severity of sanctions is negatively related to students' non-disclosure.	Not supported
H5: Perceived swiftness of sanctions is negatively related to students' non-disclosure.	Not supported
H6: Attitude toward GenAI non-disclosure is positively related to students' non-disclosure.	Supported
H7: Subjective norms toward non-disclosure are positively related to students' non-disclosure.	Supported
H8: Perceived behavioral control over non-disclosure is positively related to students' non-disclosure.	Supported for the partial outcomes
H9: Academic self-efficacy to complete assessed academic tasks without GenAI assistance is negatively related to students' non-disclosure.	Supported for the partial outcomes
H10: Policy awareness is negatively related to students' non-disclosure.	Supported for the partial outcomes

Note. n = 431.

4.6. Robustness and supplementary analyses

The analyses presented so far tested the study's hypotheses. This section reports additional analyses that go beyond those hypotheses. Some are robustness checks, examining whether the main findings hold when sample characteristics are taken into account. Others are supplementary analyses that provide additional context that were not captured by the hypothesis tests. A brief summary of each analysis is provided below, with the full results available in the appendices. None of these analyses changes the conclusions of the main models, and any differences are noted where relevant.

4.6.1. Demographic controls

The sample leaned toward master's students and women (Section 3.2.1). This raises the question of whether the results reflect who took part rather than the relationships the study proposes. To check this, the combined model was re-estimated with demographic controls. For each of the four outcomes, age, gender, and level of study were added as a second block after the eight main predictors, which left 418 cases (Appendix 15). Adding this block raised explained variance only marginally (ΔR^2 between .013 and .019). More importantly, it left every main predictor unchanged in both size and significance, so no predictor gained or lost significance once composition was accounted for.

Among the controls, only age was related to non-disclosure, and then weakly, with older students reporting slightly less non-disclosure on two of the four outcomes. Gender and level of study were non-significant throughout. The absence of a gender effect is itself informative, since it suggests that non-disclosure is shaped by how students evaluate the behavior rather than by demographic group. The main findings therefore held once sample composition was taken into account.

4.6.2. Survey language

The survey was offered in Dutch and English, which allowed a check on whether non-disclosure and its drivers differed by response language. This analysis is exploratory, and one caveat frames it: the language of response is not a reliable indicator for nationality, since some Dutch students may have chosen the English survey (Section 3.2). Comparing the two groups across the twelve constructs, most were similar, with no significant difference on eight of them, including attitude, subjective norms, and perceived behavioral control (Appendix 16). The clearest difference was in policy awareness, which English respondents reported markedly higher ($d = .65$). The remaining differences were small, with Dutch respondents reporting slightly more partial non-disclosure and English respondents rating severity a little higher.

Running the eight-predictor regression separately for each group showed that the model behaved in largely the same way for both. Attitude and subjective norms were the central drivers of non-disclosure regardless of language, mirroring the main analysis. The one notable exception was academic self-efficacy, which was negatively and significantly related to non-disclosure in all four Dutch models but in none of the English ones. This difference is worth noting, but it should be interpreted with caution. The Dutch group was smaller, a significant

result in one group but not the other does not necessarily mean the effects are different. Overall, the main findings were consistent across language groups, with academic self-efficacy the only result that may warrant closer investigation.

4.6.3. Policy awareness and the deterrence perceptions

The main models treated the deterrence perceptions and policy awareness as separate predictors of non-disclosure. A further question is whether policy awareness also shapes the deterrence perceptions themselves. In particular, it examines whether a clearer understanding of the rules influences how likely, severe, or swift students expect sanctions to be. To test this, three simple regressions were run, each with policy awareness as the sole predictor of one deterrence perception (Appendix 17). Policy awareness was unrelated to certainty and to severity, and both models explained almost none of the variance. The only significant link was with swiftness ($\beta = .141$, $p = .003$), which was positive but small and accounted for about 2% of the variance. Policy awareness therefore had no meaningful direct relationship with the deterrence perceptions, which supports treating them as distinct constructs in the main models.

4.6.4. Open-ended comments

The survey ended with an optional field for further thoughts, which 68 respondents used. These comments were examined through a light qualitative content analysis and grouped into ten themes (Appendix 18). The analysis is exploratory and supplementary. It was not designed to test the hypotheses. Instead, it intended to show, in students' own words, the reasoning behind the disclosure decisions the study models.

The most prominent theme concerned the conditions under which GenAI use is considered legitimate. Respondents did not judge AI use as simply acceptable or unacceptable. Instead, they based their views on how the tool was used: as support rather than a replacement, for peripheral rather than core tasks, and with its output carefully checked rather than accepted without question. A closely related theme was that the work was still considered the student's own. Respondents argued that the ideas, writing style, and final responsibility remained with the student. Alongside these, students described a wide range of concrete uses, from explaining concepts and generating code to summarizing material and editing text. Taken together, the comments show that GenAI use is not one behavior but many practices, each carrying different implications for whether disclosure feels necessary. This directly supports the study's decision to treat non-disclosure as more than a binary act.

A second group of comments spoke to the disclosure decision itself. Many respondents were unsure about what disclosure actually requires. They were uncertain where ordinary help ends and reportable AI use begins. They also questioned how the same declaration can cover both minor and substantial uses. Some pointed to the effort of documenting every step, while others argued that where prompting reflects their own thinking there may be little to declare. Running through these comments was a disclosure dilemma. Disclosing even careful, minor use was felt to lower the credibility of the work, yet non-disclosure raised the risk of being accused of concealment, including false accusation. A further set of comments linked the

decision to students' circumstances. They mentioned workload, insufficient feedback, and competition with peers. This suggests that non-disclosure may not always reflect a moral decision, but can also result from structural pressures.

The remaining comments addressed broader issues. These included concerns about overreliance on GenAI, calls for clearer institutional policies, and the view that GenAI use has become a normal part of student work. As a result, disclosure may be seen as less necessary.

Read as a whole, the comments triangulate the quantitative model and give its constructs a human texture. They portray non-disclosure not simply as a matter of honesty, but as a judgment students make under specific task, social, and institutional conditions. The full content analysis is reported in Appendix 18.

5. Discussion and Conclusion

This chapter interprets the results presented in Chapter 4 and draws the study to a close. It first discusses the findings in relation to the hypotheses and prior research (Section 5.1). It then answers the central research question (Section 5.2). Building on this, it sets out recommendations for institutions (Section 5.3). It then presents the study's contributions to theory and practice (Section 5.4). It closes with the limitations of the study and directions for future research (Section 5.5).

5.1. Discussion

This section interprets the findings in light of the theoretical framework and previous research. It follows the six sub-questions from Section 1.2, grouped by theme. It begins with the conceptualization of non-disclosure and the validation of its measures (Section 5.1.1). It then discusses the three groups of predictors in turn: the deterrence perceptions (Section 5.1.2), the psycho-social constructs (Section 5.1.3), and the individual-level factors (Section 5.1.4). It closes with a synthesis of why the enforcement environment mattered less than how students evaluated disclosure (Section 5.1.5).

5.1.1. Conceptualizing non-disclosure

The central conceptual move of this study was to treat non-disclosure as two related behaviors rather than a single yes-or-no act. Full non-disclosure means concealing GenAI use entirely. Partial non-disclosure means acknowledging some GenAI use while withholding important details about its extent, type, or location. Each form was measured for past behavior and for intention. This section considers whether that distinction earned its place. It draws together the evidence for the first two sub-questions: *How can students' non-disclosure of GenAI use be conceptualized and measured as more than a binary act, distinguishing full from partial non-disclosure?* (SQ1) and *To what extent do the newly developed non-disclosure measures demonstrate construct validity and internal consistency?* (SQ2)

The measurement evidence offers a first form of support. In both factor analyses, full and partial non-disclosure emerged as separate components rather than collapsing into one. This separation was in fact clearer than the one between past behavior and intention (Section 3.3.2). The scales were also internally consistent. The partial-form scales reached acceptable to strong reliability, and the two-item full non-disclosure scales showed adequate inter-item correlations. The distinction is therefore not only conceptual but measurable. Each form behaved as a coherent and separable construct.

The results showed a consistent pattern. The four outcomes were most strongly correlated within each form over time, while correlations between the full and partial forms were much weaker (Section 4.1). This suggests that each form is internally consistent, but that the forms do not reflect a single underlying behavior. They also differ in how common they are. Partial non-disclosure was reported more often than full non-disclosure, both for past behavior and, more clearly, for intention (H1). Under-reporting was therefore not a marginal case but the more prevalent tendency. A binary treatment would have folded it into outright concealment.

The two forms were also not driven by the same factors. The predictors differed across the four outcomes (H2). This difference was concentrated in the deterrence perceptions and clearest for severity. The pattern of significance also varied for the other predictors. Perceived behavioral control, academic self-efficacy, and policy awareness reached significance for the partial outcomes but not the full ones. Attitude and subjective norms, by contrast, were significant across every outcome. Collapsing the two behaviors would have averaged these form-specific effects away.

Taken together, these findings answer the first two sub-questions. Non-disclosure can be meaningfully conceptualized as more than a binary act. The measures developed to capture it showed sound construct validity and internal consistency. The distinction is also measurable in practice, and the two forms differ in both their prevalence and their predictors. That is what justifies carrying it through the remainder of the discussion.

5.1.2. Deterrence mechanisms

The third sub-question asked: *How do the perceived certainty, severity, and swiftness of sanctions relate to students' non-disclosure of GenAI use?* Deterrence Theory holds that each perception raises the expected cost of a violation and so discourages it (Gibbs, 1979). All three were therefore expected to reduce non-disclosure. The results gave this expectation only limited support. As a block, the deterrence perceptions were weak. They explained between 2.8% and 6.6% of the variance, while the Theory of Planned Behavior constructs added four to nine times as much (Section 4.3). This alone suggests that enforcement mattered less than how students evaluated concealment. The construct-level findings show why.

Hypothesis 3 proposed that perceived certainty of detection is negatively related to non-disclosure. This followed a consistent strand of deterrence research, in which a higher perceived chance of detection reduces academic misconduct (Huang et al., 2025; Zhang et al., 2025). The hypothesis was not supported as a direct effect. Certainty had no significant relationship with any outcome once the other predictors were entered. The picture was not flat, though. In the correlations, certainty was negatively related to the two full outcomes but not the partial ones (Section 4.1), and it remained a significant negative predictor of these outcomes in the first regression block, before the psycho-social constructs were added (Section 4.3). This fits how detectable each form is. Full non-disclosure is more exposed, since a single detected use contradicts a claim of none. Partial non-disclosure is harder to prove, because some use is already admitted. Detection risk therefore registered mainly for the form students believed could be caught. The effect then vanished once attitude was added, and the mediation analysis showed it ran through attitude (Section 4.4). Certainty appears to have shaped how students judged concealment, rather than deterring them directly. Two further points support this reading. Students' perceived certainty of detection was low on average ($M = 2.88$), and reliable detection of GenAI use remains difficult in practice (Ardito, 2025; Deep et al., 2025). The finding still aligns with the general deterrence prediction that detection risk lowers misconduct (Lou et al., 2025; Zhang et al., 2025). It also fits evidence

that students adapt their use to avoid detection rather than disclose it (Adnin et al., 2025). The theoretical takeaway is that certainty operated as an input to evaluation, not as a standalone deterrent.

Hypothesis 4 proposed that perceived severity of sanctions is negatively related to non-disclosure. This drew on studies reporting that harsher expected consequences lower the intention to misuse GenAI (Goyal et al., 2026; Lou et al., 2025; Zhang et al., 2025). The hypothesis was not supported, since severity showed no direct effect on any of the four outcomes. This does not make severity irrelevant. Its correlations with non-disclosure were consistently negative, and the exploratory mediation analysis showed it operated fully through attitude (Section 4.4). Harsher expected consequences were tied to a less favorable view of concealment, which was in turn tied to less of it. The absent direct effect is also revealing in light of the descriptive results. Students rated sanctions as fairly severe ($M = 3.67$), yet perceived detection as unlikely ($M = 2.88$).

A long-standing finding in deterrence research is that severity does little on its own when certainty is low, because students do not expect to face the consequences (Nagin & Pogarsky, 2004). That pattern fits the present data. A second mechanism may work in the same direction. Strict and stigmatizing sanctions can push students toward concealment rather than honesty, which would offset any deterrent effect (Nguyen & Goto, 2024; Qu et al., 2025). The result therefore contrasts with studies reporting a direct severity effect (Goyal et al., 2026). However, it fits a more cautious strand of research suggesting that informal sanctions can outweigh formal ones (Zhang et al., 2025). For the framework, this positions severity as an input to evaluation rather than a direct deterrent, which supports the study's integration logic (Nguyen-Phuoc et al., 2024).

Hypothesis 5 proposed that perceived swiftness of sanctions is negatively related to non-disclosure. This rested on evidence that swift enforcement deters, whereas delayed consequences are discounted (Dench & Joyce, 2022; Lou et al., 2025; Nagin & Pogarsky, 2004). The hypothesis was not supported, and the result ran opposite to the prediction. Swiftness was positively related to the two intention outcomes, most strongly to full non-disclosure intention ($\beta = .204, p < .001$) and weakly to partial non-disclosure intention ($\beta = .093, p = .046$). A faster expected sanction thus accompanied a stronger intention to conceal, not a weaker one. Several readings are consistent with this, and they are not mutually exclusive. One is perceived procedural threat. A system expected to act quickly may feel like one that decides before it listens, leaving little room to explain how GenAI was used. Concealing everything then becomes a way to avoid triggering it (Nguyen & Goto, 2024; Qu et al., 2025). A second explanation is that swift enforcement strengthens the incentive to conceal without changing the underlying behavior. Rather than using GenAI less, students may hide it more thoroughly, which mirrors the adaptation mechanism reported by Adnin et al. (2025). A third reading is methodological. The swiftness of sanctions scale was reduced to two items after SOS_1 was removed in the factor analysis (Section 3.3.2). The two retained items therefore capture a narrower facet of swiftness, namely how quickly a decision follows discovery, and may not represent the construct as it was originally conceptualized. The

positive coefficient may thus reflect an incompletely measured construct rather than a genuine reversal of the deterrence prediction. It is also notable that the effect appeared for intentions but not past behavior, which suits a forward-looking appraisal of a fast system rather than an established habit. The cross-sectional design cannot separate these accounts. What can be said is narrower. Swift expected enforcement did not discourage concealment in this sample, and it pointed the other way. Why it did so is best treated as an open question for future work.

Taken together, the deterrence perceptions did not act as uniform, direct deterrents. Certainty and severity reached non-disclosure only through attitude, and swiftness ran opposite to prediction on a fragile measure. Perceived enforcement was neither strong nor consistent. This weak and uneven pattern is the first clear sign that a deterrence-based account cannot stand on its own. It motivates the turn to the psycho-social predictors examined next.

5.1.3. Psycho-social mechanisms

The fourth sub-question asked: *How do attitude, subjective norms, and perceived behavioral control relate to students' non-disclosure of GenAI use?* The Theory of Planned Behavior holds that these three factors shape behavior through intention, so each was expected to predict non-disclosure (Ajzen, 1991). In contrast to the deterrence perceptions, all three performed consistently. Together they were the strongest predictors in the study.

Hypothesis 6 proposed that a more favorable attitude toward non-disclosure is positively related to it. This drew on evidence that students often see disclosure as risky or self-incriminating rather than as neutral transparency, which makes concealment look like the safer choice (Gonsalves, 2025; Luo, 2025). The hypothesis was supported. Attitude was the strongest predictor across all four outcomes, with effects ranging from $\beta = .346$ to $\beta = .431$ (all $p < .001$). Students who evaluated non-disclosure more favorably were consistently more likely to conceal their use. A likely reason is that attitude captures the personal cost-benefit judgement that is at the center of the decision. The choice was driven more by the belief that concealment is sensible or safe than by any of the enforcement signals. This fits evidence that disclosure can carry real reputational costs, since revealing AI use has been shown to lower trust in the person who discloses (Schilke & Reimann, 2025). Attitude was also the main channel through which the deterrence perceptions reached behavior (Section 4.4), which underlines its central role. The takeaway is that how students judged non-disclosure was the clearest predictor of whether they engaged in it.

Hypothesis 7 proposed that subjective norms are positively related to non-disclosure. This rested on evidence that GenAI use is widely seen as routine and that non-disclosure is often treated as normal, which can make concealment socially acceptable (Goyal et al., 2026; Nguyen & Goto, 2024; Qu et al., 2025; Zhang et al., 2025). The hypothesis was supported across all four outcomes (β between .120 and .236, all $p \leq .021$). Students who believed that peers and important others would accept non-disclosure were more likely to withhold their own use. This points to a normalization mechanism, in which the perceived normality of concealment lowers its social cost (Gonsalves, 2025; Yusuf et al., 2025). The effect was

strongest for partial non-disclosure intention. Students appear to find selectively withholding details more acceptable than concealing use altogether. The takeaway is that non-disclosure was shaped not only by private judgement but also by what students thought their peers would accept.

Hypothesis 8 proposed that perceived behavioral control is positively related to non-disclosure. This built on the observation that students often feel able to conceal GenAI use and to manage what they report (Adnin et al., 2025). The hypothesis was supported for the two partial forms, but not for the full forms. Perceived control predicted past partial non-disclosure ($\beta = .129$, $p = .006$) and partial non-disclosure intention ($\beta = .160$, $p < .001$), while showing no significant effect on either full outcome. Where it held, a stronger sense of control over disclosure was associated with greater concealment. This fits the selective disclosure logic that is central to this study. Managing how much to report is a matter of selective control, which applies to partial non-disclosure more than to hiding use entirely. The result also supports studies suggesting that unclear guidance makes non-disclosure seem easier (Gonsalves, 2025; Moorhouse et al., 2023). The takeaway is that feeling able to manage disclosure encouraged concealment, but mainly in its selective forms.

Taken together, the three Theory of Planned Behavior constructs gave a consistent and comparatively strong account of non-disclosure. Attitude was foremost, subjective norms close behind, and perceived behavioral control relevant but more situational. This contrasts sharply with the weak and uneven deterrence perceptions.

5.1.4. Individual-level factors

The fifth sub-question asked: *How do academic self-efficacy and policy awareness relate to students' non-disclosure of GenAI use?* These two predictors were added to capture differences between students that the integrated deterrence and Theory of Planned Behavior model does not otherwise represent (McCabe et al., 2001; Pajares, 1996). Both were expected to reduce non-disclosure. Both did so, but mainly for the partial forms.

Hypothesis 9 proposed that academic self-efficacy, defined as students' confidence in completing assessed work without GenAI, is negatively related to non-disclosure. The reasoning was straightforward. Students who feel able to work independently rely on GenAI less and so have less to conceal. Students with lower confidence more often turn to it, and may find disclosure uncomfortable because it exposes that reliance (Adnin et al., 2025; Zhang et al., 2024). The hypothesis was supported for the two partial outcomes ($\beta = -.105$, $p = .017$; $\beta = -.081$, $p = .041$). The effects on the two full outcomes were negative but fell just short of significance ($p = .055$; $p = .071$). This link points back to the purpose of assessment. Self-efficacy reflects confidence in mastering the skills and knowledge that assessments are meant to certify. Students who feel they can meet the learning objectives on their own have little need for GenAI and little to hide. Non-disclosure is therefore tied to whether students believe they can demonstrate genuine mastery. The finding fits Krou et al. (2021) and Leaton Gray et al. (2025), who link stronger efficacy beliefs to resistance against unethical shortcuts. It

extends that work from dishonesty in general to non-disclosure specifically. One caution applies. The survey asked which tasks students used GenAI for (Table 2 and Table 3), but not how extensively they used it. As a result, the link between self-efficacy and non-disclosure may partly reflect differences in GenAI use. More confident students may have used it less and therefore had less to withhold. The takeaway is that students who felt able to succeed without GenAI were less inclined to withhold their use, above all in its selective forms.

Hypothesis 10 proposed that policy awareness is negatively related to non-disclosure. This drew on compliance research showing that people who understand the rules are better able to judge what is acceptable and more likely to follow them (Ishwardat et al., 2024; McCabe et al., 2001). The hypothesis was supported for the two partial outcomes ($\beta = -.101$, $p = .019$; $\beta = -.112$, $p = .004$), but had no effect on either full outcome. A likely explanation is that awareness reduces concealment that arises from uncertainty, not concealment that is deliberate. When students understand what a policy expects, they know how much detail to give, which makes them less likely to under-report through vagueness or omission. This fits evidence that a clearer grasp of the rules helps students recognize which practices are prohibited (Gullifer & Tyson, 2014). Hiding use entirely is a different matter. It is a deliberate choice that understanding the rules does little to change. It also fits the argument that an ambiguous and inconsistent policy landscape limits what awareness can achieve (An et al., 2025; Moorhouse et al., 2023). The takeaway is that understanding the rules reduced partial non-disclosure, but did not stop students from concealing their use altogether.

Taken together, the two individual-level predictors added explanatory value that the enforcement and psycho-social constructs did not capture. They also shared a clear pattern. Like perceived behavioral control, both were significant for the partial forms but not the full ones. Differences between students, in how capable they felt and how well they understood the rules, mattered for selective under-disclosure but not for outright concealment. This reinforces the distinction between the two behaviors that runs throughout this study.

5.1.5. Synthesis

Across the four themes, one finding stands out. Non-disclosure was shaped far more by how students evaluated concealment than by the enforcement environment. The Theory of Planned Behavior constructs explained four to nine times as much variance as the deterrence perceptions (Section 4.3). Attitude was the strongest predictor of every outcome, with subjective norms close behind. The perceived certainty, severity, and swiftness of sanctions, by contrast, were weak and uneven.

The exploratory mediation shows why, and it answers the sixth sub-question: *Which Theory of Planned Behavior constructs mediate the relationships between deterrence perceptions, policy awareness, academic self-efficacy, and non-disclosure?* Attitude was not only the strongest direct predictor but also the main channel through which the deterrence perceptions reached behavior. Certainty and severity were related to non-disclosure almost entirely through attitude, while subjective norms formed a narrower second channel. Even the enforcement cues that mattered influenced how students viewed concealment rather than

directly affecting their behavior. Severity is the clearest case, since it was absent from the direct models and became visible only once attitude was modelled as a channel. Because the design is cross-sectional, these pathways are plausible rather than proven (Bryman, 2016).

The cues that did not shape this evaluation had little effect. Swiftness ran opposite to prediction and bypassed attitude altogether. Policy awareness worked only as information, reducing partial under-reporting alone. The enforcement environment influenced non-disclosure mainly where it fed into students' evaluations.

The resulting picture is coherent. Attitude and subjective norms drove all four outcomes, while perceived control, self-efficacy, and policy awareness influenced the partial forms only. Above all, non-disclosure was an evaluative and social decision. Deterrence had a place in that decision, but only a limited and indirect one. This conclusion sets up the answer to the central research question.

5.2. Answering the Central Research Question

This section answers the central research question: *How do deterrence perceptions, TPB-based psycho-social factors, academic self-efficacy, and policy awareness explain students' non-disclosure of GenAI use in academic assessment contexts?* Integrating the sub-question findings, non-disclosure was explained mainly by psycho-social factors and individual capability rather than the enforcement environment. Attitude toward non-disclosure was the strongest predictor across all four outcomes, followed by subjective norms. Perceived behavioral control mattered as well, but only for the partial forms. The two individual-level predictors followed the same full-versus-partial pattern. Academic self-efficacy was negatively related to non-disclosure, significantly so for the partial forms and marginally for the full forms. Policy awareness was negatively related to the partial forms only. The deterrence perceptions played a smaller and largely indirect role. Certainty and severity reached non-disclosure mainly through attitude, rather than acting on behavior directly. Swiftness showed no such path, and it ran in the opposite direction to the one predicted. Attitude was therefore the pivotal construct. It served both as the strongest direct predictor and as the channel through which certainty and severity reached behavior.

5.3. Recommendations

The findings point away from an enforcement-first response. Detection tools and harsher penalties did little to reduce non-disclosure. How students evaluated disclosure mattered far more. The recommendations below follow from this. They center on shared understanding rather than stricter control. Some students voiced the same view. One argued that "the more it will be demonized the more it will be misused," calling instead for "teaching responsible use and having realistic expectations from students" (Respondent 64).

First, institutions should rely less on detection and sanctions. The deterrence perceptions were weak predictors throughout. Certainty of detection had no direct effect once attitudes were considered, and reliable detection remains difficult in practice. Students see this unreliability first-hand. One described being wrongly accused of GenAI use on work they had

written themselves, concluding that "the current systems for assessing AI usage are not trustworthy" (Respondent 157). Detection software and tougher penalties are therefore not a dependable route to more disclosure. A more credible path runs through assessment design. Formats that make independent mastery visible, such as oral examinations, raise genuine accountability without relying on detection tools.

Second, disclosure should feel safe rather than self-incriminating. Attitude was the strongest predictor and the channel through which the sanction perceptions reached behavior. The most promising measures therefore change how students judge disclosure. Many students see it as a threat to their credibility. As one put it, "If I disclose my use of AI to lecturers, I am afraid ... all of my work is automatically discredited" (Respondent 50). Institutions can ease this by being transparent about how AI-assisted work is assessed and graded. Disclosure should not appear to carry a hidden penalty.

Third, institutions should shift perceived norms toward openness. Subjective norms predicted every outcome. What students believe others accept therefore shapes their own behavior. When concealment looks like the common practice, its social cost falls and disclosure feels unnecessary. The aim is to make openness the expected norm rather than the exception. Instructors can model transparent AI use in their own teaching. They can also state plainly that disclosure is expected and will be met without suspicion. Moving the perceived norm toward openness addresses the pressure that currently makes concealment feel normal.

Fourth, guidance should be clear and proportionate, especially for partial non-disclosure. Partial non-disclosure was the most common form. Clearer rules helped, since policy awareness reduced partial under-reporting, though the effect was modest. Guidance should state plainly how much detail students must give about their GenAI use. Students asked for exactly this. One noted that "every course should just state this very clearly, and then it will all be fine" (Respondent 55). Guidance should also stay proportionate. If disclosure becomes too burdensome, students will skip it. As one respondent warned, a precise declaration "would often run to more pages than the assignment itself" (Respondent 65).

Fifth, institutions should support students' confidence to work without GenAI. Academic self-efficacy was linked to less non-disclosure. Strengthening this confidence offers a supportive route rather than a punitive one. It may lower reliance on GenAI, and thereby the need to conceal. Students asked for this too. One argued that "teaching students to use AI responsibly would be far more useful than the strange situation of rules they have now" (Respondent 379).

Taken together, these recommendations share one logic. Non-disclosure is mainly an evaluative and social decision. What works is shaping how students judge disclosure and what they see as normal. The most sustainable response is not tighter enforcement, but a shared understanding of disclosure. That means making it clearly defined, safe to offer, and a normal part of academic work.

5.4. Contributions to Theory and Practice

This study set out to explain why students withhold their use of generative AI in assessed work. In doing so, it offers several contributions to the academic literature and a number of practical implications for higher education institutions. These are discussed in turn below.

5.4.1. Contributions to Theory

The first contribution concerns how this study conceptualized non-disclosure. Most prior research treats disclosure as a binary outcome, separating students who acknowledge their GenAI use from those who do not (Adnin et al., 2025; Gonsalves, 2025; Luo, 2025; Schilke & Reimann, 2025; Yusuf et al., 2025). The present study instead distinguishes between complete non-disclosure and partial non-disclosure, in which students acknowledge some use while withholding important details. Because no existing instrument captured this distinction, the study developed and validated new self-report measures for both forms. The factor analyses supported the intended structure, and the two forms separated more clearly than past behavior and intention did (Section 3.3.2). The scales were also internally consistent. The distinction therefore proved to be more than conceptual. The two forms did not collapse into a single behavior, they differed in prevalence, and their predictors were not identical (H2). Perceived behavioral control, academic self-efficacy, and policy awareness were significant for the partial forms but not the full ones, and the role of severity differed between them. Looking at non-disclosure as just one single action would have hidden these important differences. Separating and validating the two forms is thus both a conceptual and a measurement contribution. It also suggests that future work should not assume that under-reporting and full concealment share the same drivers.

The second, and arguably central, contribution is a refined integration of Deterrence Theory and the Theory of Planned Behavior. The study responds to explicit calls to combine enforcement-based and psycho-social accounts of disclosure behavior (Gonsalves, 2025; Nguyen-Phuoc et al., 2024; Yusuf et al., 2025). Its more specific contribution is to show that the two frameworks connect only under certain conditions. Attitude reflects a cost-benefit judgment about whether concealment is worth the risk. Only the two perceptions that feed this judgment, how likely detection is (certainty) and how costly it would be (severity), reached behavior through attitude. Severity illustrates this most clearly. It showed no direct relationship with non-disclosure, yet operated fully through attitude across every outcome. An analysis of direct effects alone would have concluded that severity was irrelevant, which helps reconcile the mixed evidence reported in earlier work (Goyal et al., 2026; Nguyen & Goto, 2024; Qu et al., 2025; Zhang et al., 2025). The same pattern repositions policy awareness. Because awareness did not operate through attitude and was related only to the partial outcomes, it appears to act as a source of information rather than a driver of how students judge concealment, which qualifies the common assumption that awareness leads directly to compliance (Lund et al., 2025).

The third contribution introduces academic self-efficacy to the disclosure literature in a new form, framed as students' confidence in completing assessed work without GenAI

support. Defined this way, it was negatively related to non-disclosure, significantly so for the two partial forms and marginally for the two full forms. Its direction was consistent across all four outcomes, and it helps explain why students under the same rules still differ in whether they conceal (Krou et al., 2021; Leaton Gray et al., 2025). One caution applies. The survey recorded the types of GenAI use students reported (Table 2 and Table 3), but not how heavily they relied on it. The relationship may therefore partly reflect lower reliance among more confident students rather than a direct effect on concealment.

5.4.2. Contributions to Practice

The central practical contribution of this study is that it helps safeguard the validity of assessment. Assessments are meant to certify that students have met the intended learning objectives, and teachers can only judge this when they know how the work was produced (Biggs et al., 2022; Moorhouse et al., 2023). When students conceal their GenAI use, that judgement breaks down. By explaining why students conceal their GenAI use, the study addresses a problem that sits at the center of academic integrity. To protect the meaning of assessment, its causes must first be understood, and this study provides that understanding.

A second contribution is that the study shows institutions where their effort is most likely to matter. It weighs enforcement-based perceptions against psycho-social evaluations and individual differences. The evidence shows that the latter carry far more weight. This helps institutions judge how to use limited resources. The value is not a single measure to adopt, but a clearer sense of which kinds of measures are worth pursuing at all.

A third contribution follows from separating full and partial non-disclosure. The study shows that partial under-reporting is the more common form, and that it responds to different factors than complete concealment. This gives institutions a more precise target. Rather than treating non-disclosure as one problem, they can recognize that selective under-reporting and outright concealment may require different responses. This distinction provided a level of precision that was unavailable when disclosure was treated as a single binary act

The wider value goes beyond individual institutions. Assessment can only be trusted if it is valid, and that trust is what gives academic qualifications their worth. It also shows that graduates have the skills their degrees claim. Protecting it therefore protects the credibility of higher education itself.

5.5. Limitations and Directions for Future Research

Before turning to the limitations, it is worth noting what gives confidence in the findings. The final sample of 431 respondents exceeded the minimum of 160 indicated by the a priori power analysis, so the models were adequately powered to detect effects of the expected size. The measurement instrument also performed well. The factor analysis produced a clean structure in which the items loaded onto their intended constructs, and the reliability analysis showed acceptable to strong internal consistency across the scales. In addition, the regression assumptions of normality, homoscedasticity, independence, and the absence of problematic multicollinearity were met in all four models, so the estimates can be interpreted with reasonable confidence. Finally, the main effects were consistent rather than isolated, since attitude was related to all four non-disclosure outcomes and severity operated through attitude in every model. Effects that appear across four separate models are less likely to be a chance result than a single significant effect in one model. The limitations discussed below therefore qualify the findings rather than undermine them.

The study was cross-sectional, so all constructs were measured at a single point in time. This suits a first mapping of the proposed relationships, but it cannot establish temporal order or support strong causal claims (Bryman, 2016). This could be addressed through longitudinal or experimental designs, in which deterrence cues are measured or manipulated before disclosure behavior is observed.

The cross-sectional design also has consequences for the mediation analysis. The psychosocial constructs were modelled as mediators, yet they were measured at the same moment as the outcomes. This is problematic for the two past-behavior outcomes, because the behavior preceded the measurement of the mediators. The assumed order, in which these constructs shape behavior, may therefore be reversed, since earlier concealment could equally have shaped students' current attitudes and perceptions. Longitudinal mediation, with the predictor, the mediator, and the outcome measured at separate moments, would establish that the mediators come before the behavior they are meant to explain.

The sample places further limits on the findings. Recruitment relied on non-probability sampling, which restricts how far the results generalize to the wider student population (Golzar et al., 2022). The composition of the sample was also uneven, as it leaned toward certain study levels and fields of study and contained a modest majority of women, as described in Section 3.2. Because disclosure norms and pressures may differ across these groups, the imbalance may have shaped the estimates. Future research could recruit through probability-based sampling across several institutions and balance the sample more deliberately across study levels and fields. Treating these characteristics as moderators would also show whether the drivers of non-disclosure hold across different student groups.

A further limitation concerns how GenAI use was measured. The survey recorded which types of use students engaged in (Table 2 and Table 3), but it did not capture how heavily they relied on these tools. This is relevant, because non-disclosure is closely related to actual use. What students withhold is best judged against what they had to disclose in the first place. The

same gap affects the interpretation of academic self-efficacy. The negative relationship between self-efficacy and non-disclosure may partly reflect that more confident students used GenAI less, rather than a direct effect of confidence on concealment. This study did not include a control for the extent of each student's use. Future research could add focused control variables that capture how much each student relied on GenAI. It could also measure the intensity of use more directly. Disclosure could then be modelled relative to that use. This would help separate two explanations: whether higher self-efficacy leads to less GenAI use and so to less concealment, or whether it reduces non-disclosure directly.

A further set of limitations stems from the reliance on self-report. Because the behavior is sensitive and sits close to academic misconduct, some under-reporting of concealment is likely, even though the anonymous and self-administered format was chosen to reduce exactly this bias (Tourangeau & Yan, 2007). Self-assessment is a particularly concerning factor for academic self-efficacy, since people tend to rate their own capabilities more favorably than their actual performance. In addition, all predictors, the mediators, and the outcomes were drawn from the same instrument, the same respondents, and the same moment, which can inflate the observed associations through common method variance. Future research could address this using several changes. Indirect techniques, such as list experiments or the randomized response method, can lower social desirability pressure. Objective indicators of performance can be used to validate self-efficacy. Finally, multi-source data or a marker variable can guard against common method bias.

A related point is that deterrence measures and policy awareness were measured as perceptions rather than as objective conditions. This was a deliberate choice, since perceptions are what the theory concerns. However, this means that the study cannot say whether the actual detection environment or the actual content of institutional policy drives behavior. Future research could combine students' perceptions with objective information. This would help determine whether behavior is driven by students' perceptions or by the actual institutional conditions. A multilevel design, in which students are grouped within courses or institutions, would be particularly suitable for examining this.

Two of the findings themselves also invite further study. The first is the swiftness reversal. Swiftness was related to a stronger intention to conceal rather than a weaker one, opposite to the deterrence prediction. Section 5.1.2 set out several possible explanations. These included a sense of procedural threat, a stronger incentive to hide use completely, and the narrow two-item measure left after the factor analysis. The cross-sectional data could not decide between them. The second finding is that swiftness and policy awareness did not operate through the psycho-social constructs at all. Where they mattered, their effect was direct. This suggests that the mechanisms linking them to non-disclosure lie outside the mediators modelled here. The models also left a substantial share of variance unexplained, which points in the same direction. Future research could test the swiftness explanations directly, for example by experimentally varying the speed of enforcement and observing how students respond. It could also examine additional predictors or mediators, such as trust in the assessment system,

perceived fairness, or moral disengagement, which may help explain the effects that the psycho-social constructs did not carry (Luo, 2025; Schilke & Reimann, 2025).

Beyond addressing these limitations, the study opens several substantive directions for future research. Some were raised by respondents themselves in the open comments.

A first direction concerns the situational pressures behind non-disclosure. This study modelled stable perceptions and traits, but some students pointed to more immediate triggers. Several described stress, deadlines, and competition as reasons for turning to GenAI and for hiding it. One noted that seeing peers use it pushes others to use it too, "to keep up with the competition" (Respondent 529). Future work could examine how such situational pressures shape disclosure, for example through diary or experience-sampling methods.

A second direction is the gap between how students genuinely view GenAI use and how they act under the threat of sanctions. One respondent suggested studying how students would use AI "if the students had no fear of consequences," so as to reveal how they actually feel rather than how they behave out of caution (Respondent 488). Designs that lower this pressure, or that compare high- and low-enforcement settings, could separate genuine evaluation from fear-driven behavior.

A third direction links non-disclosure to learning itself. This study treated non-disclosure as the outcome. Students, however, raised a further question. They asked what their GenAI use meant for their understanding and their grades (Respondent 83). Future research could connect the extent of GenAI use, its disclosure, and actual learning outcomes. This would speak directly to the assessment-validity concern at the heart of the topic.

A fourth direction concerns culture and context. Disclosure norms may differ across countries, institutions, and languages, and this sample came largely from one national context. Comparative and cross-cultural research would show whether the drivers identified here hold in other settings.

Finally, the institutional side deserves direct study. This study focused on students. Yet non-disclosure is also shaped by universities. How they define, communicate, and sanction disclosure all matter. Research could examine these institutional practices. It could also look at how teachers interpret and apply them. This would complement the student-level picture developed here.

References

- Adnin, R., Pandkar, A., Yao, B., Wang, D., & Das, M. (2025). *Examining student and teacher perspectives on undisclosed use of generative AI in academic work*. <https://doi.org/10.1145/3706598.3713393>
- Ajzen, I. (1991). The theory of planned behavior. *Organizational Behavior and Human Decision Processes*, 50, 179-211. [https://doi.org/10.1016/0749-5978\(91\)90020-T](https://doi.org/10.1016/0749-5978(91)90020-T)
- Ajzen, I. (2006). *Constructing a theory of planned behavior questionnaire*. University of Massachusetts Amherst.
- An, Y., Yu, J. H., & James, S. (2025). Investigating the higher education institutions' guidelines and policies regarding the use of generative AI in teaching, learning, research, and administration. *International Journal of Educational Technology in Higher Education*, 22(1), 10. <https://doi.org/10.1186/s41239-025-00507-3>
- Ardito, C. G. (2025). Generative AI detection in higher education assessments. *New Directions for Teaching and Learning*, 2025(182), 11-28. <https://doi.org/10.1002/tl.20624>
- Barrett, A., & Pack, A. (2023). Not quite eye to A.I.: student and teacher perspectives on the use of generative artificial intelligence in the writing process. *International Journal of Educational Technology in Higher Education*, 20(1), 59. <https://doi.org/10.1186/s41239-023-00427-0>
- Biggs, J., Tang, C., & Kennedy, G. (2022). *Teaching for quality learning at university 5e*. McGraw-hill education (UK).
- Bittle, K., & El-Gayar, O. (2025). Generative AI and academic integrity in higher education: a systematic review and research agenda. *Information*, 16(4), 296.
- Boateng, G. O., Neilands, T. B., Frongillo, E. A., Melgar-Quinonez, H. R., & Young, S. L. (2018). Best practices for developing and validating scales for health, social, and behavioral research: a primer [Review]. *Frontiers in Public Health*, Volume 6 - 2018. <https://www.frontiersin.org/journals/public-health/articles/10.3389/fpubh.2018.00149>
- Bos, N., Lammers, K.-J., & Prince van Leeuwen, A. (2025). Exploring the role of AI in higher education: a study of usage by students and teachers in the Netherlands. *Review of Artificial Intelligence in Education*, 6, e040. <https://doi.org/10.37497/rev.artif.intell.educ.v6ii.40>
- Bryman, A. (2016). *Social research methods* (Fifth edition ed.). Oxford University Press.
- Brysbaert, M., Sui, L., Duyck, W., & Dirix, N. (2021). Improving reading rate prediction with word length information: evidence from Dutch. *Q J Exp Psychol (Hove)*, 74(11), 2013-2018. <https://doi.org/10.1177/17470218211017100>
- Bulgurcu, B., Cavusoglu, H., & Benbasat, I. (2010). Information security policy compliance: an empirical study of rationality-based beliefs and information security awareness. *Management Information Systems Quarterly*, 34(3), 523-548. <https://doi.org/10.2307/25750690>
- Capili, B. (2021). Cross-sectional studies. *The American journal of nursing*, 121, 59-62. <https://doi.org/10.1097/01.NAJ.0000794280.73744.fe>
- Centraal Bureau voor de Statistiek (CBS). (2023). More women than men in higher education for 23 consecutive years. <https://www.cbs.nl/en-gb/news/2023/10/more-women-than-men-in-higher-education-for-23-consecutive-years>
- Centraal Bureau voor de Statistiek (CBS). (2026). *Hoger onderwijs; ingeschrevenen, onderwijssoort, opleidingsfase en -vorm* (CBS StatLine. <https://www.cbs.nl/nl-nl/cijfers/detail/85423NED>

- Coates, H., Croucher, G., & Calderon, A. (2025). Governing academic integrity: ensuring the authenticity of higher thinking in the era of generative artificial intelligence. *Journal of Academic Ethics*, 23(4), 2015-2028. <https://doi.org/10.1007/s10805-025-09639-7>
- Cohen, J. (1988). *Statistical power analysis for the behavioral sciences (2nd ed.)*. Routledge. <https://doi.org/10.4324/9780203771587>
- Conrad, F. G., Couper, M. P., Tourangeau, R., & Zhang, C. (2017). Reducing speeding in web surveys by providing immediate feedback. *Survey Research Methods*, 11(1), 45-61. <https://doi.org/10.18148/srm/2017.v11i1.6304>
- Costello, A. B., & Osborne, J. (2005). Best practices in exploratory factor analysis: four recommendations for getting the most from your analysis. *Practical Assessment, Research & Evaluation*, 10, 1-9.
- D'Arcy, J., Hovav, A., & Galletta, D. (2009). User awareness of security countermeasures and its impact on information systems misuse: a deterrence approach. *Information Systems Research*, 20, 79-98. <https://doi.org/10.1287/isre.1070.0160>
- Deep, P. D., Edgington, W. D., Ghosh, N., & Rahaman, M. S. (2025). Evaluating the effectiveness and ethical implications of AI detection tools in higher education. *Information*, 16(10), 905.
- den Hollander, N., Struyf, A., Georgiou, D., & Wong, J. (2025). From perception to policy: course coordinators' views on generative artificial intelligence in higher education. *International Journal of Educational Management*, 1-16. <https://doi.org/10.1108/IJEM-12-2024-0842>
- Dench, D. L., & Joyce, T. J. (2022). Information and credible sanctions in curbing online cheating among undergraduates: a field experiment. *National Bureau of Economic Research Working Paper Series*, No. 29755. <https://doi.org/10.3386/w29755>
- Digital Education Council. (2026). *AI in Higher Education Global Survey 2026*. <https://www.digitaleducationcouncil.com/resource-library-items/ai-in-higher-education-global-survey-2026>
- Dunn, A., Heggstad, E., Shanock, L., & Theilgard, N. (2018). Intra-individual response variability as an indicator of insufficient effort responding: comparison to other indicators and relationships with individual differences. *Journal of Business and Psychology*, 33. <https://doi.org/10.1007/s10869-016-9479-0>
- Eisinga, R., Grotenhuis, M. t., & Pelzer, B. (2013). The reliability of a two-item scale: Pearson, Cronbach, or Spearman-Brown? *International Journal of Public Health*, 58(4), 637-642. <https://doi.org/10.1007/s00038-012-0416-3>
- Elangovan, N., & Sundaravel, E. (2021). Method of preparing a document for survey instrument validation by experts. *MethodsX*, 8, 101326. <https://doi.org/10.1016/j.mex.2021.101326>
- Elo, S., & Kyngäs, H. (2008). The qualitative content analysis process. *Journal of Advanced Nursing*, 62(1), 107-115. <https://doi.org/10.1111/j.1365-2648.2007.04569.x>
- Erdfelder, E., Faul, F., & Buchner, A. (1996). GPOWER: a general power analysis program. *Behavior Research Methods, Instruments, & Computers*, 28(1), 1-11. <https://doi.org/10.3758/BF03203630>
- Fabrigar, L., Wegener, D., MacCallum, R. C., & Strahan, E. J. (1999). Evaluating the use of Exploratory Factor Analysis in psychological research. *Psychological Methods*, 4, 272-299. <https://doi.org/10.1037//1082-989X.4.3.272>
- Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, 39(2), 175-191. <https://doi.org/10.3758/BF03193146>

- Field, A. (2018). *Discovering statistics using IBM SPSS statistics* (Fifth edition ed.). SAGE Publications.
- Finch, W. H. (2012). Distribution of variables by method of outlier detection. *Frontiers in Psychology*, 3, 211. <https://doi.org/10.3389/fpsyg.2012.00211>
- Fleckenstein, J., Meyer, J., Jansen, T., Keller, S. D., Köller, O., & Möller, J. (2024). Do teachers spot AI? Evaluating the detectability of AI-generated texts among student essays. *Computers and Education: Artificial Intelligence*, 6, 100209. <https://doi.org/10.1016/j.caeai.2024.100209>
- Freeman, J. (2025). *Student generative AI survey 2025*. <https://www.hepi.ac.uk/wp-content/uploads/2025/02/HEPI-Kortext-Student-Generative-AI-Survey-2025.pdf>
- Freeman, J., Liossis, P., & David, N. (2006). Deterrence, defiance and deviance: an investigation into a group of recidivist drink drivers' self-reported offending behaviours. *The Australian and New Zealand Journal of Criminology*, 39(1), 1-19.
- George, D., & Mallery, P. (2019). *IBM SPSS statistics 26 step by step: a simple guide and reference*. Routledge. <https://doi.org/10.4324/9780429056765>
- Gibbs, J. P. (1979). Assessing the deterrence doctrine a challenge for the social and behavioral. *The American Behavioral Scientist* (pre-1986), 22(6), 653. <http://tudelft.on.worldcat.org/atoztitles/link?sid=ProQ:&issn=00027642&volume=22&issue=6&title=The+American+Behavioral+Scientist+%28pre-1986%29&spage=653&date=1979-07-01&title=Assessing+the+Deterrence+Doctrine+A+Challenge+for+the+Social+and+Behavioral+Sciences&au=Gibbs%2C+Jack+P&id=doi:>
- Gnambs, T., & Kaspar, K. (2015). Disclosure of sensitive behaviors across self-administered survey modes: a meta-analysis. *Behavior Research Methods*, 47(4), 1237-1259. <https://doi.org/10.3758/s13428-014-0533-4>
- Golzar, J., Tajik, O., & Noor, S. (2022). Convenience sampling. 1, 72-77. <https://doi.org/10.22034/ijels.2022.162981>
- Gonsalves, C. (2025). Addressing student non-compliance in AI use declarations: implications for academic integrity and assessment in higher education. *Assessment & Evaluation in Higher Education*, 50(4), 592-606. <https://doi.org/10.1080/02602938.2024.2415654>
- Goyal, S., Chauhan, S., & Motiwalla, L. (2026). Examining business students' intentions to misuse ChatGPT through the lens of deterrence and neutralisation theories. *Behaviour & Information Technology*, 45(4), 660-679. <https://doi.org/10.1080/0144929X.2025.2525307>
- Greitemeyer, T., & Kastenmüller, A. (2024). *A longitudinal analysis of the willingness to use chatGPT for academic cheating: Applying the theory of planned behavior* [doi:10.1037/tmb0000133]. American Psychological Association.
- Grigoryeva, M. S., & Robbins, B. G. (2024). Every forest has its shadow: the demographics of concealment in the united states. *Sociological Science*, 11(13), 340-378. <https://doi.org/10.15195/v11.a13>
- Gullifer, J. M., & Tyson, G. A. (2014). Who has read the policy on plagiarism? Unpacking students' understanding of plagiarism. *Studies in Higher Education*, 39(7), 1202-1218. <https://doi.org/10.1080/03075079.2013.777412>
- Hadra, M., Cambridge, K., & Mesbah, M. (2026). Evaluating the accuracy and reliability of AI content detectors in academic contexts. *International Journal for Educational Integrity*, 22(1), 4. <https://doi.org/10.1007/s40979-026-00213-1>

- Hair, J. F., Black, W. C., Babin, B. J., & Anderson, R. E. (2014). *Multivariate data analysis* (Seventh edition, Pearson new international edition ed.). Pearson Education Limited.
- Hashim, S., Mohamad, S. F., Lim, S., & Che Ahmat, N. H. (2022). Pretesting survey questionnaire: a guide on dissemination. *International Journal of Academic Research in Economics and Management Sciences*, 11. <https://doi.org/10.6007/IJAREMS/v11-i3/15228>
- Hayes, A. F. (2013). *Introduction to mediation, moderation, and conditional process analysis : a regression-based approach*. The Guilford Press.
- Howard, M. C. (2023). A systematic literature review of exploratory factor analyses in management. *Journal of Business Research*, 164. <https://doi.org/10.1016/j.jbusres.2023.113969>
- Huang, D., Hash, N., Cummings, J. J., & Prena, K. (2025). Academic cheating with generative AI: Exploring a moral extension of the theory of planned behavior. *Computers and Education: Artificial Intelligence*, 8, 100424. <https://doi.org/10.1016/j.caeai.2025.100424>
- Huang, J. L., Lui, M., & Bowling, N. A. (2015). Insufficient effort responding: Examining an insidious confound in survey data [doi:10.1037/a0038510]. 100, 828-845. <https://doi.org/10.1037/a0038510>
- Hutson, J. (2024). Rethinking plagiarism in the era of generative AI. *Journal of Intelligent Communication*, 3(2), 20–31. <https://doi.org/10.54963/jic.v3i2.220>
- Ishwardat, S., van Steenbergen, E., Coffeng, T., & Ellemers, N. (2024). Stimulating regulatory compliance and ethical behavior of organizations: a review. *Business Ethics and Leadership*, 8, 151-172. [https://doi.org/10.61093/bel.8\(3\).151-172.2024](https://doi.org/10.61093/bel.8(3).151-172.2024)
- Jantos, A., Brendel, A., & Kilz, L. (2024). Online vs. Face-to-face exams - deterrence theory explaining cheating behaviour in summative written exams. <https://doi.org/10.21125/inted.2024.1739>
- Krou, M. R., Fong, C. J., & Hoff, M. A. (2021). Achievement motivation and academic dishonesty: a meta-analytic investigation. *Educational Psychology Review*, 33(2), 427-458. <https://doi.org/10.1007/s10648-020-09557-7>
- Kumar, R., & Mindzak, M. (2024). Who Wrote This?: Detecting Artificial Intelligence–Generated Text from Human-Written Text. *Canadian Perspectives on Academic Integrity*, 7. <https://doi.org/10.55016/ojs/cpai.v7i1.77675>
- Kusmaryono, I., Wijayanti, D., & Maharani, H. R. (2022). Number of response options, reliability, validity, and potential bias in the use of the likert scale education and social science research: a literature review. *International Journal of Educational Methodology*, 8(4), 625-637. <https://doi.org/10.12973/ijem.8.4.625>
- Kwak, S. K., & Kim, J. H. (2017). Statistical data preparation: management of missing values and outliers. *Korean Journal of Anesthesiology*, 70(4), 407-411. <https://doi.org/10.4097/kjae.2017.70.4.407>
- Larson, D. G., & Chastain, R. L. (1990). Self-concealment: Conceptualization, measurement, and health implications [doi:10.1521/jscp.1990.9.4.439]. *Journal of Social and Clinical Psychology*, 9(4), 439-455. <https://doi.org/10.1521/jscp.1990.9.4.439>
- Leaton Gray, S., Edsall, D., & Parapadakis, D. (2025). AI-based digital cheating at university, and the case for new ethical pedagogies. *Journal of Academic Ethics*, 23(4), 2069-2086. <https://doi.org/10.1007/s10805-025-09642-y>
- Lou, L., Jiao, Y., Koh, J., & Dai, W. (2025). Exploring the drivers of content entrepreneurs' compliance with generative AI policies: a mixed-methods approach. *Journal of Theoretical and Applied Electronic Commerce Research*, 20(4), 284.

- Lund, B., Mannuru, N. R., Teel, Z. A., Lee, T. H., Ortega, N. J., Simmons, S., & Ward, E. (2025). Student perceptions of AI-assisted writing and academic integrity: ethical concerns, academic misconduct, and use of generative AI in higher education. *AI in Education*, 1(1), 2.
- Luo, J. (2024). A critical review of GenAI policies in higher education assessment: a call to reconsider the “originality” of students’ work. *Assessment & Evaluation in Higher Education*, 49(5), 651-664. <https://doi.org/10.1080/02602938.2024.2309963>
- Luo, J. (2025). How does GenAI affect trust in teacher-student relationships? Insights from students’ assessment experiences. *Teaching in Higher Education*, 30(4), 991-1006. <https://doi.org/10.1080/13562517.2024.2341005>
- Malhotra, N. (2008). Completion time and response order effects in web surveys. *The Public Opinion Quarterly*, 72(5), 914-934. <http://www.jstor.org.tudelft.idm.oclc.org/stable/25548051>
- McCabe, D., Trevino, L., & Butterfield, K. (2001). Cheating in academic institutions: a decade of research. *Ethics & Behavior*, 11. https://doi.org/10.1207/S15327019EB1103_2
- McCabe, D. L., Feghali, T., & Abdallah, H. (2008). Academic dishonesty in the middle east: individual and contextual factors. *Research in Higher Education*, 49(5), 451-467. <https://doi.org/10.1007/s11162-008-9092-9>
- Moorhouse, B. L., Yeo, M. A., & Wan, Y. (2023). Generative AI tools and assessment: guidelines of the world's top-ranking universities. *Computers and Education Open*, 5, 100151. <https://doi.org/10.1016/j.caeo.2023.100151>
- Nagin, D. S., & Pogarsky, G. (2004). Time and punishment: delayed consequences and criminal behavior. *Journal of Quantitative Criminology*, 20(4), 295-317. <https://doi.org/10.1007/s10940-004-5866-1>
- Nguyen-Phuoc, D. Q., Xuan Mai, N., Kim, I., & Oviedo-Trespalacios, O. (2024). Questioning penalties and road safety Policies: Are they enough to deter risky motorcyclist Behavior? *Accident Analysis & Prevention*, 207, 107756. <https://doi.org/10.1016/j.aap.2024.107756>
- Nguyen, H. M., & Goto, D. (2024). Unmasking academic cheating behavior in the artificial intelligence era: Evidence from Vietnamese undergraduates. *Education and Information Technologies*, 29(12), 15999-16025. <https://doi.org/10.1007/s10639-024-12495-4>
- Nielsen, T., Dammeyer, J., Vang, M. L., & Makransky, G. (2018). Gender fairness in self-efficacy? A rasch-based validity study of the general academic self-efficacy scale (GASE). *Scandinavian Journal of Educational Research*, 62(5), 664-681. <https://doi.org/10.1080/00313831.2017.1306796>
- Nunnally, J. C. (1978). An overview of psychological measurement. In B. B. Wolman (Ed.), *Clinical Diagnosis of Mental Disorders: A Handbook* (pp. 97-146). Springer US. https://doi.org/10.1007/978-1-4684-2490-4_4
- O'Cathain, A., & Thomas, K. J. (2004). "Any other comments?" Open questions on questionnaires - a bane or a bonus to research? *BMC Medical Research Methodology*, 4(25). <https://doi.org/10.1186/1471-2288-4-25>
- Organisation for Economic Co-operation Development (OECD). (2025). *Education GPS: Netherlands - overview of the education system (Education at a Glance 2025)*. <https://gpseducation.oecd.org/CountryProfile?primaryCountry=NLD&treshold=5&topic=EO>
- Pajares, F. (1996). Self-efficacy beliefs in academic settings. *Review of Educational Research*, 66(4), 543-578.

- Perkins, M., Roe, J., Vu, B. H., Postma, D., Hickerson, D., McGaughan, J., & Khuat, H. Q. (2024). Simple techniques to bypass GenAI text detectors: implications for inclusive education. *International Journal of Educational Technology in Higher Education*, 21(1), 53. <https://doi.org/10.1186/s41239-024-00487-w>
- Peytcheva, E., & Yan, T. (2025). *A practical guide to survey questionnaire design and evaluation* (1 ed.). Routledge. <https://doi-org.tudelft.idm.oclc.org/10.4324/9781003367826>
- Plonsky, L., & Ghanbar, H. (2018). Multiple regression in L2 research: a methodological synthesis and guide to interpreting R2 values. *The Modern Language Journal*, 102(4), 713-731. <https://doi.org/10.1111/modl.12509>
- Qu, Y., Loo, H. E., & Wang, J. (2025). Generative artificial intelligence in higher education: Emotional tensions and ethical declaration. *BRITISH JOURNAL OF EDUCATIONAL TECHNOLOGY*. <https://doi.org/10.1111/bjet.70029>
- Ravšelj, D., Keržič, D., Tomaževič, N., Umek, L., Brezovar, N., A. Iahad, N., Abdulla, A. A., Akopyan, A., Aldana Segura, M. W., AlHumaid, J., Allam, M. F., Alló, M., Andoh, R. P. K., Andronic, O., Arthur, Y. D., Aydın, F., Badran, A., Balbontín-Alvarado, R., Ben Saad, H.,...Aristovnik, A. (2025). Higher education students' perceptions of ChatGPT: A global study of early reactions. *PLOS ONE*, 20(2), e0315011. <https://doi.org/10.1371/journal.pone.0315011>
- Rayner, K., Schotter, E. R., Masson, M. E. J., Potter, M. C., & Treiman, R. (2016). So much to read, so little time: how do we read, and can speed reading help? *Psychological Science in the Public Interest*, 17(1), 4-34. <https://doi.org/10.1177/1529100615623267>
- Samonte, B. M. V., & Acedera, A. P. (2025). Artificial intelligence literacy, awareness of institutional policies and self-regulation skills as predictors of students' adherence to academic integrity. *International Journal of All Research Writings*, 6(11), 245 - 254. <https://ijarw.com/PublishedPaper/IJARW2557.pdf>
- Schilke, O., & Reimann, M. (2025). The transparency dilemma: How AI disclosure erodes trust. *Organizational Behavior and Human Decision Processes*, 188, 104405. <https://doi.org/10.1016/j.obhdp.2025.104405>
- Shrestha, N. (2021). Factor analysis as a tool for survey analysis. *American Journal of Applied Mathematics and Statistics*, 9, 4-11. <https://doi.org/10.12691/ajams-9-1-2>
- Smiti, A. (2020). A critical overview of outlier detection methods. *Computer Science Review*, 38, 100306. <https://doi.org/10.1016/j.cosrev.2020.100306>
- Spiegel, M. R., & Stephens, L. J. (2013). *Schaum's outline of statistics* (5th ed.). McGraw-Hill.
- Stephenson, R., & Armstrong, C. (2026). *Student generative AI survey 2026* (HEPI Report, Issue. Oxford. <https://www.hepi.ac.uk/wp-content/uploads/2026/03/HEPI-Report-199-Gen-AI-Survey-2026.pdf>
- Stevens, J. (2002). *Applied multivariate statistics for the social sciences* (4th ed.). L. Erlbaum.
- Tan, C., Alhammad, M., & Stelmaszak, M. (2024). *Teachers' perceptions of students' use of generative AI in summative assessments at higher education institutions: an exploratory study*.
- Ting, H., Memon, M., Ramayah, T., & Cheah, J.-H. (2025). Snowball sampling: a review and guidelines for survey research. *Asian Journal of Business Research*, 15, 1-15. <https://doi.org/10.14707/ajbr.250186>
- Tourangeau, R., & Yan, T. (2007). Sensitive questions in surveys. *Psychological Bulletin*, 133, 859-883. <https://doi.org/10.1037/0033-2909.133.5.859>

- Wang, H., Bah, M. J., & Hammad, M. (2019). Progress in outlier detection techniques: A survey. *IEEE Access*, 7, 107964-108000.
<https://doi.org/10.1109/ACCESS.2019.2932769>
- Weaver, K. D. (2024). The artificial intelligence disclosure (AID) framework: an introduction. *arXiv preprint arXiv:2408.01904*.
- Weber-Wulff, D., Anohina-Naumeca, A., Bjelobaba, S., Foltýnek, T., Guerrero-Dib, J., Popoola, O., Šigut, P., & Waddington, L. (2023). Testing of detection tools for AI-generated text. *International Journal for Educational Integrity*, 19(1), 26.
<https://doi.org/10.1007/s40979-023-00146-z>
- Worthington, R., & Whittaker, T. (2006). Scale development research. *Counseling Psychologist*, 34, 806-838. <https://doi.org/10.1177/0011000006288127>
- Yusuf, A., Mouas, S., & Al-khresheh, M. H. (2025). Understanding the psychological factors associated with (non) disclosure behaviour after genAI usage. *Journal of Academic Ethics*, 24(1), 34. <https://doi.org/10.1007/s10805-025-09712-1>
- Zellner, A. (1962). An Efficient Method of Estimating Seemingly Unrelated Regressions and Tests for Aggregation Bias. *Journal of the American Statistical Association*, 57(298), 348-368. <https://doi.org/10.2307/2281644>
- Zhang, C., & Conrad, F. (2014). Speeding in Web Surveys: The tendency to answer very fast and its association with straightlining. *Survey Research Methods*, 8, 127-135.
<https://doi.org/10.18148/srm/2014.v8i2.5453>
- Zhang, L., Amos, C., & Pentina, I. (2025). Interplay of rationality and morality in using ChatGPT for academic misconduct. *Behaviour & Information Technology*, 44(3), 491-507. <https://doi.org/10.1080/0144929X.2024.2325023>
- Zhang, S., Zhao, X., Zhou, T., & Kim, J. H. (2024). Do you have AI dependency? The roles of academic self-efficacy, academic stress, and performance expectations on problematic AI usage behavior. *International Journal of Educational Technology in Higher Education*, 21(1), 34. <https://doi.org/10.1186/s41239-024-00467-0>
- Zijlstra, W. P., van der Ark, L. A., & Sijtsma, K. (2010). Outliers in questionnaire data. *Journal of Educational and Behavioral Statistics*, 36(2), 186-212.
<https://doi.org/10.3102/1076998610366263>

Appendix 1 – Survey Questions

The survey questions for this study measure the main constructs included in the conceptual model: perceived certainty of detection, perceived severity of sanctions, perceived swiftness of sanctions, attitudes toward non-disclosure, subjective norms, perceived behavioral control, academic self-efficacy to complete assessed academic work without GenAI assistance, policy awareness, and students' non-disclosure of GenAI use in assessed academic work.

As a condition for participation, all respondents are asked to provide informed consent (in the form of an opening statement) for their responses to be used for academic research. In order to reduce confusion regarding the terms used in this survey, a short instruction block will provide a clear explanation. The instruction block will show the distinction between partial and full non-disclosure behavior thereby reducing interpretation noise. Table 12 provides an overview of the survey questions and their corresponding constructs, and Table 13 provides an overview of the demographic questions included in the survey.

Table 12 – Survey questions

<i>Construct Variable</i>	Previous Full Non-Disclosure Dependent Variable	
<i>Header</i>	During your last year of studies, how often have you done the following in assessed work?	
<i>5-Point Frequency Scale</i>	Please indicate how often each of the following statements applied to you during this academic year, on a scale from 1 to 5: 1 = Never, 5 = Always.	
<i>Item</i>	<i>Statement</i>	<i>Source</i>
PFD1	I did not disclose any of my GenAI use when I was expected to.	
PFD2	I deliberately left out any mention of having used GenAI in my work.	-
<i>Construct Variable</i>	Previous Partial Non-Disclosure Dependent Variable	
<i>Header</i>	During your last year of studies, how often have you done the following in assessed work?	
<i>5-Point Frequency Scale</i>	Please indicate how often each of the following statements applied to you during this academic year, a scale from 1 to 5: 1 = Never, 5 = Always.	
<i>Item</i>	<i>Statement</i>	<i>Source</i>
PPD1	I disclosed some of my GenAI use but not all of it.	
PPD2	I left out details about how I used GenAI.	
PPD3	I made my GenAI use seem less extensive than it actually was.	
PPD4	I did not make clear how much my final work depended on GenAI output.	-
PPD5	I only reported minor uses of GenAI but left out more substantial contributions.	
PPD6	I described my GenAI use in general terms rather than in specific detail.	

<i>Construct Variable</i>	Full Non-Disclosure Intention Dependent Variable	
<i>Header</i>	For your next round of assessments, how likely are you to do the following?	
<i>5-Point Likert Scale</i>	Please rate how likely you are to do each of the following on a scale from 1 to 5: 1 = Very unlikely, 5 = Very likely.	
<i>Item</i>	<i>Statement</i>	<i>Source</i>
FDI1	I will keep all of my GenAI use undisclosed when disclosure is expected.	-
FDI2	I will deliberately leave out any mention of having used GenAI in my work.	
<i>Construct Variable</i>	Partial Non-Disclosure Intention Dependent Variable	
<i>Header</i>	For your next round of assessments, how likely are you to do the following?	
<i>5-Point Likert Scale</i>	Please rate how likely you are to do each of the following on a scale from 1 to 5: 1 = Very unlikely, 5 = Very likely.	
<i>Item</i>	<i>Statement</i>	<i>Source</i>
PDI1	I will disclose some of my GenAI use but not all of it.	-
PDI2	I will leave out details about how I used GenAI.	
PDI3	I will make my GenAI use seem less extensive than it actually was	
PDI4	I will give limited information about how much my final work depended on GenAI output.	
PDI5	I will only report minor uses of GenAI but leave out more substantial contributions.	
PDI6	I will describe my GenAI use in general terms rather than in specific detail.	
<i>Construct Variable</i>	Certainty of Sanctions Independent Variable	
<i>5-Point Likert Scale</i>	Please indicate how much you agree with the following statements. These statements are about using GenAI in assessed work without disclosing the extent of your use. Measured on a scale of 1 to 5: 1 = strongly disagree, 5 = strongly agree.	
<i>Item</i>	<i>Statement</i>	<i>Source</i>
CER1	It is likely that undisclosed GenAI use will be detected.	(Nguyen-Phuoc et al., 2024)
CER2	My examiner would detect the extent of my GenAI use.	
CER3	AI detection tools would detect the extent of my GenAI use.	
CER4	An oral defense or follow-up interview would reveal the extent of my GenAI use.	
<i>Construct Variable</i>	Severity of Sanctions Independent Variable	

5-Point Likert Scale	Please imagine that you used GenAI in assessed work but did not disclose the extent of that use. The following statements are about what you think would happen if this was discovered Measured on a scale of 1 to 5: 1 = strongly disagree, 5 = strongly agree.	
Item	Statement	Source
SER1	If this was discovered, the formal consequences would be serious.	(Jantos et al., 2024)
SER2	If this was discovered, my academic progress would be hindered.	
SER3	If this was discovered, I would have to go through a difficult disciplinary process.	
SER4	If this was discovered, my reputation with lecturers would be harmed.	
SER5	If this was discovered, people whose opinion I value would lose trust in me.	
Construct Variable	Swiftness of Sanctions Independent Variable	
5-Point Likert Scale	Please imagine that you used GenAI in assessed work but did not fully report how you used it. The following statements are about what you think would happen if this was discovered Measured on a scale of 1 to 5: 1 = strongly disagree, 5 = strongly agree.	
Item	Statement	Source
SOS1	If this was discovered, the university would immediately take action.	(Freeman et al., 2006)
SOS2	If this was discovered, a final decision would be reached in a very short time.	
SOS3	If this was discovered, it would not take long for me to receive a sanction.	
Construct Variable	Attitudes Independent Variable	
5-Point Semantic Differential Scale	Please indicate how you personally evaluate the following behavior. For me, not disclosing the extent of my GenAI use in assessed work would be ...	
Item	Scale	Source
ATT1	Bad – Good	(Greitemeyer & Kastenmüller, 2024)
ATT2	Foolish – Wise	
ATT3	Useless – Useful	
ATT4	Harmful – Beneficial	
Construct Variable	Subjective Norms Independent Variable	
5-Point Likert Scale	Please indicate how much you agree with the following statements. These statements are about not disclosing the full extent of your GenAI use in assessed work. Measured on a scale of 1 to 5: 1 = strongly disagree, 5 = strongly agree.	
Item	Statement	Source

SUN1	Most people who are important to me would find this acceptable.	
SUN2	Most people who are important to me would understand this behavior.	(Greitemeyer & Kastenmüller, 2024)
SUN3	My peers would find this acceptable.	
SUN4	My peers would understand this behavior.	
Construct Variable	Perceived Behavioral Control Independent Variable	
5-Point Likert Scale	Please indicate how much you agree with the following statements. These statements are about not disclosing the full extent of your GenAI use in assessed work. Measured on a scale of 1 to 5: 1 = strongly disagree, 5 = strongly agree.	
Item	Statement	Source
PBC1	I can easily report my GenAI use less extensively than required.	
PBC2	I am confident I can control what I report about my GenAI use.	(Nguyen-Phuoc et al., 2024)
PBC3	It is entirely up to me how much I disclose about my GenAI use.	
Construct Variable	Policy awareness Independent Variable	
5-Point Likert Scale	Please indicate how much you agree with the following statements. These statements are about the rules for using GenAI in assessed work. Measured on a scale of 1 to 5: 1 = strongly disagree, 5 = strongly agree.	
Item	Statement	Source
PA1	I understand the GenAI rules that apply to my courses and assessments.	
PA2	I understand when GenAI use is allowed in assessed work.	(Bulgurcu et al., 2010;
PA3	I know what types of GenAI use need to be reported in assessed work.	McCabe et al., 2008)
PA4	It is clear to me how much detail I should provide about my GenAI use.	
Construct Variable	Academic Self-Efficacy Independent Variable	
5-Point Likert Scale	Please indicate your agreement with the following statements. Measured on a scale of 1 to 5: 1 = strongly disagree, 5 = strongly agree.	
Item	Statement	Source
ASE1	I can solve difficult academic tasks without using GenAI tools if I try hard enough.	
ASE2	I know I can stick to my study goals and achieve them without relying on GenAI tools.	
ASE3	I am confident I can complete difficult assignments without GenAI.	(Nielsen et al., 2018)
ASE4	If my peers can complete study tasks without GenAI, so can I.	

Note. Item wording was adapted from the cited sources to align with the context of students' GenAI use and non-disclosure in assessed academic work.

Table 13 – Demographic questions

Demographic question		Source
What is your age?		
Please enter your age in years: ____		-
What is your gender?		
1. Woman	2. Man	-
3. Other	4. Prefer not to say	
Which of the following best describes your field of study?		
1. Arts & Humanities	2. Business, Finance, & Law	
3. Medicine & Health	4. STEM (Science, Technology, Engineering, Math)	-
5. Social Sciences & Education	6. Other (Please specify)	
What is your current level of study?		
1. First-year bachelor student	2. Second-year bachelor student	-
3. Third-year bachelor student	4. Master's student	
5. Other		
How have you used AI for graded assignments during your last year of studies? (Select all that apply)		
1. Used AI to suggest ideas for research	2. Used AI to explain concepts to me	(Stephenson & Armstrong, 2026)
3. Used AI to structure my thoughts	4. Used GenAI to search for sources, literature, or references	
5. Used AI to summarize a relevant article	6. Used AI to generate visual aids, audio or other media	
7. Used AI to check the grammar, spelling, or flow of my own text	8. Included AI-generated text in assessment after editing it myself	
9. Included AI-generated text in assessment after editing it with a digital tool	10. Included AI-generated text directly in a piece of assessed work	
11. Other	12. None of the above	
Final comments		
Having completed this survey and understanding its focus, are there any additional insights or thoughts you think are important for our research to consider? (optional)		-

Appendix 2 – Measurements of the Constructs

This appendix provides a more detailed overview of the measurement of the constructs included in the survey instrument. The questionnaire was developed to measure the principal constructs in the proposed conceptual model, including classical deterrence theory perceptions, Theory of Planned Behavior variables, academic self-efficacy, policy awareness, students' GenAI use, and non-disclosure behavior and intention in assessed academic work. Each construct was operationalized using adaptations of established items from prior peer-reviewed studies. Using previously developed measures supports content validity, while contextual adaptation is appropriate when instruments are transferred to a new setting (Boateng et al., 2018).

Although some of the original scales used different response scales (such as 7 point likert), all measures in this questionnaire were adjusted to a five-point scales. This decision was made to improve consistency across constructs, simplify the response process, and reduce respondent burden. Methodological literature indicates that five-point scales are widely used and provide a practical balance between response differentiation and ease of completion, especially in longer questionnaires (Kusmaryono et al., 2022). Appendix 1 provides an overview of the survey questions.

2.1 Non-Disclosure

The constructs previous full non-disclosure, previous partial (selective) non-disclosure, full non-disclosure intention, and partial (selective) non-disclosure intention are measured using parallel sets of items. These items were developed through a theory-driven adaptation process, rather than taken from existing GenAI scales. This, because existing research on GenAI disclosure does not yet distinguish between full and partial/selective forms of non-disclosure.

The operationalization builds on the proposal's conceptual distinction between full non-disclosure and partial/selective non-disclosure. Full non-disclosure refers to situations in which GenAI use is not acknowledged at all. Partial/selective non-disclosure refers to situations in which some GenAI use is acknowledged, while relevant details about the type, extent, or location of use are omitted. This distinction is supported by disclosure research showing that disclosure and concealment are separate rather than opposite constructs, and that individuals may engage in both processes at the same time (Larson & Chastain, 1990). In this sense, partial/selective non-disclosure can be understood as a strategic form of disclosure in which some information is revealed while other relevant information is withheld.

The intention items were formulated in line with Theory of Planned Behavior guidance, meaning that they were written as behavior-specific, self-directed, and future-oriented statements (Ajzen, 2006). The previous-behavior items were informed by behavior-specific non-disclosure research, which suggests that withholding is best measured in relation to specific types of information rather than through one broad general question (Yusuf et al., 2025). Additionally, behavior items were restricted to a predefined timeframe, namely the current academic year and the next round of assessment, in order to reduce ambiguity and support more accurate recall (Grigoryeva & Robbins, 2024).

As an example, one item for partial/selective non-disclosure intention is: *“I made my GenAI use seem less extensive than it actually was.”*

2.2 Certainty of Sanctions

The certainty of detection construct was measured using items adapted from Nguyen-Phuoc et al. (2024) and contextualized to the setting of GenAI use in assessed academic work. The original items were adapted to reflect students’ perceived likelihood that undisclosed or partially disclosed GenAI use would be identified in academic assessment. Because the present study distinguishes between full and partial/selective non-disclosure, the wording was changed to focus on not disclosing the extent of GenAI use. Further, the measurement items were shortened and consolidated to improve readability.

The final items capture three main sources of perceived detection: detection by examiners, detection through scanners and tools and procedural mechanisms such as interviews and oral defenses. Together, these items operationalize certainty of sanctions as students’ belief that undisclosed or incompletely disclosed GenAI use can be identified through channels available to institutions.

As an example, one measurement item includes: *“It is likely that undisclosed GenAI use would be detected.”*

2.3 Severity of Sanctions

The severity of sanctions construct is measured using five items adapted from Jantos et al. (2024). These items assess the perceived seriousness and personal consequences of being caught not disclosing the extent of one’s GenAI use in assessed academic work. The first item broadly focuses on the severity of formal consequences, while the second item narrows this focus to potential harm to academic progress. The third item relates to disciplinary process which students will likely have to navigate when being caught. The fourth and fifth items capture more informal or reputational consequences, including damage to one’s standing with lecturers or supervisors and reduced trust or respect from others.

In this way, the items capture a broad range of possible consequences and assess the extent to which these sanctions are perceived as costly and consequential. A sample item is: *“If this was discovered, the formal consequences would be serious.”*

2.4 Swiftiness of Sanctions

The swiftiness of sanctions construct is measured using three items adapted from Freeman et al. (2006). These items focus on the expected speed with which sanctions would follow once undisclosed GenAI use is discovered. More specifically, the items measure respondents’ beliefs about how quickly sanctions would be imposed, the perceived time between discovery and sanctioning, and the delay time by which consequences would follow. Together, these items reflect the operationalization of swiftiness as the perceived promptness of institutional response after detection. The first sample item includes: *“If this was discovered, the university would immediately take action.”*

2.5 Attitudes Toward Non-Disclosure

The attitude construct is measured using four semantic differential items adapted from Greitemeyer and Kastenmüller (2024). These items assess respondents' overall evaluation of the behavior of not fully disclosing GenAI use in assessed academic work. The scale combines various dimensions of respondents' attitude towards the act of disclosure. Items ask whether the behavior is perceived as bad or good, foolish or wise, useless or useful, and harmful or beneficial. In combination, these items capture how positively or negatively respondents view non-disclosure, which is central to how students evaluate disclosure versus concealment. One of the measurement items is: *"For me, not fully disclosing my use of GenAI in assessed academic work would be bad-good."*

2.6 Subjective Norms

The subjective norms construct is measured using four items which are adapted from Greitemeyer and Kastenmüller (2024). These items assess the perceived views of important others regarding the non-disclosure of GenAI use in assessed academic work. The first item captures whether significant others would consider the behavior as acceptable. The second item measures the extent to which they would approve. The third and fourth items shift these measures towards peers instead of significant others. Together, these items operationalize subjective norms as perceived social pressure surrounding non-disclosure. The first item includes: *"Most people who are important to me would find this acceptable."*

2.7 Perceived Behavioral Control

The perceived behavioral control construct was measured using three items adapted from Nguyen-Phuoc et al. (2024). These items assess students' perceived ease, confidence, and control in relation to the disclosure of GenAI use in academic assessments. These items capture the extent to which students believe they are capable of withholding full disclosure if they choose to do so. One item is formulated as: *"I can easily report less of my GenAI use than is required."*

2.8 Policy awareness

Policy awareness was measured using four items that capture students' subjective understanding of the rules governing GenAI use in their courses and assessed work. The construct builds on prior work in information systems research, where policy awareness reflects an individual's understanding of the rules and requirements of a relevant policy (Bulgurcu et al., 2010), and on measures of students' perceived understanding of academic integrity policies (McCabe et al., 2008). The items were adapted to the GenAI assessment context. The first captures general understanding of the rules that apply to courses and assessments, and the second focuses on knowing when GenAI use is permitted. The third addresses awareness of which types of GenAI use must be reported, while the fourth captures clarity about the level of detail required when disclosing GenAI use. Together, the items represent policy awareness as students' subjective understanding of GenAI rules, permitted use, and disclosure requirements. A sample item is: *"I understand the GenAI rules that apply to my courses and assessments."*

2.9 Academic Self-Efficacy

The academic self-efficacy construct was measured using four items adapted from Nielsen et al. (2018), who developed the General Academic Self-Efficacy scale (GASE). These measures were also contextualized to specifically capture students' confidence in completing academic work without using GenAI. The first item focuses on the ability to handle difficult academic tasks, the second on maintaining study goals, the third on being able to succeed through effort, and the fourth on social comparison with other students. Together, these items capture students' perceived confidence in completing assessed academic work independently of GenAI tools. One of the items includes: *"I can solve difficult academic tasks without using GenAI tools if I try hard enough."*

Appendix 3 - Survey Questions, Dutch translation

This appendix presents the Dutch translation of the questionnaire used in this study. The translation was written in clear and accessible Dutch to ensure that participants could easily understand the questions and provide accurate responses.

Table 14 – Survey questions translated to Dutch

<i>Construct Variable</i>	Previous Full Non-Disclosure Dependent Variable	
<i>Header</i>	Hoe vaak heb je afgelopen collegejaar het volgende gedaan bij beoordeelde opdrachten?	
<i>5-Point Frequency Scale</i>	Geef aan hoe vaak elke uitspraak op jou van toepassing was tijdens dit academisch jaar, op een schaal van 1 tot 5: 1 = Nooit, 5 = Altijd. 6 = zeg ik liever niet.	
<i>Item</i>	<i>Statement</i>	<i>Source</i>
PFD1	Ik heb niets vermeld over mijn GenAI-gebruik, terwijl dit wel verwacht werd.	-
PFD2	Ik heb mijn gebruik van GenAI bewust weggelaten in een beoordeelde opdracht.	
<i>Construct Variable</i>	Previous Partial Non-Disclosure Dependent Variable	
<i>Header</i>	Hoe vaak heb je afgelopen collegejaar het volgende gedaan bij beoordeelde opdrachten?	
<i>5-Point Frequency Scale</i>	Geef aan hoe vaak elke uitspraak op jou van toepassing was tijdens dit academisch jaar, op een schaal van 1 tot 5: 1 = Nooit, 5 = Altijd. 6 = zeg ik liever niet.	
<i>Item</i>	<i>Statement</i>	<i>Source</i>
PPD1	Ik heb een deel van mijn GenAI-gebruik vermeld, maar niet alles.	-
PPD2	Ik heb details weggelaten over hoe ik GenAI heb gebruikt.	
PPD3	Ik heb mijn GenAI-gebruik minder uitgebreid laten lijken dan het eigenlijk was.	
PPD4	Ik heb niet duidelijk gemaakt hoeveel van mijn eindresultaat gebaseerd was op GenAI.	
PPD5	Ik heb alleen kleinschalig GenAI-gebruik vermeld en grotere bijdragen weggelaten.	
PPD6	Ik heb mijn GenAI-gebruik in algemene lijnen beschreven in plaats van in detail.	
<i>Construct Variable</i>	Full Non-Disclosure Intention Dependent Variable	
<i>Header</i>	Hoe waarschijnlijk is het dat je de onderstaande dingen doet bij je eerstvolgende beoordeelde opdracht?	
<i>5-Point Likert Scale</i>	Geef aan hoe waarschijnlijk het is dat je elk van de volgende dingen doet, op een schaal van 1 tot 5: 1 = Zeer onwaarschijnlijk, 5 = Zeer waarschijnlijk	
<i>Item</i>	<i>Statement</i>	<i>Source</i>
FDI1	Ik ben van plan niets over mijn GenAI-gebruik te vermelden wanneer dit verwacht wordt.	-

FDI2	Ik ben van plan mijn gebruik van GenAI bewust weg te laten in een beoordeelde opdracht.	
<i>Construct Variable</i>	Partial Non-Disclosure Intention Dependent Variable	
<i>Header</i>	Hoe waarschijnlijk is het dat je de onderstaande dingen doet bij je eerstvolgende beoordeelde opdracht?	
<i>5-Point Likert Scale</i>	Geef aan hoe waarschijnlijk het is dat je elk van de volgende dingen doet, op een schaal van 1 tot 5: 1 = Zeer onwaarschijnlijk, 5 = Zeer waarschijnlijk	
<i>Item</i>	<i>Statement</i>	<i>Source</i>
PDI1	Ik ben van plan een deel van mijn GenAI-gebruik te vermelden, maar niet alles.	-
PDI2	Ik ben van plan details weg te laten over hoe ik GenAI heb gebruikt.	
PDI3	Ik ben van plan mijn GenAI-gebruik minder uitgebreid te laten lijken dan het eigenlijk was.	
PDI4	Ik ben van plan niet duidelijk te maken hoeveel van mijn eindresultaat gebaseerd is op GenAI.	
PDI5	Ik ben van plan alleen kleinschalig GenAI-gebruik te vermelden en grotere bijdragen weg te laten.	
PDI6	Ik ben van plan mijn GenAI-gebruik in algemene lijnen te beschrijven in plaats van in detail.	
<i>Construct Variable</i>	Certainty of Sanctions Independent Variable	
<i>5-Point Likert Scale</i>	In hoeverre ben je het eens met de volgende stellingen? Deze stellingen gaan over het gebruiken van GenAI in beoordeelde opdrachten zonder dit volledig te vermelden. Gemeten op een schaal van 1 tot 5: 1 = Helemaal mee oneens, 5 = Helemaal mee eens	
<i>Item</i>	<i>Statement</i>	<i>Source</i>
CER1	De kans is groot dat niet-vermeld GenAI-gebruik wordt ontdekt.	(Nguyen-Phuoc et al., 2024)
CER2	Mijn docent zal ontdekken hoe uitgebreid ik GenAI heb gebruikt.	
CER3	AI-detectietools zullen laten zien hoe uitgebreid ik GenAI heb gebruikt.	
CER4	Een mondelinge verdediging of vervolginterview zal laten zien hoe uitgebreid ik GenAI heb gebruikt.	
<i>Construct Variable</i>	Severity of Sanctions Independent Variable	
<i>5-Point Likert Scale</i>	Stel je voor dat je niet helemaal eerlijk bent geweest over je GenAI-gebruik bij een beoordeelde opdracht. De volgende stellingen gaan over wat er volgens jou zou gebeuren als dit wordt ontdekt. Gemeten op een schaal van 1 tot 5: 1 = Helemaal mee oneens, 5 = Helemaal mee eens	
<i>Item</i>	<i>Statement</i>	<i>Source</i>
SER1	Als dit wordt ontdekt, zullen de officiële gevolgen ernstig zijn.	(Jantos et al., 2024)

SER2	Als dit wordt ontdekt, zal mijn studievoortgang negatief beïnvloed worden.	
SER3	Als dit wordt ontdekt, zal ik een lastige disciplinaire procedure moeten doorlopen.	
SER4	Als dit wordt ontdekt, zal mijn reputatie bij docenten beschadigd worden.	
SER5	Als dit wordt ontdekt, zullen anderen hun vertrouwen in mij verliezen	
<i>Construct Variable</i>	Swiftness of Sanctions Independent Variable	
<i>5-Point Likert Scale</i>	Stel je voor dat je niet helemaal eerlijk bent geweest over je GenAI-gebruik bij een beoordeelde opdracht. De volgende stellingen gaan over wat er volgens jou zou gebeuren als dit wordt ontdekt. Gemeten op een schaal van 1 tot 5: 1 = Helemaal mee oneens, 5 = Helemaal mee eens	
<i>Item</i>	<i>Statement</i>	<i>Source</i>
SOS1	Als dit wordt ontdekt, zal mijn docent gelijk actie ondernemen.	(Freeman et al., 2006)
SOS2	Als dit wordt ontdekt, zal er snel een definitief besluit worden genomen.	
SOS3	Als dit wordt ontdekt, zal het niet lang duren voordat een consequentie volgt.	
<i>Construct Variable</i>	Attitudes Independent Variable	
<i>5-Point Semantic Differential Scale</i>	Geef aan hoe jij persoonlijk kijkt naar het volgende gedrag. Voor mij is het niet volledig vermelden van mijn GenAI-gebruik:	
<i>Item</i>	<i>Scale</i>	<i>Source</i>
ATT1	Slecht – Goed	(Greitemeyer & Kastenmüller, 2024)
ATT2	Onverstandig – Verstandig	
ATT3	Nutteloos – Nuttig	
ATT4	Schadelijk – Gunstig	
<i>Construct Variable</i>	Subjective Norms Independent Variable	
<i>5-Point Likert Scale</i>	In hoeverre ben je het eens met de volgende stellingen? Deze stellingen gaan over het niet volledig vermelden van je GenAI-gebruik bij beoordeelde opdrachten. Gemeten op een schaal van 1 tot 5: 1 = Helemaal mee oneens, 5 = Helemaal mee eens	
<i>Item</i>	<i>Statement</i>	<i>Source</i>
SUN1	De meeste mensen die belangrijk voor mij zijn, zullen dit gedrag goedkeuren.	(Greitemeyer & Kastenmüller, 2024)
SUN2	De meeste mensen die belangrijk voor mij zijn, zullen dit gedrag begrijpen.	
SUN3	Mijn medestudenten zullen dit gedrag goedkeuren.	
SUN4	Mijn medestudenten zullen dit gedrag begrijpen.	

<i>Construct Variable</i>	Perceived Behavioral Control Independent Variable	
<i>5-Point Likert Scale</i>	In hoeverre ben je het eens met de volgende stellingen? Deze stellingen gaan over het niet volledig vermelden van je GenAI-gebruik bij beoordeelde opdrachten. Gemeten op een schaal van 1 tot 5: 1 = Helemaal mee oneens, 5 = Helemaal mee eens	
<i>Item</i>	<i>Statement</i>	<i>Source</i>
PBC1	Ik kan mijn GenAI-gebruik makkelijk minder uitgebreid vermelden dan verplicht is.	(Nguyen-Phuoc et al., 2024)
PBC2	Ik heb er vertrouwen in dat ik zelf kan bepalen wat ik vermeld over mijn GenAI-gebruik.	
PBC3	Het is volledig aan mij hoeveel ik vermeld over mijn GenAI gebruik.	
<i>Construct Variable</i>	Policy awareness Independent Variable	
<i>5-Point Likert Scale</i>	In hoeverre ben je het eens met de volgende stellingen? Deze stellingen gaan over de regels voor het gebruik van GenAI bij beoordeelde opdrachten. Gemeten op een schaal van 1 tot 5: 1 = Helemaal mee oneens, 5 = Helemaal mee eens	
<i>Item</i>	<i>Statement</i>	<i>Source</i>
PA1	Ik begrijp de GenAI-regels die gelden voor mijn vakken en opdrachten.	(Bulgurcu et al., 2010; McCabe et al., 2008)
PA2	Ik weet wanneer het gebruik van GenAI is toegestaan bij beoordeelde opdrachten.	
PA3	Ik weet welke soorten GenAI-gebruik vermeld moeten worden bij beoordeelde opdrachten.	
PA4	Het is duidelijk hoeveel details ik moet geven over mijn GenAI-gebruik.	
<i>Construct Variable</i>	Academic Self-Efficacy Independent Variable	
<i>5-Point Likert Scale</i>	In hoeverre ben je het eens met de onderstaande stellingen? Gemeten op een schaal van 1 tot 5: 1 = Helemaal mee oneens, 5 = Helemaal mee eens	
<i>Item</i>	<i>Statement</i>	<i>Source</i>
ASE1	Ik kan moeilijke studietaken meestal zonder GenAI oplossen als ik mijn best doe.	(Nielsen et al., 2018)
ASE2	Ik kan mijn studiedoelen behalen zonder afhankelijk te zijn van GenAI.	
ASE3	Ik heb er vertrouwen in dat ik moeilijke opdrachten zonder GenAI kan maken.	
ASE4	Als mijn medestudenten studietaken zonder GenAI kunnen afronden, dan kan ik dat ook.	

Note. Item wording was adapted from the cited sources to align with the context of students' GenAI use and non-disclosure in assessed academic work.

Table 15 – Demographic questions translated to Dutch

Demographic question		Source
<i>Wat is je leeftijd?</i> Vul hier je leeftijd in jaren in: ____		-
<i>Wat is je geslacht?</i>		-
1. Vrouw	2. Man	-
3. Anders	4. Zeg ik liever niet	
<i>Wat is je studierichting?</i>		-
1. Kunst & Geesteswetenschappen	2. Economie, Bedrijfskunde & Recht	-
3. Geneeskunde & Gezondheid	4. Techniek en Exacte Wetenschappen	
5. Sociale Wetenschappen & Onderwijs	6. Anders, namelijk:	
<i>In welk jaar van je studie zit je nu?</i>		-
1. Eerstejaars bachelorstudent	2. Tweedejaars bachelorstudent	-
3. Derdejaars (of oudere) bachelorstudent	4. Masterstudent	
5. Anders, namelijk:		
<i>Hoe heb jij afgelopen collegejaar AI gebruikt voor beoordeelde opdrachten?</i> (Meerdere antwoorden mogelijk)		(Stephenson & Armstrong, 2026)
1. AI gebruikt om ideeën voor onderzoek te bedenken	2. AI gebruikt om concepten aan mij uit te leggen	(Stephenson & Armstrong, 2026)
3. AI gebruikt om mijn gedachten te structureren	4. GenAI gebruikt om bronnen, literatuur of referenties te zoeken	
5. AI gebruikt om een relevant artikel samen te vatten	6. AI gebruikt om afbeeldingen, audio of andere media te maken	
7. AI gebruikt om mijn eigen tekst te controleren	8. AI-gegenereerde tekst zelf bewerkt en daarna gebruikt in een opdracht	
9. AI-gegenereerde tekst met een digitale tool bewerkt en gebruikt in een opdracht.	10. AI-gegenereerde tekst direct overgenomen en gebruikt in een opdracht.	
11. Anders	12. Geen van de bovenstaande opties	
<i>Opmerkingen</i> <i>Nu je deze vragenlijst hebt afgerond en weet waar deze over gaat, zijn er nog aanvullende inzichten of gedachten die volgens jou belangrijk zijn voor dit onderzoek? (optioneel)</i>		-

Appendix 4 – Sample Size Calculation

An a priori power analysis was conducted to determine the minimum sample size required for the planned quantitative analyses. This study uses a cross-sectional survey design and multiple linear regression models to examine predictors of students' GenAI non-disclosure in assessed academic work. The main regression model includes eight predictors: certainty of sanctions, severity of sanctions, swiftness of sanctions, attitudes toward non-disclosure, subjective norms, perceived behavioral control, academic self-efficacy, and policy awareness.

The required sample size was calculated using G*Power 3.1, originally developed by Erdfelder et al. (1996) and further extended by Faul et al. (2007). G*Power is suitable for a priori power calculations across several statistical test families, including F tests and regression analyses. In G*Power, an a priori power analysis was performed under the F-test family using the linear multiple regression option for a fixed model with R^2 deviation from zero. The calculation assumed a medium effect size of $f^2 = .15$, an α error probability of .05, a desired power of .95, and eight predictors.

The R^2 deviation from zero option was selected because the study tests whether the full predictor set explains significant variance in non-disclosure (Plonsky & Ghanbar, 2018). The medium effect size was based on Cohen's benchmark for multiple regression (Cohen, 1988). This choice is further supported by related GenAI academic misconduct studies showing meaningful explained variance in comparable intention and behavior models (Greitemeyer & Kastenmüller, 2024; Zhang et al., 2025). The analysis indicated a minimum required sample size of $n = 160$.




Figure 5 – G* Power 3.1, sample size calculation

Appendix 5 – Data cleaning Process

5.1 Data cleaning criteria

Before detailing the full data cleaning process, this section establishes the specific criteria used to evaluate and filter the dataset.

Uncompleted responses

Prior to importing the data into SPSS, the dataset was cleaned of incomplete submissions. This included respondents who exited the survey before completion, leaving their sessions in an "in-progress" state. Because these responses lacked necessary data for the analysis, they were excluded from the final dataset. Based on this criterion, 139 respondents were removed, reducing the sample from 728 to 589 respondents.

Missing data

Missing data were examined prior to the main analyses. The survey was configured so that respondents could not skip questions. This was done to reduce the likelihood of incomplete responses, both intentional and accidental. Consequently, no missing data were found in the final dataset.

For questions related to past behavior, respondents had the option to select “prefer not to say.” This option was included for items measuring past full and partial non-disclosure behavior (e.g. 8 items in total). Although these responses were not technically missing, they could not be used in the relevant analyses. Therefore, to maintain a consistent dataset, respondents who selected “prefer not to say” on these items were removed from the dataset.

Based on this criterion, 23 respondents were removed, reducing the sample from 589 to 566 respondents.

Response time

A response-time criterion was used to identify potential insufficient effort responding (IER). IER occurs when respondents answer survey items with limited effort. This may mean that they do not carefully follow instructions, interpret item content, or provide accurate responses (Huang et al., 2015).

To identify potential careless or insufficient-effort responding, a minimum completion-time threshold was calculated based on the word count of each questionnaire version. This approach aligns with web survey research, where faster completion times often indicate satisficing and reduced response quality (Malhotra, 2008; Zhang & Conrad, 2014). Furthermore, minimum response-time thresholds have successfully been used to flag speeding and evaluate interventions aimed at reducing rapid responding (Conrad et al., 2017). Therefore, completion time was considered a suitable criterion for flagging potentially insufficient effort responses.

A threshold of 300 words per minute was used for the calculation of the minimum completion time. Rayner et al. (2016) note that skilled or college-educated readers typically read at approximately 200 to 400 words per minute. This threshold was applied to both the Dutch and English versions. This was considered appropriate because Brysbaert et al. (2021) show that reading-rate predictions generalize well from English to Dutch.

The Dutch questionnaire contained a total of 800 words, resulting in a minimum completion time of 160 seconds. The English questionnaire contained 988 words, resulting in a minimum completion time of 198 seconds. In addition, a maximum completion-time threshold of 3600 seconds, or 1 hour, was applied. This threshold was chosen to retain responses from respondents who may have been distracted or interrupted while completing the survey, while still excluding unusually long completion times. Respondents who completed the questionnaire faster than the minimum thresholds or slower than the maximum threshold were flagged as potential IER cases. Based on this criterion, 116 respondents were removed, reducing the sample from 566 to 450 respondents.

Response variability

Zero variance responses were removed from the dataset to reduce the likelihood of careless or low-quality responses. Following Dunn et al. (2018), the Intra-individual Response Variability (IRV) index was used to identify respondents who showed no variation in their answers across the survey. A score of 0.00 variability indicated that the respondent gave identical answers throughout, which was considered a marker of insufficient response engagement.

Respondents with zero variance were therefore excluded from the dataset. Specifically, respondents 338, 347, and 42 showed 0.00 response variability, two of which were already removed by the response time criterion. Based on this criterion, one respondent was removed, reducing the sample from 450 to 449 respondents.

Impossible answers

Impossible answers were removed from the dataset to ensure the quality and consistency of the data used for analysis. First, responses with impossible age values were screened. One respondent reported an age of 99, which was considered implausible for the target population of this study and was therefore removed from the dataset.

In addition, the AI use question was checked for contradictory response patterns. This question allowed respondents to select from 12 possible answer options, including “none of the above.” Respondents who selected “none of the above” while also selecting one or more AI use cases were considered to have provided inconsistent answers. Based on this criterion, respondents 21, 39, and 575 were removed. Based on this criterion, 4 respondents were removed, reducing the sample from 449 to 445 respondents.

Outliers

Outliers were assessed to identify unusually extreme responses that could disproportionately influence the statistical analyses and the validity of the results. The full outlier detection procedure is presented in Appendix 6. In this appendix the analysis is described for each of the twelve constructs: PFD, PPD, FDI, PDI, CER, SER, SOS, ATT, SUN, PBC, PA, and ASE. For this procedure, mean scores were calculated for each construct by averaging the corresponding measurement items, after which boxplots were used to detect potential outliers.

The analysis identified eighteen outliers in total, which occurred only within the SER, SOS, SUN, PBC and PA constructs. Each outlier was examined individually to determine whether the response appeared valid or whether it may have reflected low-quality responding. This assessment was based on respondents' survey completion times. Respondents with outlying scores who also completed the questionnaire below the median response time were considered insufficiently engaged and were therefore removed from the dataset. Based on this criterion, 8 respondents were removed, reducing the sample from 445 to 437 respondents.

Categorizing open-ended responses (other)

Data recategorization was conducted for two demographic questions: field of study and level of study. These questions included open-ended responses that needed to be reassigned to the existing survey categories where possible. For field of study, responses were categorized using the predefined categories, and no new categories were created. For level of study, two additional categories were created: “PhD” and “Other”. In addition, six responses were removed from the dataset because they were illegible or non-compliant, reducing the sample from 437 to 431 respondents.

Field of study

The open-ended responses were categorized using the six predefined study-field categories used in the survey. First, each original response was translated into English where necessary and assigned to the most appropriate existing category. No new categories were created, because introducing additional categories after data collection could have affected how other respondents might have categorized themselves. To make the coding process transparent, two tables are presented. The first table (Table 16) shows the respondent number, translated response, and assigned category. The second table (Table 17) provides a short justification for each category assignment. Six responses were removed from the dataset because they were illegible or non-compliant answers. These are respondents 135, 328, 7, 398, 496, and 560.

Table 16 – Open-ended responses recategorization, Field of study

Respondent nr.	Filled out response	Placed in existing category
87	Applied Psychology	Social Sciences & Education (number 5)
83	Artificial Intelligence	STEM (number 4)
176	Psychology	Social Sciences & Education (number 5)
315	Psychology	Social Sciences & Education (number 5)
36	Urban and Regional Planning (<i>Planologie</i>)	Social Sciences & Education (number 5)
38	Military	Social Sciences & Education (number 5)
233	Marketing and Communication	Business, Finance, & Law (number 2)
209	Marketing & Communication	Business, Finance, & Law (number 2)
189	Marketing	Business, Finance, & Law (number 2)
292	Marketing	Business, Finance, & Law (number 2)
224	Marketing	Business, Finance, & Law (number 2)
17	Management	Business, Finance, & Law (number 2)
20	Information Technology	STEM (number 4)
74	Industrial Design	STEM (number 4)
120	Hospitality	Business, Finance, & Law (number 2)
173	Geosciences	STEM (number 4)
29	Physiotherapy and Sports & Exercise	Medicine & Health (number 3)
231	Engineering	STEM (number 4)
326	Education	Social Sciences & Education (number 5)
127	Economics combined with Social Sciences	Business, Finance, & Law (number 2)
135	Don't currently have one	<u>Removed</u>
122	Animal Management	Business, Finance, & Law (number 2)
126	Communication/Media	Social Sciences & Education (number 5)
188	Communication and Information Sciences	Social Sciences & Education (number 5)
155	Communication	Social Sciences & Education (number 5)
55	Artificial Intelligence	STEM (number 4)

198	Animal Science	STEM (number 4)
328	0000	<u>Removed</u>
7	'blank'	<u>Removed</u>
381	Marketing, Communication & Creativity	Business, Finance, & Law (number 2)
388	Marketing, Communication & Creativity	Business, Finance, & Law (number 2)
379	Design	Arts & Humanities (number 1)
401	Computer science (IT)	STEM (number 4)
366	Computer science	STEM (number 4)
398	'blank'	<u>Removed</u>
411	Marketing and Communications	Business, Finance, & Law (number 2)
416	Media	Arts & Humanities (number 1)
420	Toerism3	Business, Finance, & Law (number 2)
467	Science Communication	STEM (number 4)
474	Communications	Business, Finance, & Law (number 2)
486	Psychology	Social Sciences & Education (number 5)
496	NA	<u>Removed</u>
515	Marketing	Business, Finance, & Law (number 2)
541	Psychology	Social Sciences & Education (number 5)
544	Economies & creative	Business, Finance, & Law (number 2)
550	Hospitality	Business, Finance, & Law (number 2)
551	Psychology	Social Sciences & Education (number 5)
556	Technology	STEM (number 4)
560	PhD	<u>Removed</u>
568	Psychology	Social Sciences & Education (number 5)
571	informatics and cognitive science	STEM (number 4)

Table 17 – Field of study recategorization, justification

Respondent nr.	Justification
87	Applied psychology studies human behavior and mental processes, which fits social sciences.
83	Refers to technical/exact sciences and Artificial Intelligence, clearly within science/technology.
176	Although psychology can be housed in different faculties, the discipline is generally coded as social science here.
315	Psychology focuses on behavior and cognition, fitting social sciences.
36	Spatial planning/urban planning is commonly linked to geography, policy, and social science.
38	Although “Military” is not a direct academic field, it is most closely related to social science areas such as security studies, governance and international relations. Marketing is primarily business-related; communication is secondary here.
233	Marketing is a business discipline
209	Marketing is a business discipline.
189	Marketing is a business discipline.
292	Marketing is a business discipline.
224	Marketing is a business discipline.
17	Management is part of business/organizational studies.
20	Information technology belongs to technology/computing.
74	Industrial design combines engineering, technology, and design; most fitting under STEM in this scheme.
120	Hospitality is a business-oriented professional field.

173	Geosciences are part of natural/earth sciences.
29	Physiotherapy and sport/movement studies are health-related fields.
231	Engineering is explicitly included in STEM.
326	Education directly matches this category.
127	Economics fits best under business/finance, even though the response also mentions social sciences.
135	The respondent does not provide a field of study.
122	Animal management appears to be an applied/professional management field rather than pure animal science.
126	Communication and media studies commonly examine society, media use, and human communication.
188	Communication and information sciences are usually social-scientific fields focused on information, media, and communication.
155	Communication studies fits social sciences.
55	AI is part of computer science/technology.
198	Animal science is a biological/agricultural science field.
328	“0000” is not a meaningful study-field response.
7	This answer was left blank, therefore it was removed
381	Marketing is a business discipline
388	Marketing is a business discipline
379	Design involves creative, visual, and cultural practices, which aligns most closely with arts and humanities.
401	Computer science and IT are technical disciplines focused on computing, technology, and information systems, which clearly fall within STEM.
366	Computer science is a STEM category.
398	This answer was left blank, therefore it was removed
411	Marketing is primarily a business discipline.
416	Media aligns most closely with arts and humanities.
420	Tourism is commonly classified as a business field.
467	Science communication is closely linked to STEM disciplines.
474	Communication is classified here as a business-related field.
486	Psychology is generally classified as a social science.
496	No valid field of study was provided.
515	Marketing is a business discipline.
541	Psychology is generally classified as a social science.
544	Economics is the primary field mentioned and fits business.
550	Hospitality is a business-oriented field.
551	Psychology is generally classified as a social science.
556	Technology clearly falls within STEM.
560	PhD refers to a degree level, not a field of study.
568	Psychology is generally classified as a social science.
571	Informatics is a computing discipline within STEM.

Level of study

The open-ended responses for level of study were categorized using the original survey categories where possible: first-year bachelor student, second-year bachelor student, third-year bachelor student, and master's student. Two new categories were created, namely PhD and Other. The PhD category was added because several respondents indicated that they were PhD or doctoral students. The Other category was used for responses that did not fit the original categories or the PhD category, such as employed respondents, graduates, fourth-year bachelor students, or unclear study levels. These responses were still kept because they provide meaningful information and these respondents may still have valuable experiences regarding their past use and disclosure of AI. Therefore, they were retained in the dataset but treated separately under the category "Other." To make the recoding process transparent, two tables are presented. The first table (Table 18) shows the respondent number, original response, and assigned category. The second table (Table 19) provides a short justification for the categorized responses.

Table 18 – Open-ended responses recategorization, Level of study

Respondent nr.	Filled out response	Placed in category
253	Employed / Working	Other (number 6)
102	Working as a primary school teacher	Other (number 6)
38	Working / Employed	Other (number 6)
353	Undergraduate	Other (number 6)
169	Second-year Master's student	Masters (number 4)
45	Recently graduated with a Bachelor's degree	Other (number 6)
238	Recently graduated with a Bachelor's degree	Other (number 6)
197	PhD student	PhD (number 5)
220	PhD student	PhD (number 5)
261	PhD	PhD (number 5)
313	PhD	PhD (number 5)
332	PhD	PhD (number 5)
254	Graduated	Other (number 6)
326	Graduate Master's	Other (number 6)
301	Graduate	Other (number 6)
305	Fourth-year Bachelor student	Other (number 6)
260	Doctoral	PhD (number 5)
204	DBA (Doctor of Business Administration)	PhD (number 5)
348	Fourth-year Bachelor's student	Other (number 6)
94	Graduated	Other (number 6)
277	Graduated	Other (number 6)
75	Graduated from my Bachelor's degree last year	Other (number 6)
389	Graduated	Other (number 6)
585	One-year accelerated bachelor's after a first bachelor's degree	Other (number 6)
570	Starting my second bachelors	Other (number 6)
536	PhD student	PhD (number 5)
509	PhD	PhD (number 5)
551	Honour's 4th Year	Other (number 6)
413	Fresh off PhD	PhD (number 5)
556	fourth year	Other (number 6)
568	doctoral student	PhD (number 5)
573	College	Other (number 6)

Table 19 – Level of study recategorization, justification

Respondent nr.	Justification
253	Employment status, not a study level.
102	Employment status, not a study level.
38	Employment status, not a study level.
353	Bachelor level is implied, but the year is not specified.
169	Second-year master's student, fits the master's student category.
45	Graduated from bachelor's, not currently in a listed study level.
238	Graduated from bachelor's, not currently in a listed study level.
197	The respondent explicitly identifies as a PhD student.
220	The respondent explicitly identifies as a PhD student.
261	"PhD" directly refers to doctoral-level study.
313	"PhD" directly refers to doctoral-level study.
332	"PhD" directly refers to doctoral-level study.
254	Graduated, but no current study level is given.
326	Graduated from Master's, not currently in a listed study level.
301	Graduated, but no current study level is given.
305	Fourth-year bachelor is not one of the listed options.
260	"Doctoral" refers to doctoral-level education, which is grouped under PhD.
204	A DBA, or Doctor of Business Administration, is a doctoral-level qualification and is therefore grouped under PhD.
348	Fourth-year bachelor is not one of the listed options.
94	Graduated, but no current study level is given.
277	Graduated, but no current study level is given.
75	Graduated from bachelor's, not currently in a listed study level.
389	Graduated, but no current study level is given.
585	This does not fit the standard bachelor/master categories.
570	The respondent is pursuing a second bachelor's degree rather than indicating a study year.
536	The respondent is a PhD student.
509	The respondent indicates PhD level study.
551	A fourth-year honors programme does not fit the standard categories.
413	The respondent indicates PhD-level education.
556	A fourth-year student does not fit the available bachelor-year categories.
568	A doctoral student is classified as PhD.
573	The response does not specify a study year or degree level matching the available categories.

5.2 Data cleaning process

This section shows the stepwise process of data cleaning. It shows the process of each step as discussed previously in Appendix 5.1.

Table 20 – Data cleaning process

Importing data into SPSS		
Columns removed in the SPSS file	StartDate	Finished
	EndDate	RecordedDate
	Status	Responseld
	IPAddress	RecipientFirstName
	Progress	RecipientEmail
	ExternalReference	LocationLatitude
	DistributionChannel	LocationLongitude
	UserLanguage	
Each respondent was assigned a unique numerical identifier ranging from 1 to 589.		# respondents at start: 589
Response time.		Changed view from 2 decimals to 0 decimals
Cleaning variable settings		
Changed variable name	VAR0001 → Respondent_number	
	VAR0002 → Duration_seconds	
	VAR0003 → Selected_language	
	VAR0004 → PFD_1	
	VAR0005 → PFD_2	
	VAR0006 → PPD_1	
	VAR0007 → PPD_2	
	VAR0008 → PPD_3	
	VAR0009 → PPD_4	
	VAR0010 → PPD_5	
	VAR0011 → PPD_6	
	VAR0012 → FDI_1	
	VAR0013 → FDI_2	
	VAR0014 → PDI_1	
	VAR0015 → PDI_2	
	VAR0016 → PDI_3	
	VAR0017 → PDI_4	
	VAR0018 → PDI_5	
	VAR0019 → PDI_6	
	VAR0020 → CER_1	
	VAR0021 → CER_2	
	VAR0022 → CER_3	
	VAR0023 → CER_4	
	VAR0024 → SER_1	
	VAR0025 → SER_2	
	VAR0026 → SER_3	
Note. The table is continued on the next page.		

Table 20 Continued

Changed variable name	VAR0027 → SER_4 VAR0028 → SER_5 VAR0029 → SOS_1 VAR0030 → SOS_2 VAR0031 → SOS_3 VAR0032 → ATT_1 VAR0033 → ATT_2 VAR0034 → ATT_3 VAR0035 → ATT_4 VAR0036 → SUN_1 VAR0037 → SUN_2 VAR0038 → SUN_3 VAR0039 → SUN_4 VAR0040 → PBC_1 VAR0041 → PBC_2 VAR0042 → PBC_3 VAR0043 → PA_1 VAR0044 → PA_2 VAR0045 → PA_3 VAR0046 → PA_4 VAR0047 → ASE_1 VAR0048 → ASE_2 VAR0049 → ASE_3 VAR0050 → ASE_4 VAR0051 → Age VAR0052 → Gender VAR0053 → Field_of_Study VAR0054 → Level_of_Study VAR0055 → AI_use_1 VAR0056 → AI_use_2 VAR0057 → AI_use_3 VAR0058 → AI_use_4 VAR0059 → AI_use_5 VAR0061 → AI_use_6 VAR0062 → AI_use_7 VAR0063 → AI_use_8 VAR0064 → AI_use_9 VAR0065 → AI_use_10 VAR0066 → AI_use_11 VAR0067 → AI_use_12 VAR0068 → Additional_insights
Changed variable label	Respondent_number → Respondent number Duration_seconds → Duration in seconds Selected_language → Language in which the survey was administered

Note. The table is continued on the next page.

Table 20 Continued

Changed variable label	PPD_1 → I disclosed some of my GenAI use but not all of it. PPD_2 → I left out details about how I used GenAI. PPD_3 → I made my GenAI use seem less extensive than it actually was. PPD_4 → I did not make clear how much my final work depended on GenAI output. PPD_5 → I only reported minor uses of GenAI but left out more substantial contributions. PPD_6 → I described my GenAI use in general terms rather than in specific detail. FDI_1 → I will keep all of my GenAI use undisclosed when disclosure is expected. FDI_2 → I will deliberately leave out any mention of having used GenAI in my work. PDI_1 → I will disclose some of my GenAI use but not all of it. PDI_2 → I will leave out details about how I used GenAI. PDI_3 → I will make my GenAI use seem less extensive than it actually was. PDI_4 → I will give limited information about how much my final work depended on GenAI output. PDI_5 → I will only report minor uses of GenAI but leave out more substantial contributions. PDI_6 → I will describe my GenAI use in general terms rather than in specific detail. CER_1 → It is likely that undisclosed GenAI use will be detected. CER_2 → My examiner would detect the extent of my GenAI use. CER_3 → AI detection tools would detect the extent of my GenAI use. CER_4 → An oral defense or follow-up interview would reveal the extent of my GenAI use. SER_1 → If this was discovered, the formal consequences would be serious. SER_2 → If this was discovered, my academic progress would be hindered. SER_3 → If this was discovered, I would have to go through a difficult disciplinary process. SER_4 → If this was discovered, my reputation will be harmed. SER_5 → If this was discovered, people important to me would lose trust in me.
------------------------	---

Note. The table is continued on the next page.

Table 20 Continued

Changed variable label	SOS_1 → If this was discovered, my examiner would immediately take action. SOS_2 → If this was discovered, a final decision would be reached in a very short time. SOS_3 → If this was discovered, it would not take long for me to receive a sanction. ATT_1 → For me, not disclosing the extent of my GenAI use would be ... - Bad:Good ATT_2 → For me, not disclosing the extent of my GenAI use would be ... - Foolish:Wise ATT_3 → For me, not disclosing the extent of my GenAI use would be ... - Useless:Useful ATT_4 → For me, not disclosing the extent of my GenAI use would be ... - Harmful:Beneficial SUN_1 → Most people who are important to me would find this acceptable. SUN_2 → Most people who are important to me would understand this behavior. SUN_3 → My peers would find this acceptable. SUN_4 → My peers would understand this behavior. PBC_1 → I can easily report my GenAI use less extensively than required. PBC_2 → I am confident I can control what I report about my GenAI use. PBC_3 → It is entirely up to me how much I disclose about my GenAI use. PA_1 → I understand the GenAI rules that apply to my courses and assessments. PA_2 → I understand when GenAI use is allowed in assessed work. PA_3 → I know what types of GenAI use need to be reported in assessed work. PA_4 → It is clear to me how much detail I should provide about my GenAI use. ASE_1 → I can solve difficult academic tasks without using GenAI tools if I try hard enough. ASE_2 → I know I can stick to my study goals and achieve them without relying on GenAI tools. ASE_3 → I am confident I can complete difficult assignments without GenAI. ASE_4 → If my peers can complete study tasks without GenAI, so can I.
------------------------	--

Note. The table is continued on the next page.

Table 20 Continued

Changed variable label	Age → Please enter your age in years. Gender → What is your gender. Field_of_Study → What is your field of study. Level_of_Study → What is your current level of study. AI_use_1 → Used AI to suggest ideas for research AI_use_2 → Used AI to explain concepts to me AI_use_3 → Used AI to structure my thoughts AI_use_4 → Used GenAI to search for sources, literature, or references AI_use_5 → Used AI to summarize a relevant article AI_use_6 → Used AI to generate visual aids, audio or other media AI_use_7 → Used AI to check the grammar, spelling, or flow of my own text AI_use_8 → Included AI-generated text in assessment after editing it myself AI_use_9 → Included AI-generated text in assessment after editing it with a digital tool AI_use_10 → Included AI-generated text directly in a piece of assessed work AI_use_11 → Other AI_use_12 → None of the above Additional_insights → Are there any additional insights or thoughts you think are important for our research to consider?
Updated values for: PFD_1, PFD_2, PPD_1, PPD_2, PPD_3, PPD_4, PPD_5, PPD_6	1= Never 2 = Sometimes 3 = About half the time 4 = Most of the time 5 = Always 6 = Prefer not to say
Updated values for: FDI_1, FDI_2, PDI_1, PDI_2, PDI_3, PDI_4, PDI_5, PDI_6	1= Extremely unlikely 2 = Somewhat unlikely 3 = Neutral 4 = Somewhat likely 5 = Extremely likely
Updated values for: CER_1, CER_2, CER_3, CER_4, SER_1, SER_2, SER_3, SER_4, SER_5, SOS_1, SOS_2, SOS_3, SUN_1, SUN_2, SUN_3, SUN_4, PBC_1, PBC_2, PBC_3, PA_1, PA_2, PA_3, PA_4, ASE_1, ASE_2, ASE_3, ASE_4	1= Strongly disagree 2 = Somewhat disagree 3 = Neutral 4 = Somewhat agree 5 = Strongly agree

Note. The table is continued on the next page.

Table 20 Continued

Updated values for: ATT_1	1 = Bad 2 = 2 3 = 3 4 = 4 5 = Good
Updated values for: ATT_2	1 = Foolish 2 = 2 3 = 3 4 = 4 5 = Wise
Updated values for: ATT_3	1 = Useless 2 = 2 3 = 3 4 = 4 5 = Useful
Updated values for: ATT_4	1 = Harmful 2 = 2 3 = 3 4 = 4 5 = Beneficial
Updated values for: Selected language	1 = EN (English) 2 = NL (Dutch)
Changed measure type to scale for	“Respondent_number”, “Duration_seconds”, and “Age”
Changed the measure type to ordinal for	"Selected_language", "PFD_1", "PFD_2", "PPD_1", "PPD_2", "PPD_3", "PPD_4", "PPD_5", "PPD_6", "FDI_1", "FDI_2", "PDI_1", "PDI_2", "PDI_3", "PDI_4", "PDI_5", "PDI_6", "CER_1", "CER_2", "CER_3", "CER_4", "SER_1", "SER_2", "SER_3", "SER_4", "SER_5", "SOS_1", "SOS_2", "SOS_3", "ATT_1", "ATT_2", "ATT_3", "ATT_4", "SUN_1", "SUN_2", "SUN_3", "SUN_4", "PBC_1", "PBC_2", "PBC_3", "PA_1", "PA_2", "PA_3", "PA_4", "ASE_1", "ASE_2", "ASE_3" and "ASE_4",
Changed the measure type to nominal for	"Gender", "Field_of_Study", "Level_of_Study", "AI_use_1", "AI_use_2", "AI_use_3", "AI_use_4", "AI_use_5", "AI_use_6", "AI_use_7", "AI_use_8", "AI_use_9", "AI_use_10", "AI_use_11" and "AI_use_12"
<i>Removing respondents due to missing data, (i.e. not filling in all answers in the survey)</i>	
There was no missing data present in this survey.	None of the respondents were removed, due to missing data.
Respondents removed due to answering “prefer not to say, for questions related to past behavior.	Respondent number: 369 → PFD_1, PFD_2, PPD_1, PPD_2, PPD_3, PPD_4, PPD_5, PPD_6 480 → PFD_1, PFD_2, PPD_1, PPD_2, PPD_3, PPD_4, PPD_5, PPD_6
<i>Note. The table is continued on the next page.</i>	

Table 20 Continued

Total removed	583 → PFD_1, PFD_2, PPD_1, PPD_2, PPD_3, PPD_4, PPD_5, PPD_6	
	325 → PFD_1, PFD_2, PPD_1, PPD_3, PPD_4, PPD_6	
	428 → PFD_1, PPD_1, PPD_3, PPD_4, PPD_5	
	463 → PFD_1, PFD_2, PPD_1, PPD_2, PPD_3	
	249 → PFD_1, PFD_2, PPD_3	
	561 → PFD_1, PFD_2, PPD_2	
	577 → PFD_1, PFD_2, PPD_2	
	557 → PPD_2, PPD_4, PPD_5	
	391 → PFD_1	
	240 → PFD_1	
	305 → PFD_1	
	394 → PPD_1	
	306 → PFD_1, PPD_4	
	497 → PPD_1, PPD_2	
	582 → PFD_2, PPD_4	
	268 → PPD_3	
	206 → PPD_4	
	422 → PFD_2	
	487 → PFD_2	
	526 → PFD_2	
	584 → PFD_2	
	23 respondents	
Subtotal respondents	Subtotal number of respondents before removal: 589	Subtotal number of respondents after removal: 566
<i>Insufficient effort analysis, Survey completion time</i>		
Number of survey questions		
Construct items	6	
Demographic questions	47	
The minimum completion time for survey in seconds	59 seconds	
The maximum completion time for survey in seconds	408947 seconds	
Average completion time taken by respondents	1797,14 seconds	
Median time completion time taken by respondents	331 seconds	
<i>Minimum time constraints</i>		
Minimum time constraint for English survey	198 seconds	
Minimum time constraint for Dutch survey	160 seconds	
<i>Note.</i> The table is continued on the next page		

Table 20 Continued

<i>Minimum time constraints</i>	
Maximum time constraint for both English and Dutch surveys	3600 seconds
Respondent_number removed from English survey due to < 198 sec	Respondent number:
	N=564 → 59 seconds
	N=219 → 69 seconds
	N=274 → 71 seconds
	N=164 → 75 seconds
	N=134 → 76 seconds
	N=330 → 81 seconds
	N=383 → 82 seconds
	N=518 → 83 seconds
	N=227 → 87 seconds
	N=190 → 88 seconds
	N=284 → 88 seconds
	N=358 → 91 seconds
	N=542 → 93 seconds
	N=307 → 95 seconds
	N=350 → 96 seconds
	N=538 → 97 seconds
	N=288 → 97 seconds
	N=291 → 98 seconds
	N=338 → 99 seconds
	N=437 → 101 seconds
	N=461 → 101 seconds
	N=265 → 104 seconds
	N=552 → 107 seconds
	N=312 → 107 seconds
	N=425 → 113 seconds
	N=531 → 114 seconds
	N=483 → 116 seconds
	N=247 → 116 seconds
	N=309 → 120 seconds
	N=317 → 120 seconds
	N=414 → 126 seconds
	N=263 → 127 seconds
	N=370 → 128 seconds
	N=215 → 131 seconds
	N=280 → 131 seconds
	N=424 → 133 seconds
	N=412 → 134 seconds
	N=419 → 135 seconds
	N=476 → 136 seconds

Note. The table is continued on the next page

Table 20 Continued	
Respondent_number removed from English survey due to < 198 sec	Respondent number:
	N=205 → 139 seconds
	N=499 → 142 seconds
	N=511 → 143 seconds
	N=524 → 144 seconds
	N=415 → 145 seconds
	N=374 → 147 seconds
	N=417 → 152 seconds
	N=385 → 152 seconds
	N=459 → 155 seconds
	N=286 → 155 seconds
	N=310 → 156 seconds
	N=257 → 159 seconds
	N=404 → 160 seconds
	N=256 → 160 seconds
	N=454 → 161 seconds
	N=349 → 161 seconds
	N=408 → 163 seconds
	N=354 → 163 seconds
	N=289 → 164 seconds
	N=520 → 166 seconds
	N=281 → 168 seconds
	N=479 → 169 seconds
	N=347 → 169 seconds
	N=156 → 172 seconds
	N=214 → 175 seconds
	N=311 → 175 seconds
	N=324 → 176 seconds
	N=339 → 176 seconds
	N=148 → 177 seconds
	N=183 → 177 seconds
	N=588 → 182 seconds
	N=161 → 182 seconds
	N=278 → 184 seconds
	N=336 → 184 seconds
	N=225 → 186 seconds
	N=345 → 186 seconds
	N=453 → 188 seconds
	N=392 → 188 seconds
	N=527 → 189 seconds
	N=346 → 189 seconds
	N=563 → 190 seconds
	N=507 → 191 seconds
	N=273 → 192 seconds
	N=452 → 193 seconds

Note. The table is continued on the next page.

Table 20 Continued

Respondent_number removed from English survey due to < 198 sec	Respondent number: N=455 → 193 seconds N=222 → 193 seconds N=460 → 195 seconds N=170 → 195 seconds N=335 → 196 seconds
Total removed	88 respondents
Respondent_number removed from Dutch survey due to < 160 sec	Respondent number: N=213 → 97 seconds N=540 → 100 seconds N=212 → 105 seconds N=504 → 107 seconds N=505 → 113 seconds N=565 → 116 seconds N=546 → 121 seconds N=136 → 122 seconds N=502 → 133 seconds N=85 → 143 seconds N=105 → 149 seconds N=106 → 152 seconds N=112 → 155 seconds N=80 → 156 seconds N=143 → 156 seconds N=294 → 158 seconds N=501 → 159 seconds
Total removed	17 respondents
Respondent_number removed from English and Dutch surveys due to >3600 seconds	Respondent number: N= 92 → 408947 seconds N= 46 → 215227 seconds N= 37 → 81995 seconds N= 69 → 39213 seconds N=174 → 28368 seconds N=302 → 12951 seconds N=320 → 7608 seconds N=250 → 7026 seconds N= 96 → 5650 seconds N= 10 → 5134 seconds N= 71 → 3775 seconds
Total removed	11 respondents
<i>Note.</i> The table is continued on the next page.	

Table 20 Continued

Table 2b Continued			
Subtotal number of respondents removed due to outside of time range:	116	Respondents left after removing all responses that were outside of the set time range:	450
<i>Insufficient effort analysis, Intra-individual Response Variability</i>			
Respondents removed due to zero Intra-individual Response Variability (IRV)	Respondent number: 338 → IRV = 0.00 (already removed) 347 → IRV = 0.00 (already removed) 42 → IRV = 0.00		
Subtotal number of respondents removed due to zero variability	1	Respondents left after removing all responses that had zero variability	449
<i>Impossible answers</i>			
Respondents removed due to impossible age	Respondent number: 236 → AGE = 99		
Respondents removed due to impossible AI use answer (None of the above combined with other use cases).	Respondent number: 21 39 575		
Subtotal number of respondents removed due to impossible answers	4	Respondents left after removing all responses that had impossible answers	445
<i>Outlier detection</i>			
Respondents removed based on outlier detection (as conducted in Appendix 6 – Outlier Detection)	Respondent number: 438, 477, 581, 272, 323, 318, 111, 110		
Subtotal number of respondents removed based on outlier detection	8	Respondents left after removing all responses that had impossible answers	437
<i>Open-ended response cleaning</i>			
Amount of respondents answering an open answer on the field of study question	29 respondents		
Respondent number answering an open answer on the field of study question	87, 83, 176, 315, 36, 38, 233, 209, 189, 292, 224, 17, 20, 74, 120, 173, 29, 231, 326, 127, 135, 122, 126, 188, 155, 55, 198, 328, 7, 381, 388, 379, 401, 366, 398, 411, 416, 420, 467, 474, 486, 496, 515, 541, 544, 550, 551, 556, 560, 568, 571		
	Total: 51 respondents		
<i>Note.</i> The table is continued on the next page.			

Table 20 Continued			
Respondent number answering an open answer on the level of study question	253, 102, 38, 353, 169, 45, 238, 197, 220, 261, 313, 332, 254, 326, 301, 305, 260, 204, 348, 94, 277, 75, 389, 585, 570, 536, 509, 551, 413, 556, 568, 573		
	Total: 32		
Respondents removed due to ineligible answers	Respondent number: 135, 328, 7, 398, 496, 560		
Subtotal number of respondents removed due to impossible answers	6	Respondents left after removing all responses that had impossible answers	431
Open responses categorized. Consult Table 16 to Table 19 for elucidation			
End of the data cleaning process.			

Appendix 6 – Outlier Detection

This appendix presents the outlier detection process conducted for the twelve constructs; PFD, PPD, FDI, PDI, CER, SER, SOS, ATT, SUN, PBC, PA, and ASE. The objective of this analysis was to improve the quality of the dataset, as outliers may have a disproportionate influence on statistical analyses and the validity of the results (Wang et al., 2019)

To identify the outliers of this dataset, mean scores were calculated for each construct in SPSS by averaging the corresponding measurement items. These mean scores were subsequently used for the outlier analysis (Finch, 2012).

The analysis identified 21 outliers occurrences across 18 unique respondents. As presented in Table 21, these outliers were observed for the SER, SOS, SUN, PBC and PA constructs. There were no outliers identified for the remaining constructs, indicating that responses for the remaining constructs showed relatively homogeneous distributions without unusually extreme observations. Moreover, none of the identified observations were classified as extreme outliers (Spiegel & Stephens, 2013). Figure 6 through Figure 17 display the corresponding boxplots, which provide a visual representation of the detected outliers (Smiti, 2020).

Each identified outlier was examined individually to determine whether it represented an invalid, biased, or legitimate response (Zijlstra et al., 2010). This assessment was based on the respondents' survey completion times. Eight of the eighteen outliers were associated with respondents who completed the questionnaire below the median of the response time of the whole dataset. As these respondents were considered insufficiently engaged with the survey, their responses respondents number 438, 477, 581, 272, 323, 318, 111, and 110 were removed from the dataset. These outliers can result from various factors such as participants response errors or data entry errors (Kwak & Kim, 2017). The remaining ten outliers were kept as completion times were considered appropriate.

Table 21 – Outlier detection overview

Variable		Outlier					
		1	2	3	4	5	6
Severity of sanctions (SER)	Row	263	270	329	359		
	Respondent	344	357	438	477		
Swiftiness of sanctions (SOS)	Row	326	329	390	359		
	Respondent	434	438	523	477		
Subjective norms (SUN)	Row	431	433	439	437		
	Respondent	581	586	323	272		
Perceived behavioral control (PBC)	Row	438	436	431	381	416	264
	Respondent	318	111	581	512	559	348
Policy awareness (PA)	Row	409	95	41			
	Respondent	550	110	48			




Figure 6 – Boxplot PFD_mean




Figure 7 – Boxplot PPD_mean




Figure 8 – Boxplot FDI_mean




Figure 9 – Boxplot PDI_mean




Figure 10 – Boxplot CER_mean




Figure 11 – Boxplot SER_mean




Figure 12 – Boxplot *SOS_mean*




Figure 13 – Boxplot *ATT_mean*




Figure 14 – Boxplot *SUN_mean*




Figure 15 – Boxplot PBC_mean




Figure 16 – Boxplot PA_mean




Figure 17 – Boxplot ASE_mean

Appendix 7 – Exploratory Factor Analysis (EFA)

This appendix documents the exploratory factor analysis (EFA) in detail and reports the output that supports the decisions described in Section 3.3.2. The analysis was conducted as two separate models. The first model combined the disclosure intention constructs with the predictor constructs, since together these form the theoretical model derived from the Theory of Planned Behavior and Deterrence Theory. The second model contained the past behavior constructs, namely past full non-disclosure and past partial non-disclosure.

For each model, this appendix reports the communalities and the total variance explained for the final solution, followed by the rotated component matrix of the initial solution, before any items were removed. The item-level decisions are explained alongside these tables.

7.1 EFA Model 1: Intentions and predictors

The final solution for the first model was reached after two items were removed, as explained below. It retained 37 items across ten components. Table 22 reports the KMO measure of sampling adequacy (0.866) and Bartlett's test of sphericity ($\chi^2(741) = 9906.45$, $p < .001$), which together confirmed that the data were suitable for factor analysis. Table 23 reports the communalities of the retained items, which ranged from 0.489 to 0.838. With the exception of CER_4, which fell marginally below the 0.50 guideline, these values indicate that the retained items were adequately represented by the extracted components. Table 24 reports the total variance explained. The ten retained components together accounted for 72.94% of the variance, which is a satisfactory level for this type of analysis (Hair et al., 2014).

Table 22 – KMO and Bartlett's Test Results for EFA Model

Kaiser-Meyer-Olkin Measure of Sampling Adequacy.		.866
Bartlett's Test of Sphericity	Approx. Chi-Square	9906.454
	df	741
	Sig.	.000

Table 23 – Communalities, final solution, Model 1

Communalities	Initial	Extraction
FDI_1: I will keep all of my GenAI use undisclosed when disclosure is expected.	1.000	.778
FDI_2: I will deliberately leave out any mention of having used GenAI in my work.	1.000	.804
PDI_1: I will disclose some of my GenAI use but not all of it.	1.000	.657
PDI_2: I will leave out details about how I used GenAI.	1.000	.707
PDI_3: I will make my GenAI use seem less extensive than it actually was	1.000	.791
PDI_4: I will give limited information about how much my final work depended on GenAI output.	1.000	.763
PDI_5: I will only report minor uses of GenAI but leave out more substantial contributions.	1.000	.711
PDI_6: I will describe my GenAI use in general terms rather than in specific detail.	1.000	.679
CER_1: It is likely that undisclosed GenAI use will be detected.	1.000	.659
CER_2: My examiner would detect the extent of my GenAI use.	1.000	.776
CER_3: AI detection tools would detect the extent of my GenAI use.	1.000	.689
CER_4: An oral defense or follow-up interview would reveal the extent of my GenAI use.	1.000	.489
SER_1: If this was discovered, the formal consequences would be serious.	1.000	.727
SER_2: If this was discovered, my academic progress would be hindered.	1.000	.803
SER_3: If this was discovered, I would have to go through a difficult disciplinary process.	1.000	.747
SER_4: If this was discovered, my reputation will be harmed.	1.000	.626
SOS_2: If this was discovered, a final decision would be reached in a very short time.	1.000	.838
SOS_3: If this was discovered, it would not take long for me to receive a sanction.	1.000	.830
ATT_1: For me, not disclosing the extent of my GenAI use would be – Bad : Good	1.000	.701
ATT_2: For me, not disclosing the extent of my GenAI use would be – Foolish : Wise	1.000	.737
ATT_3: For me, not disclosing the extent of my GenAI use would be – Useless : Useful	1.000	.708
ATT_4: For me, not disclosing the extent of my GenAI use would be – Harmful : Beneficial	1.000	.776
SUN_1: Most people who are important to me would find this acceptable.	1.000	.684
SUN_2: Most people who are important to me would understand this behavior.	1.000	.740
SUN_3: My peers would find this acceptable.	1.000	.776
SUN_4: My peers would understand this behavior.	1.000	.763
PBC_1: I can easily report my GenAI use less extensively than required.	1.000	.690
PBC_2: I am confident I can control what I report about my GenAI use.	1.000	.766
PBC_3: It is entirely up to me how much I disclose about my GenAI use.	1.000	.681
PA_1: I understand the GenAI rules that apply to my courses and assessments.	1.000	.750
PA_2: I understand when GenAI use is allowed in assessed work.	1.000	.761
PA_3: I know what types of GenAI use need to be reported in assessed work.	1.000	.803
PA_4: It is clear to me how much detail I should provide about my GenAI use.	1.000	.697
ASE_1: I can solve difficult academic tasks without using GenAI tools if I try hard enough.	1.000	.681
ASE_2: I know I can stick to my study goals and achieve them without relying on GenAI tools.	1.000	.770
ASE_3: I am confident I can complete difficult assignments without GenAI.	1.000	.806
ASE_4: If my peers can complete study tasks without GenAI, so can I	1.000	.622

Note. Extraction Method: Principal Component Analysis.

Table 24 – Total variance explained, final solution, Model 1.

Component	Initial Eigenvalues			Extraction Sums of Squared Loadings			Rotation Sums of Squared Loadings ^a
	Total	% of Variance	Cumulative %	Total	% of Variance	Cumulative %	Total
1	9.125	24.662	24.662	9.125	24.662	24.662	6.223
2	3.735	10.094	34.756	3.735	10.094	34.756	4.191
3	3.291	8.896	43.652	3.291	8.896	43.652	3.421
4	2.386	6.448	50.100	2.386	6.448	50.100	3.186
5	2.075	5.609	55.709	2.075	5.609	55.709	2.962
6	1.654	4.470	60.179	1.654	4.470	60.179	5.517
7	1.427	3.857	64.035	1.427	3.857	64.035	3.093
8	1.278	3.453	67.488	1.278	3.453	67.488	5.814
9	1.135	3.066	70.555	1.135	3.066	70.555	2.154
10	.881	2.382	72.936	.881	2.382	72.936	2.185
11	.771	2.084	75.020				
12	.657	1.776	76.796				
13	.638	1.726	78.522				
14	.584	1.577	80.099				
15	.575	1.554	81.654				
16	.512	1.383	83.037				
17	.489	1.322	84.358				
18	.475	1.285	85.643				
19	.446	1.207	86.850				
20	.411	1.110	87.960				
21	.398	1.074	89.034				
22	.377	1.020	90.054				
23	.373	1.009	91.063				
24	.354	.957	92.020				
25	.320	.866	92.886				
26	.304	.823	93.709				
27	.275	.744	94.453				
28	.263	.712	95.165				
29	.257	.695	95.860				
30	.242	.653	96.513				
31	.224	.605	97.118				
32	.210	.567	97.685				
33	.202	.547	98.232				
34	.194	.524	98.756				
35	.174	.471	99.227				
36	.156	.421	99.648				
37	.130	.352	100.000				

Note. Extraction Method: Principal Component Analysis. a. When components are correlated, sums of squared loadings cannot be added to obtain a total variance.

Table 25 presents the pattern matrix of the final solution for the first model, after the removal of SER_5 and SOS_1. Each retained item loaded on its intended component, and the ten components correspond to the ten constructs in the theoretical model. Loadings below 0.30 are suppressed for readability. The one secondary loading that remained, for FDI_2, is discussed in the following section.

Table 25 – Pattern matrix, final solution, Model 1

	Component									
	1	2	3	4	5	6	7	8	9	10
FDI_1									-.720	
FDI_2	.327								-.668	
PDI_1	.729									
PDI_2	.764									
PDI_3	.771									
PDI_4	.760									
PDI_5	.696									
PDI_6	.846									
CER_1					.750					
CER_2					.874					
CER_3					.806					
CER_4					.586					
SER_1		.795								
SER_2		.870								
SER_3		.856								
SER_4		.738								
SOS_2										-.762
SOS_3										-.757
ATT_1								.769		
ATT_2								.811		
ATT_3								.761		
ATT_4								.802		
SUN_1						.777				
SUN_2						.820				
SUN_3						.844				
SUN_4						.848				
PBC_1							-.695			
PBC_2							-.839			
PBC_3							-.776			
PA_1			.862							
PA_2			.853							
PA_3			.901							
PA_4			.821							
ASE_1				.819						
ASE_2				.873						
ASE_3				.896						
ASE_4				.645						

Note. Extraction Method: Principal Component Analysis; Rotation Method: Oblimin with Kaiser Normalization; Rotation converged in 10 iterations.

Table 26 presents the pattern matrix of the initial solution, before any items were removed. This matrix formed the basis for the item-level decisions. Across the initial solution, most items loaded cleanly on their intended component, but two items behaved problematically and were removed. Throughout the analysis, an item was retained on its primary component when the difference between its highest and next-highest loading exceeded 0.30 (Costello & Osborne, 2005; Hair et al., 2014). Items that violated this criterion were considered for removal.

Table 26 - Pattern matrix, initial solution, Model 1.

Item	Component									
	1	2	3	4	5	6	7	8	9	10
FDI_1									-.713	
FDI_2	.337								-.663	
PDI_1	.727									
PDI_2	.766									
PDI_3	.773									
PDI_4	.762									
PDI_5	.699									
PDI_6	.843									
CER_1					.747					
CER_2					.865					
CER_3					.802					
CER_4					.619					
SER_1		.521								.383
SER_2		.458								.519
SER_3		.429								.548
SER_4										.831
SER_5										.688
SOS_1		.621								
SOS_2		.859								
SOS_3		.825								
ATT_1								.770		
ATT_2								.804		
ATT_3								.758		
ATT_4								.803		
SUN_1					.772					
SUN_2					.809					
SUN_3					.826					
SUN_4					.817					
PBC_1									-.709	
PBC_2									-.838	
PBC_3									-.768	
PA_1			.856							
PA_2			.849							
PA_3			.892							
PA_4			.822							
ASE_1				.824						
ASE_2				.871						
ASE_3				.896						
ASE_4				.646						

Note. Extraction Method: Principal Component Analysis; Rotation Method: Oblimin with Kaiser Normalization; Rotation converged in 29 iterations; items in *italics* were removed.

The first item removed was SER_5. In the initial solution, the severity items did not form a single, clearly separated component, and severity did not cleanly distinguish itself from the swiftness (timing) items. To resolve this, alternative solutions were compared in which individual severity items were removed in turn, and the clearest and most interpretable structure was obtained when SER_5 was excluded (Costello & Osborne, 2005; Howard, 2023). This outcome is consistent with the content of the item. SER_5 states, "If this was discovered, people important to me would lose trust in me," which captures the social and relational consequences of discovery rather than the formal severity of a sanction. Because the remaining severity items refer to formal penalties and academic consequences, SER_5 was conceptually the least representative indicator of the intended construct, which explains why its removal produced the cleanest separation. After SER_5 was removed, the four remaining severity items loaded cleanly on a single component, with loadings of .795 for SER_1, .870 for SER_2, .856 for SER_3, and .738 for SER_4, and no further cross-loadings.

The second item removed was SOS_1. The swiftness construct was intended to capture the speed with which consequences would follow. After SER_5 was removed, however, SOS_1 loaded .541 on Component 2, the same component as the severity items, together with a negative cross-loading of $-.389$ on Component 10, the small, separate timing component on which SOS_2 and SOS_3 loaded. SOS_1 therefore did not belong cleanly to the swiftness construct, since it aligned more closely with severity than with the two remaining timing items.

The behavior of SOS_1 can be traced to its wording. The item refers to the university immediately taking action, which respondents appear to have interpreted as the certainty that enforcement would occur rather than the speed of the process, causing it to overlap with the severity component. SOS_1 was therefore removed. After its removal, the two remaining items loaded distinctly on a single component, with SOS_2 at $-.762$ on Component 10, and SOS_3 at $-.757$ on Component 10.

One other item also showed minor cross-loadings, specifically FDI_2. In this case, the item loaded most strongly on its intended component, and the difference between the magnitude of the primary and secondary loading exceeded 0.30. This item was therefore retained on its primary component, in line with the criterion stated above.

7.2 EFA Model 2: Past behavior

The second model contained the past behavior constructs, measured with two items for past full non-disclosure and six items for past partial non-disclosure. The final solution was reached after three items were removed, as explained below. It retained five items across two components. Table 27 reports the KMO measure (0.887) and Bartlett's test of sphericity ($\chi^2(28) = 1659.24$, $p < .001$), which together confirmed that the data were suitable for factor analysis. Table 28 reports the communalities of the retained items, which ranged from 0.610 to 0.820, indicating that each item was well represented by the extracted components. Table 29 reports the total variance explained. Together, the two components accounted for 73.78% of the variance, which is a satisfactory level for this type of analysis (Hair et al., 2014).

Table 27 – KMO and Bartlett's Test Results for EFA Model 2.

Kaiser-Meyer-Olkin Measure of Sampling Adequacy.			.887
Bartlett's Test of Sphericity	Approx. Chi-Square	1659.239	
	df	28	
	Sig.	.000	

Table 28 – Communalities, final solution, Model 2

Communalities			
		Initial	Extraction
PFD_1	I did not disclose any of my GenAI use when I was expected to.	1.000	.820
PFD_2	I deliberately left out any mention of having used GenAI in my work.	1.000	.817
PPD_1	I disclosed some of my GenAI use but not all of it.	1.000	.610
PPD_5	I only reported minor uses of GenAI but left out more substantial contributions.	1.000	.718
PPD_6	I described my GenAI use in general terms rather than in specific detail.	1.000	.724

Note. Extraction Method: Principal Component Analysis.

Table 29 – Total variance explained, final solution, Model 2

Total Variance Explained							
Component	Initial Eigenvalues			Extraction Sums of Squared Loadings			Rotation Sums of Squared Loadings ^a
	Total	% of Variance	Cumulative %	Total	% of Variance	Cumulative %	
1	2.710	54.193	54.193	2.710	54.193	54.193	2.314
2	.979	19.583	73.776	.979	19.583	73.776	2.105
3	.570	11.400	85.176				
4	.399	7.985	93.161				
5	.342	6.839	100.000				

Note. Extraction Method: Principal Component Analysis. a. When components are correlated, sums of squared loadings cannot be added to obtain a total variance.

Table 30 presents the pattern matrix of the final solution for the second model, after the removal of PPD_2, PPD_3, and PPD_4. Each retained item loaded on its intended component, and the 2 components correspond to the 2 constructs in the theoretical model. Loadings below 0.30 are suppressed for readability. No cross-loadings remained.

Table 30 – Pattern matrix, Final solution, Model 2

Item	Component 1	Component 2
PFD_1		-.898
PFD_2		-.901
PPD_1	.755	
PPD_5	.744	
PPD_6	.901	

Note. Extraction Method: Principal Component Analysis; Rotation Method: Oblimin with Kaiser Normalization; Rotation converged in 5 iterations; PFD = Past Full Disclosure; PPD = Past Partial disclosure; items in *italics* were removed.

Table 31 presents the rotated component matrix of the initial solution, before any items were removed. In this solution, three items showed problematic cross-loadings. Following common practice, the items were removed one at a time, starting with the most problematic and re-running the analysis after each removal (Howard, 2023).

Table 31 – Pattern matrix, initial solution, Model 2

Item	Component 1	Component 2
PFD_1	.890	
PFD_2	.922	
PPD_1		.754
<i>PPD_2</i>	.562	.360
<i>PPD_3</i>	.538	.444
<i>PPD_4</i>	.611	
PPD_5		.698
PPD_6		.893

Note. Extraction Method: Principal Component Analysis; Rotation Method: Oblimin with Kaiser Normalization; Rotation converged in 9 iterations; PFD = Past Full Disclosure; PPD = Past Partial disclosure; items in *italics* were removed.

The first item removed was PPD_3. Although it was intended to measure past partial non-disclosure, it loaded more strongly on Component 1, whereas most of the other partial non-disclosure items loaded more strongly on Component 2. The item was therefore removed, and the analysis was re-run. The second item removed was PPD_4, which continued to load more heavily on Component 1 after the first re-run. Following its removal and a further re-run, PPD_2 showed the same pattern and was removed as well.

After these three removals, the component structure resolved cleanly. Component 1 now corresponded to past partial non-disclosure, and Component 2 to past full non-disclosure. No cross-loadings remained.

Appendix 8 – Regression Assumption: Linearity and Normality

This appendix reports the assessment of the linearity and normality assumptions for the four regression models estimated in this study, namely past full non-disclosure (PFD), past partial non-disclosure (PPD), full non-disclosure intention (FDI) and partial non-disclosure intention (PDI). Linear regression (as used in the hierarchical regression analysis) assumes that the relationship between the predictors and the outcome is linear and that the residuals are approximately normally distributed (Field, 2018). When these conditions hold, estimates and their significance can be interpreted with confidence. Two complementary strategies were used to evaluate this assumption. The first relied on graphical inspection of the residuals, and the second relied on statistical methods. Using both approaches reduces the risk of drawing conclusions from a single, potentially misleading source of evidence (Field, 2018).

8.1 Graphical assessment

The graphical assessment was based on histograms of the standardized residuals and normal probability (P-P) plots for each model. The histogram shows whether the distribution of the residuals approximates the symmetric, bell-shaped form expected under normality. The P-P plot compares the observed cumulative distribution of the residuals against the distribution that would be expected if they were perfectly normal (Field, 2018). Residuals that track the diagonal reference line closely indicate little departure from normality.

Across the four models, the histograms of the residuals were reasonably symmetrical and bell-shaped, with no pronounced skew or heavy tails. The corresponding P-P plots followed the diagonal closely, and the points showed only small and unsystematic deviations from the reference line. The same plots were also inspected for curvature, since a systematic, non-random pattern would point to a violation of linearity. No such pattern was present, which supports the assumption that the relationships between the predictors and each outcome were approximately linear. Taken together, the graphical evidence suggested approximately normal residuals for all models.




Figure 18 – Normal P-P plot and histogram of standardized residuals, PFD model




Figure 21 – Normal P-P plot and histogram of standardized residuals, PPD model




Figure 20 – Normal P-P plot and histogram of standardized residuals, FDI model




Figure 19 – Normal P-P plot and histogram of standardized residuals, PDI model

8.2 Statistical assessment

The graphical inspection was supported by statistical indicators. Skewness and kurtosis were examined first. Skewness describes the symmetry of a distribution, and kurtosis describes the weight of its tails, with values near zero indicating a close approximation to normality (George & Mallery, 2019). As a common guideline, values within the range of ± 1.0 are treated as reflecting only minor departures from normality (George & Mallery, 2019). For all four models, the skewness values ranged from $-.539$ to $.527$ and the kurtosis values ranged from $-.027$ to $.483$. None of these values approached the ± 1.0 boundary, which indicates that the residual distributions were close to normal.

The formal tests of normality, the Kolmogorov-Smirnov test and the Shapiro-Wilk test, were examined next. Both tests check whether the residuals follow a normal distribution. A significant result ($p < .05$) indicates that they do not. For most of the models, both tests were significant, which would point to non-normal residuals. This result should be read with care. Both tests are very sensitive in large samples, where even tiny deviations from normality can turn out significant (Field, 2018). With 431 respondents, a significant result is therefore likely even when the actual departure from normality is very small. For this reason, the formal tests were not treated as decisive on their own but were weighed against the graphical evidence and the skewness and kurtosis values.

Table 32 – Skewness and kurtosis of standardized residuals

	Skewness		Kurtosis	
	Statistic	Std. Error	Statistic	Std. Error
PFD_res	.527	.140	.286	.279
PPD_res	-.382	.140	.115	.279
FDI_res	.387	.140	-.027	.279
PDI_res	-.539	.140	.483	.279

Note. $n = 431$, PFD_res = Past Full Disclosure residuals; PPD_res = Past Partial Disclosure residuals; FDI_res = Future Disclosure Intention residuals; PDI_res = Past Disclosure Intention residuals. All residual variables represent standardized residuals obtained from their respective regression models.

Table 33 – Kolmogorov-Smirnov and Shapiro-Wilk tests of normality

	Kolmogorov-Smirnov ^a			Shapiro-Wilk		
	Statistic	df	Sig.	Statistic	df	Sig.
PFD_res	.058	431	.015	.983	431	.001
PPD_res	.048	431	.091	.990	431	.030
FDI_res	.064	431	.004	.987	431	.008
PDI_res	.067	431	.002	.981	431	<.001

Note. $n = 431$; a. Lilliefors Significance Correction; PFD_res = Past Full Disclosure residuals; PPD_res = Past Partial Disclosure residuals; FDI_res = Future Disclosure Intention residuals; PDI_res = Past Disclosure Intention residuals. All residual variables represent standardized residuals obtained from their respective regression models.

8.3 Conclusion

The three sources of evidence pointed in the same direction. The histograms and P-P plots showed roughly normal and linear residual patterns. The skewness and kurtosis values stayed well within the accepted range for all four models. The significant results of the Kolmogorov-Smirnov and Shapiro-Wilk tests were most likely due to the sensitivity of these tests in large samples, rather than a real problem with normality. In addition, linear regression copes well with mild non-normality when the sample is large, so any small departures were unlikely to distort the estimates at a sample size of 431 (Hair et al., 2014). Overall, the residuals were considered close enough to normal and the relationships close enough to linear. The linearity and normality assumption was therefore regarded as satisfied across all four models.

Appendix 9 – Regression Assumption: Homoscedasticity

This appendix reports the assessment of the homoscedasticity assumption for the four regression models: past full non-disclosure (PFD), past partial non-disclosure (PPD), full non-disclosure intention (FDI) and partial non-disclosure intention (PDI). Homoscedasticity means that the variance of the residuals stays roughly constant across all levels of the predicted values (Field, 2018). If this assumption is violated, the spread of the residuals changes systematically, which can distort the standard errors and the significance tests.

The assumption was checked by plotting the standardized residuals (ZRESID on the Y-axis) against the standardized predicted values (ZPRED on the X-axis) for each model. When the assumption holds, the points scatter randomly around zero and form a fairly even band, with no clear shape or trend (Field, 2018). A funnel shape or any other systematic pattern would instead signal a problem.

For all four models, the residuals were spread evenly across the range of predicted values. The points clustered around zero without widening, narrowing, or forming a recognizable pattern. This indicates that the variance of the residuals was stable and that the assumption of constant variance was met. The homoscedasticity assumption was therefore regarded as satisfied across all four models. The corresponding scatterplots are shown in Figure 22 to Figure 25 (p. 124 to 125).




Figure 22 – Residual scatterplot, PFD model




Figure 23 – Residual scatterplot, PPD model




Figure 24 – Residual scatterplot, FDI model




Figure 25 – Residual scatterplot, PDI model

Appendix 10 – Regression Assumption: Multicollinearity

This appendix reports the assessment of the multicollinearity assumption for the four hierarchical regression models: past full non-disclosure (PFD), past partial non-disclosure (PPD), full non-disclosure intention (FDI), and partial non-disclosure intention (PDI). Multicollinearity occurs when predictors are strongly correlated with one another, which makes it difficult to separate the individual effect of each predictor and can inflate the standard errors and destabilize the estimates (Field, 2018).

The assumption was checked using two related indicators, the variance inflation factor (VIF) and tolerance. The VIF shows how strongly each predictor is related to the others, while tolerance expresses the same information on a 0 to 1 scale. Following Field (2018), a VIF above 10, or an average VIF well above 1, would signal a possible problem, as would a tolerance value below 0.2. The statistics were read from the full model, in which all eight predictors are entered together (Model 3). Because the four models share this same predictor set, their collinearity statistics are identical, since these values depend only on the relationships among the predictors and not on the outcome.

For all four models, tolerance values ranged from 0.617 to 0.929 and VIF values from 1.076 to 1.620. All tolerance values stayed well above 0.2, and no VIF came close to 10. The average VIF across the predictors was approximately 1.36, which further indicates that the estimates were not distorted by collinearity (Field, 2018).

Taken together, these results indicate that the predictors were not strongly correlated with one another. The multicollinearity assumption was therefore regarded as satisfied across all four models. The full collinearity statistics are reported in Table 34.

Table 34 – Collinearity statistics, non-disclosure models

Collinearity Statistics ^a		
Variable	Tolerance	VIF
CER	.893	1.120
SER	.627	1.596
SOS	.645	1.551
ATT	.617	1.620
SUN	.652	1.533
PBC	.795	1.258
PA	.929	1.076
ASE	.893	1.120

Note. a. for all four disclosure models, $n = 431$, CER = Certainty, SER = Severity, SOS = Swiftiness, ATT = Attitude, SUN = Subjective Norms, PBC = Perceived Behavioral Control, PA = Policy awareness, ASE = Academic Self-Efficacy.

Appendix 11 – Regression Assumption: Independence

This appendix reports the assessment of the independence assumption for the four regression models: past full non-disclosure (PFD), past partial non-disclosure (PPD), full non-disclosure intention (FDI), and partial non-disclosure intention (PDI). Independence means that the residuals are not correlated with one another (Field, 2018). If the residuals are related, the model may understate the standard errors and overstate the significance of the results.

The assumption was checked using the Durbin-Watson statistic, which tests for correlation between neighboring residuals. Values close to 2.0 indicate no autocorrelation, while values below 1.0 or above 3.0 are a cause for concern (Field, 2018).

For the four models, the Durbin-Watson values ranged from 1.823 to 2.065. All values fell close to 2.0 and well within the acceptable range. This indicates that the residuals were independent in each model. The independence assumption was therefore regarded as satisfied across all models. The Durbin-Watson values, together with the model summaries, are reported in Table 35 to Table 38.

Table 35 – Model summary and Durbin-Watson statistic, PFD model

Model	R	R Square	Adjusted R Square	Std. Error of the Estimate	Durbin-Watson
1	.559	.313	.300	1.01192	1.823

Note. Dependent Variable = PFD, Predictors: CER, SER, SOS, ATT, SUN, PBC, PA, ASE.

Table 36 – Model summary and Durbin-Watson statistic, PPD model

Model	R	R Square	Adjusted R Square	Std. Error of the Estimate	Durbin-Watson
1	.522	.272	.258	.95381	1.973

Note. Dependent Variable = PPD, Predictors: CER, SER, SOS, ATT, SUN, PBC, PA, ASE.

Table 37 – Model summary and Durbin-Watson statistic, FDI model

Model	R	R Square	Adjusted R Square	Std. Error of the Estimate	Durbin-Watson
1	.557	.310	.297	1.03757	2.065

Note. Dependent Variable = FDI, Predictors: CER, SER, SOS, ATT, SUN, PBC, PA, ASE.

Table 38 – Model summary and Durbin-Watson statistic, PDI model

Model	R	R Square	Adjusted R Square	Std. Error of the Estimate	Durbin-Watson
1	.638	.407	.396	.82800	1.875

Note. Dependent Variable = PDI, Predictors: CER, SER, SOS, ATT, SUN, PBC, PA, ASE.

Appendix 12 – Comparing Predictors Across the Forms of Non-Disclosure

This appendix reports the full set of tests behind the analysis of Hypothesis 2 in Section 4.2.1. Hypothesis 2 concerns whether the predictors of non-disclosure differ across the four outcomes. Because the four outcomes were measured in the same respondents, their coefficients were not independent. The outcomes were therefore analyzed jointly rather than in four separate regressions.

The four equations were therefore estimated together as one system using seemingly unrelated regression, with all eight predictors entered in each equation (Zellner, 1962). Because the four equations share the same predictors, the SUR coefficients are identical to those from the separate regressions reported in Table 7 to Table 10 (p. 44 & 45). The advantage of the joint model is that it accounts for the correlations between the outcomes, allowing valid statistical comparisons of coefficients across outcomes.

Coefficient equality was tested with Wald tests (Zellner, 1962). Four families of tests were run:

- A global test examining whether all predictor coefficients were equal across the four outcomes.
- Predictor-specific tests examining whether each individual predictor had the same coefficient across the four outcomes.
- Full versus partial non-disclosure contrasts examining whether each predictor differed between full and partial non-disclosure, averaged across past behavior and intention.
- Difference-in-differences tests examining whether the difference between full and partial non-disclosure varied between past behavior and intention.

Because all predictors share a common 1 to 5 response scale, their coefficients are directly comparable. To account for multiple testing, Holm-adjusted p-values are reported alongside the raw values, and effects that are significant only before correction are treated as indicative.

The global test was significant, so the predictors of non-disclosure differ across its forms. When each predictor was examined separately, the differences were found only for the deterrence perceptions. The psycho-social and individual-level predictors showed similar effects across all four outcomes. Averaged across time, the full and partial forms were not driven by different predictors. The one effect to survive correction for multiple testing was for severity, whose full-versus-partial gap differed between past behavior and intention. The remaining differences were significant only before correction and are treated as indicative.

Table 39 – Global test of coefficient equality across all four outcomes

Test	χ^2	df	<i>p</i>
All predictors equal across all four outcomes	39.46	24	.024

Note. *n* = 431. A significant result indicates that at least one predictor differs across the forms of non-disclosure.

Table 40 – Per-predictor tests of coefficient equality across the four outcomes

Predictor	χ^2	df	p	p (Holm)
Certainty	6.31	3	.097	.585
Severity	9.30	3	.026	.205
Swifttness	8.81	3	.032	.224
Attitude	3.88	3	.275	.826
Subjective norms	5.82	3	.121	.605
Perceived behavioral control	4.88	3	.181	.724
Academic self-efficacy	0.47	3	.925	1.000
Policy awareness	2.07	3	.559	1.000

Note. n = 431, Each row tests whether that predictor's coefficient is equal across the four outcomes.

Table 41 – Averaged full-versus-partial contrast, by predictor

Predictor	Diff.	95% CI	p	p (Holm)
Certainty	-.130	[-.232, -.027]	.013	.105
Severity	-.018	[-.140, .105]	.778	1.000
Swifttness	.067	[-.039, .173]	.215	.932
Attitude	.092	[-.020, .203]	.108	.648
Subjective norms	.006	[-.118, .130]	.923	1.000
Perceived behavioral control	-.116	[-.241, .009]	.068	.479
Academic self-efficacy	.005	[-.118, .128]	.939	1.000
Policy awareness	.064	[-.031, .160]	.186	.932

Note. n = 431. Diff. = mean(full) – mean(partial), pooled across past behavior and intention, on the 1 to 5 scale.

Table 42 – Difference-in-Differences: Full vs. Partial Across Past and Intention

Predictor	Diff.	95% CI	p	p (Holm)
Certainty	.002	[-.119, .123]	.980	1.000
Severity	.218	[.074, .363]	.003	.025
Swifttness	-.150	[-.275, -.025]	.019	.132
Attitude	.083	[-.049, .215]	.219	1.000
Subjective norms	.062	[-.084, .208]	.407	1.000
Perceived behavioral control	.060	[-.087, .208]	.422	1.000
Academic self-efficacy	.033	[-.113, .178]	.658	1.000
Policy awareness	.007	[-.105, .120]	.900	1.000

Note. n = 431, Diff. = (full – partial | past) – (full – partial | intention). A confidence interval excluding zero indicates that the full-versus-partial gap differs between past behavior and intention.

Appendix 13 – Theory individual models

The hierarchical regression in Chapter 4.3 entered the predictors in blocks and reported the full model. This appendix presents the same data in a different way, as three standalone models per outcome. One model contains only the deterrence perceptions, one contains only the Theory of Planned Behavior constructs, and one combines both sets. This analysis will help to make the theoretical comparison explicit. The hierarchical regression already hints at it, but separate models show more directly how much each framework explains on its own, and what is gained by combining them. There is therefore some overlap with Chapter 4.3. The models are repeated here in full for transparency and ease of interpretation.

Each model was estimated for all four outcomes, giving twelve regressions in total. The deterrence model used certainty, severity, and swiftness. The Theory of Planned Behavior model used attitude, subjective norms, and perceived behavioral control. The combined model used all six. The two individual-level predictors, academic self-efficacy and policy awareness, were not included, so the combined model here is not identical to the final models in Chapter 4.2. All models were estimated on the full sample ($n = 431$).

Explained variance across the three models

The clearest result is the difference in explanatory power. Table 43 reports the model fit. The deterrence model explained very little of the variance in non-disclosure, with R^2 between .028 and .066. The Theory of Planned Behavior model explained far more, with R^2 between .239 and .376. The combined model barely improved on it. The gain in R^2 ranged from .008 to .029, and only for full non-disclosure intention did it reach a noteworthy size. Almost all of the explained variance therefore came from the psycho-social constructs. Adding the deterrence perceptions contributed little beyond them.

Deterrence theory model

On its own, the deterrence model produced a few significant effects, but the overall fit was weak (Table 44). Certainty was negatively related to the two full outcomes. Severity was negatively related to three of the four outcomes. Swiftness was significant only for full non-disclosure intention, and there in the positive direction, against the deterrence prediction. These effects are consistent with the correlations in Section 4.1. They show that the deterrence perceptions are not unrelated to non-disclosure. Their associations are simply small.

Theory of Planned Behavior model

The Theory of Planned Behavior model was much stronger (Table 45). Attitude was the dominant predictor for every outcome, with standardized coefficients between .360 and .440. Subjective norms were significant across all four outcomes as well. Perceived behavioral control was significant for the two partial outcomes only. This mirrors the pattern found in the full models in Chapter 4.3.

Combined model

When both sets were entered together, the picture changed in a telling way (Table 46). The Theory of Planned Behavior constructs kept their strength, while the deterrence perceptions lost almost all of theirs. Certainty and severity were no longer significant for any outcome. Only swiftness remained significant, and again only for full non-disclosure intention. The deterrence effects seen in the deterrence-only model were therefore largely absorbed once the psycho-social constructs were included. The exploratory analysis in Appendix 13 offers a likely reason. The effects of certainty and severity ran mainly through attitude and subjective norms, so they weakened once those constructs were in the model.

Why this comparison matters

Together, the three models answer a question that the combined model alone cannot. They show that the enforcement environment, as students perceived it, accounts for only a small share of non-disclosure. In contrast, how students evaluate the behavior accounts for a much larger share. Combining the two theories added little explained variance, but it clarified the structure of the effects. Deterrence perceptions matter, but their influence appears small and largely indirect, rather than a strong and direct deterrent. This supports the study's broader argument that non-disclosure is driven more by psycho-social evaluation than by the threat of sanctions.

Table 43 – Model fit for the deterrence, TPB, and combined regression models

Outcome	Deterrence R ² (adj.)	TPB R ² (adj.)	Combined R ² (adj.)
Past full non-disclosure	.041 (.034)	.294 (.289)	.304 (.295)
Past partial non-disclosure	.028 (.021)	.239 (.233)	.247 (.237)
Full non-disclosure intention	.066 (.060)	.272 (.267)	.301 (.291)
Partial non-disclosure intention	.035 (.028)	.376 (.371)	.385 (.377)

Note. n = 431. R² = explained variance; adj. = adjusted R². All twelve models were significant (p ≤ .007).

Table 44 – Deterrence theory regression model

Predictor	Past full	Past partial	Full intention	Partial intention
Certainty	-.157**	-.043	-.168***	-.092
Severity	-.107	-.164**	-.217***	-.143*
Swiftness	.011	.024	.131*	-.002
R ² (adj.)	.041 (.034)	.028 (.021)	.066 (.060)	.035 (.028)
F(3, 427)	6.09***	4.05**	10.13***	5.16**

Note. n = 431. β = standardized coefficient. * p < .05, ** p < .01, *** p < .001. Swiftness was positive for full non-disclosure intention, against the hypothesized direction.

Table 45 – Theory of Planned Behavior regression model

Predictor	Past full	Past partial	Full intention	Partial intention
Attitude	.440***	.360***	.387***	.381***
Subjective norms	.127*	.120*	.178***	.229***
Perceived behavioral control	.054	.106*	.038	.143***
R ² (adj.)	.294 (.289)	.239 (.233)	.272 (.267)	.376 (.371)
F(3, 427)	59.35***	44.61***	53.20***	85.69***

Note. n = 431. β = standardized coefficient. * p < .05, ** p < .01, *** p < .001.

Table 46 – Combined deterrence and TPB regression model

Predictor	Past full	Past partial	Full intention	Partial intention
Certainty	-.050	.060	-.068	.039
Severity	.035	-.037	-.081	.020
Swiftness	.079	.085	.197***	.078
Attitude	.454***	.376***	.393***	.406***
Subjective norms	.144**	.132*	.192***	.249***
Perceived behavioral control	.049	.114*	.030	.148***
R ² (adj.)	.304 (.295)	.247 (.237)	.301 (.291)	.385 (.377)
F(6, 424)	30.92***	23.24***	30.37***	44.30***

Note. n = 431. β = standardized coefficient. * p < .05, ** p < .01, *** p < .001. This model excludes academic self-efficacy and policy awareness, so its coefficients differ slightly from the eight-predictor models in Chapter 4.

Appendix 14 – Exploratory Analysis of Indirect Effects

The regression models in Chapter 4.2 estimated the direct associations between the predictors and the four non-disclosure outcomes. This appendix adds one further question. Where a predictor is related to non-disclosure, does that relationship run partly through the psycho-social constructs of the Theory of Planned Behavior?

All constructs were measured at a single point in time. The analysis therefore cannot establish the temporal order that an indirect pathway implies. It is treated as exploratory. The results show whether the observed associations are consistent with an indirect pathway, not that such a pathway operates over time (Bryman, 2016). For this reason, no formal hypotheses are attached to the analysis. The findings are read as associational rather than causal.

Analytical approach

Indirect effects were estimated with the PROCESS macro for SPSS, using Model 4 for parallel multiple mediation (Hayes, 2013). Each predictor was entered as the independent variable in a separate model. Attitude, subjective norms, and perceived behavioral control were entered together as parallel mediators. The remaining predictors were kept in the outcome model as covariates, so the specification matched the main regression in Chapter 4. Every model was run against all four outcomes: past full non-disclosure, past partial non-disclosure, full non-disclosure intention, and partial non-disclosure intention. Bias-corrected bootstrap confidence intervals were based on 5,000 resamples. An indirect effect was treated as significant when its 95% interval did not include zero (Hayes, 2013). The full results appear in Table 47 to Table 51.

Five predictors were examined: certainty of sanctions, severity of sanctions, swiftness of sanctions, policy awareness, and academic self-efficacy. Three Theory of Planned Behavior constructs were considered as possible channels. Perceived behavioral control was included for completeness, not on strong theoretical grounds. As explained in Section 3.5, its items concern control over one's own reporting, rather than an evaluative or normative response to sanctions, which is why it was not treated as a channel in the main analysis. It is tested here so that all three Theory of Planned Behavior constructs are examined and reported, and so that its limited role is documented rather than assumed. Figure 26 shows the relationships examined here.




Figure 26 - Conceptual model, tested relationships, exploratory mediation analysis

Certainty of sanctions

Certainty of sanctions reached non-disclosure mainly through indirect paths. The indirect effect through attitude was negative and significant for all four outcomes (indirect effects from $-.074$ to $-.101$). The indirect effect through subjective norms was also significant throughout, though smaller (from $-.019$ to $-.037$). All bootstrap intervals excluded zero. Perceived behavioral control carried less. It transmitted a significant effect only for the two partial outcomes ($-.027$ and $-.032$), not for the full ones. Even there, the effect only just reached significance. Both confidence intervals stopped just short of zero at the upper end ($[-.061, -.004]$ and $[-.066, -.009]$). The path is therefore weak and should be read with caution. It points to a possible role for perceived behavioral control in selective under-reporting, rather than firm evidence of one. The direct effect of certainty was non-significant in every model (c' from $-.100$ to $.055$).

Students who saw detection as more likely evaluated non-disclosure less favorably. They also perceived less peer approval of it. Both perceptions were, in turn, linked to less concealment.

Severity of sanctions

Severity showed the clearest indirect pattern. Its effect ran through attitude for all four outcomes (indirect effects from $-.093$ to $-.126$). It also ran through subjective norms throughout (from $-.030$ to $-.057$). Every bootstrap interval excluded zero. Perceived behavioral control showed no significance, since all its intervals included zero. The direct effect of severity was non-significant in each model. This is worth noting. The direct regression in Chapter 4.2 found no significant effect of severity. Here, severity was clearly related to non-disclosure, but only through how students evaluated the behavior and what they believed their peers accept (Nguyen-Phuoc et al., 2024).

Swiftness of sanctions

Swiftness produced no significant indirect effects. Every indirect path, through attitude, subjective norms, and perceived behavioral control, was small and non-significant (indirect effects roughly $-.002$ to $-.049$, all intervals including zero). Its associations with non-disclosure were direct instead. Swiftness kept a significant positive direct effect on the two intention outcomes (full intention $c' = .234$, $p < .001$; partial intention $c' = .092$, $p < .05$). Its direct effects on the two past-behavior outcomes were non-significant. The positive sign is unexpected. Under a deterrence logic, faster enforcement should reduce concealment, not raise it. This mirrors the direct regression and is taken up in the discussion. The point here is simpler: swiftness did not work through any psycho-social channel.

Policy awareness

Policy awareness also showed no significant indirect effects. None of the three paths were significant, as all confidence intervals included zero. Its relationship with non-disclosure was direct. It was also limited to the two partial outcomes, where it carried a significant negative direct effect (past partial $c' = -.113$, $p < .05$; partial intention $c' = -.120$, $p < .01$). Awareness

of the rules was linked to less selective under-reporting. That link did not run through attitude or subjective norms.

Academic self-efficacy

Academic self-efficacy worked mainly through attitude. The indirect effect through attitude was negative and significant for all four outcomes (from $-.103$ to $-.140$, all intervals excluding zero). Subjective norms and perceived behavioral control did not reach significance. The pattern differed by outcome type. For the two full outcomes, the direct effect was non-significant, so the association ran entirely through attitude. For the two partial outcomes, a significant negative direct effect remained (past partial $c' = -.148$, $p < .05$; partial intention $c' = -.111$, $p < .05$). Attitude therefore explained part, but not all, of the relationship there. Students with more confidence in their academic ability held a less favorable attitude toward non-disclosure. That attitude was, in turn, linked to less concealment.

Overview

Together, the results point to a clear structure. Attitude was the main channel. It carried certainty, severity, and academic self-efficacy across the outcomes. Subjective norms formed a second channel, but a narrower one. It carried certainty and severity only, not academic self-efficacy. Perceived behavioral control played almost no role. Where it did, the effect was weak. It transmitted just one path, for certainty, and only on the partial outcomes. Even there, both intervals only just cleared zero. Swiftiness and policy awareness did not work through any channel. Where they mattered, their effect was direct. This fits the study's integration logic (Nguyen-Phuoc et al., 2024). Deterrence perceptions that carry evaluative or cost content reach non-disclosure partly through students' evaluations. Perceptions about the timing of enforcement, or about knowledge of the rules, do not. Figure 27 summarizes which paths held and which did not. The dashed line from certainty to perceived behavioral control marks the one path that held for the partial outcomes only, rather than across all four.




Figure 27 – Conceptual model, showing which indirect pathways were supported

Table 47 – Indirect effects of certainty of sanctions on non-disclosure

Certainty of sanctions			
Outcome / mediator	Indirect effect	95% CI [LL, UL]	Result
<i>Past full non-disclosure</i> ($c = -.213$; $c' = -.074$)			
Attitude	-.101*	[-.163, -.048]	Full mediation
Subjective norms	-.025*	[-.056, -.003]	Full mediation
Perceived behavioral control	-.014	[-.044, .008]	No mediation
<i>Past partial non-disclosure</i> ($c = -.065$; $c' = .055$)			
Attitude	-.074*	[-.125, -.034]	Full mediation
Subjective norms	-.019*	[-.045, -.002]	Full mediation
Perceived behavioral control	-.027*	[-.061, -.004]	Full mediation
<i>Full non-disclosure intention</i> ($c = -.233$; $c' = -.100$)			
Attitude	-.089*	[-.144, -.041]	Full mediation
Subjective norms	-.034*	[-.068, -.008]	Full mediation
Perceived behavioral control	-.009	[-.037, .012]	No mediation
<i>Partial non-disclosure intention</i> ($c = -.118$; $c' = .030$)			
Attitude	-.079*	[-.129, -.037]	Full mediation
Subjective norms	-.037*	[-.071, -.009]	Full mediation
Perceived behavioral control	-.032*	[-.066, -.009]	Full mediation

Note. $n = 431$. Coefficients are unstandardized (PROCESS Model 4; Hayes, 2013); Bootstrap confidence intervals based on 5,000 samples; an indirect effect is significant (*) when its interval excludes zero. c = total effect ($c' +$ total indirect); c' = direct effect controlling for all mediators.

Table 48 – Indirect effects of severity of sanctions on non-disclosure

Severity of Sanctions			
Outcome / mediator	Indirect effect	95% CI [LL, UL]	Result
<i>Past full non-disclosure</i> ($c = -.131$; $c' = .041$)			
Attitude	-.126*	[-.196, -.059]	Full mediation
Subjective norms	-.037*	[-.072, -.008]	Full mediation
Perceived behavioral control	-.008	[-.028, .005]	No mediation
<i>Past partial non-disclosure</i> ($c = -.188$; $c' = -.051$)			
Attitude	-.093*	[-.150, -.043]	Full mediation
Subjective norms	-.030*	[-.061, -.005]	Full mediation
Perceived behavioral control	-.015	[-.040, .003]	No mediation
<i>Full non-disclosure intention</i> ($c = -.280$; $c' = -.112$)			
Attitude	-.111*	[-.176, -.052]	Full mediation
Subjective norms	-.052*	[-.092, -.019]	Full mediation
Perceived behavioral control	-.005	[-.024, .007]	No mediation

Note. The table is continued on the next page.

Table 48 Continued

Partial non-disclosure intention ($c = -.158$; $c' = .015$)

Attitude	-.098*	[-.153, -.047]	Full mediation
Subjective norms	-.057*	[-.098, -.024]	Full mediation
Perceived behavioral control	-.018	[-.046, .004]	No mediation

Note. $n = 431$. Coefficients are unstandardized (PROCESS Model 4; Hayes, 2013); Bootstrap confidence intervals based on 5,000 samples; an indirect effect is significant (*) when its interval excludes zero. c = total effect (c' + total indirect); c' = direct effect controlling for all mediators.

Table 49 – Indirect effects of swiftness of sanctions on non-disclosure

Swiftness of Sanctions			
Outcome / mediator	Indirect effect	95% CI [LL, UL]	Result
<i>Past full non-disclosure</i> ($c = .025$; $c' = .094$)			
Attitude	-.049	[-.106, .004]	No mediation
Subjective norms	-.017	[-.045, .002]	No mediation
Perceived behavioral control	-.003	[-.019, .007]	No mediation
<i>Past partial non-disclosure</i> ($c = .046$; $c' = .101$)			
Attitude	-.036	[-.082, .002]	No mediation
Subjective norms	-.013	[-.038, .002]	No mediation
Perceived behavioral control	-.006	[-.028, .012]	No mediation
<i>Full non-disclosure intention</i> ($c = .166$; $c' = .234^*$)			
Attitude	-.043	[-.096, .003]	Direct only
Subjective norms	-.023	[-.058, .002]	Direct only
Perceived behavioral control	-.002	[-.015, .007]	Direct only
<i>Partial non-disclosure intention</i> ($c = .022$; $c' = .092^*$)			
Attitude	-.038	[-.085, .001]	Direct only
Subjective norms	-.025	[-.058, .003]	Direct only
Perceived behavioral control	-.007	[-.030, .013]	Direct only

Note. $n = 431$. Coefficients are unstandardized (PROCESS Model 4; Hayes, 2013); Bootstrap confidence intervals based on 5,000 samples; an indirect effect is significant (*) when its interval excludes zero. c = total effect (c' + total indirect); c' = direct effect controlling for all mediators.

Table 50 – Indirect effects of policy awareness on non-disclosure

Policy Awareness			
Outcome / mediator	Indirect effect	95% CI [LL, UL]	Result
<i>Past full non-disclosure</i> ($c = -.059$; $c' = -.045$)			
Attitude	-.002	[-.048, .046]	No mediation
Subjective norms	-.010	[-.031, .005]	No mediation
Perceived behavioral control	-.002	[-.014, .007]	No mediation
<i>Past partial non-disclosure</i> ($c = -.126$; $c' = -.113^*$)			
Attitude	-.001	[-.036, .033]	Direct only
Subjective norms	-.008	[-.026, .005]	Direct only
Perceived behavioral control	-.004	[-.020, .011]	Direct only
<i>Full non-disclosure intention</i> ($c = -.077$; $c' = -.060$)			
Attitude	-.001	[-.043, .041]	No mediation
Subjective norms	-.014	[-.041, .007]	No mediation
Perceived behavioral control	-.001	[-.011, .006]	No mediation
<i>Partial non-disclosure intention</i> ($c = -.142$; $c' = -.120^*$)			
Attitude	-.001	[-.037, .035]	Direct only
Subjective norms	-.016	[-.044, .008]	Direct only
Perceived behavioral control	-.005	[-.023, .013]	Direct only

Note. $n = 431$. Coefficients are unstandardized (PROCESS Model 4; Hayes, 2013); Bootstrap confidence intervals based on 5,000 samples; an indirect effect is significant (*) when its interval excludes zero. c = total effect (c' + total indirect); c' = direct effect controlling for all mediators.

Table 51 – Indirect effects of academic self-efficacy on non-disclosure

Academic Self-Efficacy			
Outcome / mediator	Indirect effect	95% CI [LL, UL]	Result
<i>Past full non-disclosure</i> ($c = -.282$; $c' = -.127$)			
Attitude	-.140*	[-.215, -.073]	Full mediation
Subjective norms	-.019	[-.047, .001]	No mediation
Perceived behavioral control	.004	[-.005, .020]	No mediation
<i>Past partial non-disclosure</i> ($c = -.259$; $c' = -.148^*$)			
Attitude	-.103*	[-.164, -.052]	Partial mediation
Subjective norms	-.015	[-.038, .001]	Direct only
Perceived behavioral control	.007	[-.010, .028]	Direct only
<i>Full non-disclosure intention</i> ($c = -.269$; $c' = -.122$)			
Attitude	-.123*	[-.190, -.063]	Full mediation
Subjective norms	-.026	[-.057, .002]	No mediation
Perceived behavioral control	.003	[-.006, .016]	No mediation

Note. The table is continued on the next page.

Table 51 Continued

Partial non-disclosure intention ($c = -.240$; $c' = -.111^*$)

Attitude	-.109*	[-.168, -.059]	Partial mediation
Subjective norms	-.028	[-.062, .002]	Direct only
Perceived behavioral control	.008	[-.012, .030]	Direct only

Note. $n = 431$. Coefficients are unstandardized (PROCESS Model 4; Hayes, 2013); Bootstrap confidence intervals based on 5,000 samples; an indirect effect is significant (*) when its interval excludes zero. c = total effect ($c' +$ total indirect); c' = direct effect controlling for all mediators.

Appendix 15 – Robustness Check with Demographic Controls

This appendix reports whether the study's findings held when age, gender, and level of study were included as control variables.

Purpose

The sample was uneven in its composition, leaning toward master's students and containing a majority of women (see Section 3.2.1). This raises a fair question. Could the results reflect who was in the sample rather than the relationships the study proposes? To check this, the combined model was re-estimated with demographic controls added. A robustness check of this kind does not look for new effects. It tests whether the main predictors survive once composition is accounted for.

Analysis

For each of the four outcomes, a hierarchical regression was run. The eight main predictors were entered in the first block, and four demographic controls were entered in the second: age, gender (0 = male, 1 = female), and two dummy variables for level of study, with master's as the reference category (bachelor, and a combined PhD-and-other group). Gender was reduced to a binary variable, and the small "other" and "prefer not to say" groups were set to missing, which left 418 cases. The key figure is the change in explained variance (ΔR^2) when the demographic block was added.

Results

Adding the demographic block produced only a small increase in explained variance. The change ranged from $\Delta R^2 = .013$ to $.019$ across the four outcomes, as reported in Table 52. Three of the four increases were statistically significant, but all were negligible in size, adding at most two percentage points of explained variance to models that already accounted for 26% to 40%.

More importantly, the main predictors were essentially unchanged. Attitude, subjective norms, and perceived behavioral control kept their coefficients and their significance in every model. Academic self-efficacy and policy awareness likewise retained their pattern, remaining significant for the two partial outcomes. No main predictor gained or lost significance when the controls were added.

Among the controls themselves, only age reached significance, and only for some outcomes. It was negatively related to past full non-disclosure ($\beta = -.097$, $p = .033$) and to partial non-disclosure intention ($\beta = -.124$, $p = .003$), and approached significance for the other two. Older students reported slightly less non-disclosure, but the effect was small. Gender was non-significant in all four models, which is itself informative. Level of study was non-significant throughout.

Table 52 – Change in explained variance when demographic controls were added

Outcome	Block 1 R ² (main predictors)	Block 2 R ² (with controls)	ΔR^2	p (ΔF)
Past full non-disclosure	.306	.323	.017	.042
Past partial non-disclosure	.261	.274	.013	.135
Full non-disclosure intention	.301	.320	.019	.024
Partial non-disclosure intention	.399	.418	.019	.010

Note. n = 418. R² = explained variance; ΔR^2 = increase from adding the demographic block (age, gender, and two study-level dummies).

Table 53 – Demographic controls in the full model

Control	Past full	Past partial	Full intention	Partial intention
Age	-.097*	-.090	-.086	-.124**
Gender (female)	.054	.009	.064	.043
Bachelor (vs master's)	.047	-.002	.078	.031
Other (vs master's)	-.012	-.058	.048	-.002

Note. n = 418. β = standardized coefficient from the full model (Block 2). Gender coded 0 = male, 1 = female. Study level dummies use master's as the reference. * p < .05, ** p < .01.

Summary

The findings were robust to demographic composition. Adding age, gender, and level of study left the main predictors intact and explained only a marginal amount of additional variance. The one consistent control effect was age, where older students reported slightly less non-disclosure, though the effect was small. The absence of a gender effect is worth noting. It suggests that non-disclosure is shaped by how students evaluate the behavior rather than by demographic group, which is consistent with the study's wider argument.

Appendix 16 – Response Language

This appendix reports a supplementary analysis of whether students' non-disclosure differed between those who completed the survey in Dutch and those who completed it in English. It also examines whether the factors associated with non-disclosure varied between these groups.

Purpose

The survey was offered in Dutch and English, and respondents chose one. This raised a particular question. Do Dutch and English respondents differ in their non-disclosure, and are the drivers of non-disclosure the same in both groups? This appendix reports two supplementary analyses that address this. It sits outside the hypotheses and is exploratory.

One caveat frames everything that follows. The language of response is not a reliable indicator of nationality. Some Dutch students may have chosen the English survey, so any difference reflects the language people responded in, not their country or culture (see Section 3.2). The results should be read in that light.

Group differences on the study constructs

An independent-samples t-test compared the two groups on all twelve constructs (Table 54). The groups were similar on most of them. Eight of the twelve showed no significant difference, including attitude, subjective norms, and perceived behavioral control. The clearest difference was in policy awareness. English respondents scored well above Dutch respondents, and this was the only medium-sized effect ($d = .65$). Dutch respondents reported somewhat more partial non-disclosure, and English respondents rated severity a little higher. All of these remaining effects were small ($d \leq .29$).

Table 54 - Group Differences by Survey Language (Independent-Samples t-Tests)

Construct	English M (SD)	Dutch M (SD)	t	p	d
Past full non-disclosure	2.48 (1.28)	2.57 (1.09)	-0.79	.429	-.08
Past partial non-disclosure	2.69 (1.13)	3.00 (1.05)	-2.90	.004 **	-.29
Full non-disclosure intention	2.56 (1.26)	2.47 (1.20)	0.75	.454	.07
Partial non-disclosure intention	3.06 (1.09)	3.29 (1.01)	-2.18	.029 *	-.22
Certainty	2.94 (1.02)	2.79 (0.81)	1.76	.080	.17
Severity	3.78 (1.00)	3.51 (0.83)	2.98	.003 **	.28
Swiftness	3.31 (1.13)	3.14 (0.98)	1.68	.093	.16
Attitude	2.72 (1.13)	2.77 (0.88)	-0.52	.602	-.05
Subjective norms	3.54 (0.99)	3.63 (0.78)	-1.00	.318	-.09
Perceived behavioral control	3.64 (0.88)	3.63 (0.73)	0.17	.864	.02
Academic self-efficacy	4.10 (0.82)	3.96 (0.72)	1.85	.065	.18
Policy awareness	3.79 (0.95)	3.17 (0.94)	6.56	< .001 ***	.65

Note. $n = 431$ (English = 264, Dutch = 167). M = mean; SD = standard deviation; d = Cohen's d. Two-sided p reported. Welch's correction (equal variances not assumed) was applied where Levene's test was significant. * $p < .05$, ** $p < .01$, *** $p < .001$.

Do the drivers differ by language?

To see whether the predictors worked differently in the two groups, the full eight-predictor regression was run separately for each group. Table 55 and Table 56 (p. 143 and 144) report the standardized coefficients. Here, three points stand out.

First, the core of the model was stable. Attitude was the strongest predictor in every model, for both groups and all four outcomes. Subjective norms were also broadly relevant, though for Dutch respondents they reached significance mainly for the intention outcomes. This is the same central pattern found in Chapter 4.3, and it held regardless of language.

Second, one predictor diverged sharply. Academic self-efficacy was negatively and significantly related to non-disclosure in all four Dutch models, but in none of the English models, where its coefficients were close to zero. This is the most notable difference in the analysis, and it was consistent across every outcome.

Third, the Dutch models explained somewhat more variance and contained a few more significant effects. Certainty and swiftness were significant for Dutch past full non-disclosure, and perceived behavioral control and policy awareness for Dutch partial non-disclosure, whereas the English models leaned almost entirely on attitude and subjective norms. The positive effect of swiftness on full non-disclosure intention, against the deterrence prediction, appeared in both groups, which suggests that this anomaly is not specific to one language.

Table 55 – Group-wise regression, past behavior outcomes

Predictor	Past full, EN	Past full, NL	Past partial, EN	Past partial, NL
Certainty	-.010	-.140*	.078	.001
Severity	.060	.025	.004	-.085
Swiftness	-.002	.212**	.119	.048
Attitude	.436***	.409***	.390***	.269**
Subjective norms	.132*	.138	.190**	.004
Perceived behavioral control	.066	.054	.079	.244**
Academic self-efficacy	-.007	-.212***	-.047	-.201**
Policy awareness	-.027	-.056	-.005	-.172*
R ² (adjusted)	.283 (.261)	.439 (.410)	.277 (.254)	.320 (.286)

Note. English n = 264, Dutch n = 167. β = standardized coefficient. * p < .05, ** p < .01, *** p < .001.

Table 56 – Group-wise regression, intention outcomes

Predictor	Full intention, EN	Full intention, NL	Partial intention, EN	Partial intention, NL
Certainty	-.062	-.097	.061	-.001
Severity	-.106	-.066	.018	.017
Swiftness	.183**	.219**	.093	.076
Attitude	.385***	.341***	.378***	.412***
Subjective norms	.174**	.200*	.264***	.182*
Perceived behavioral control	.031	.043	.132*	.219***
Academic self-efficacy	-.005	-.208**	-.022	-.176**
Policy awareness	-.116*	.006	-.095	-.072
R ² (adjusted)	.299 (.277)	.387 (.356)	.372 (.352)	.502 (.477)

Note. English n = 264, Dutch n = 167. β = standardized coefficient. * p < .05, ** p < .01, *** p < .001.

Summary

Taken together, the two analyses tell a consistent story. The two groups differed on the level of a few constructs, most clearly policy awareness, but the model itself behaved in largely the same way for both. Attitude and subjective norms drove non-disclosure regardless of language. The one meaningful exception was academic self-efficacy, which mattered for Dutch respondents but not for English ones. This difference is worth noting, but it should not be over-read. The Dutch group was smaller, and a predictor being significant in one group but not the other does not necessarily mean that the effects differ between groups. In addition, survey language is not a direct indicator of nationality. Overall, the main findings appear consistent across survey-language groups, with academic self-efficacy being the only result that may warrant further investigation.

Appendix 17 – Direct Relationships, Policy Awareness and Deterrence Perceptions

Purpose

The main models treated the deterrence perceptions and policy awareness as separate predictors of non-disclosure. This raised a further question. Is policy awareness also related to the deterrence perceptions themselves? In other words, does a clearer understanding of the rules shape how likely, how severe, or how swift students expect sanctions to be? This appendix reports a supplementary test of that question. It sits outside the hypotheses and is included for completeness.

Analysis

Three simple linear regressions were estimated, one for each deterrence construct. In each model, policy awareness was the sole predictor, and one of the three perceptions, certainty, severity, or swiftness, was the outcome (Field, 2018). All 431 respondents were included. The results are reported in Table 57.

Results

Policy awareness was unrelated to two of the three perceptions. It had no significant relationship with certainty ($\beta = -.014$, $p = .766$) or with severity ($\beta = .069$, $p = .151$). Both models were non-significant and explained almost none of the variance.

The only significant link was with swiftness ($\beta = .141$, $p = .003$). The relationship was positive, so students who were more aware of the rules perceived sanctions as slightly swifter. The effect was small, however. The model explained about 2% of the variance in swiftness ($R^2 = .020$), which is negligible in practical terms.

Taken together, policy awareness showed no meaningful direct relationship with the deterrence perceptions. The one significant association, with swiftness, was small and accounts for little of the variance.

Table 57 – Simple linear regressions of policy awareness on each deterrence perception

Outcome	B	SE	β	t	p	R^2	F(1, 429)
Certainty	-.014	.046	-.014	-0.298	.766	.000	0.089
Severity	.066	.046	.069	1.438	.151	.005	2.068
Swiftness	.153	.052	.141	2.958	.003	.020	8.752

Note. $n = 431$. Each row is a separate simple linear regression with policy awareness as the sole predictor. B = unstandardized coefficient; SE = standard error; β = standardized coefficient; R^2 = explained variance.

Appendix 18 – Qualitative Analysis of the Open-Ended Comment Field

This appendix presents an exploratory qualitative analysis of the responses provided in the survey's optional open-ended comment field. It offers additional insight into the motivations, perceptions, and experiences underlying students' disclosure decisions.

18.1 Purpose and scope

The survey closed with a single optional open-ended field inviting respondents to add any further thoughts, insights, or comments about the research (see appendix 1). This appendix reports a qualitative analysis of the comments left in that field.

The analysis is exploratory and supplementary. It was not designed to test the hypotheses or to produce generalizable results. Its purpose is to show, in students' own words, the reasoning that sits behind the disclosure decisions that are modeled quantitatively in the main study. The regression results describe how strongly certain factors relate to non-disclosure. By contrast, the comments offer insight into the mechanisms and perceptions behind those relationships and reveal themes that were not originally anticipated in the survey design. Open questions of this kind are a recognized way to add depth and unexpected insight to survey research (O'Cathain & Thomas, 2004).

18.2 Data and translation

In total 68 respondents left a usable comment. The comments varied considerably in length, ranging from a few words to short paragraphs. They also varied in language, as the survey was administered in both English and Dutch. Dutch comments were translated into English prior to the analysis. Translated quotations are marked accordingly. Respondents are identified only by their anonymous response number.

18.3 Analytical approach

The comments were examined through a light form of qualitative content analysis (Elo & Kyngäs, 2008). As this qualitative analysis served a supplementary purpose, a full second round of coding and formal inter-coder reliability testing were not conducted. The aim was a transparent and defensible reading rather than an exhaustive one.

The procedure had four steps. First, all comments were read once in full to gain familiarity. Second, an initial set of descriptive categories was developed inductively from the data. Third, related categories were grouped into ten overarching themes, which were further refined into subthemes where additional distinction was useful. Fourth, every comment was coded against the resulting framework.

Two features of the coding are worth noting. Comments were multi-coded: a comment was assigned to every theme it genuinely addressed, not to one theme alone. Many comments raise several ideas at once, and forcing a single label would discard much of their content. As a result, the total number of theme assignments exceeds the number of comments. The counts

should therefore be interpreted as the frequency of themes in the dataset, not as percentages of respondents.

18.4 The thematic framework

The analysis produced ten themes. Six describe how students think about GenAI use and its disclosure. Three describe the consequences and the institutional context. One holds non-substantive feedback regarding the study itself. Table 58 shows these themes, their definitions, and their main subthemes.

Table 58 – Thematic framework for the open-ended comments

Theme	Definition	Subthemes / facets
1. Disclosure: ambiguity, burden and opposition	How hard it is to define, judge, or record GenAI use; the effort of logging it; reasons given for not disclosing; and disclosure in group work.	What counts as disclosure; proportionality of the omission; documentation burden; rejection of the need to disclose; behavioral effect of disclosure rules; group-work accountability.
2. Risk, stigma, credibility, and false accusation	Concerns about what disclosure costs: fear of sanctions, loss of credibility from admitting even responsible use, false accusation, and distrust of detection.	Fear of sanction; credibility loss from disclosing; false accusation; distrust of detection; defensive or strategic behavior.
3. Responsible use and task-dependent boundaries	The view that GenAI is acceptable as support rather than substitute, with acceptability depending on the task, its cognitive demand, whether it is graded, and whether output is verified.	Support not substitute; task-dependent legitimacy (core vs peripheral, study vs graded, text vs code); verification of output; conditional acceptance.
4. Authorship, ownership, and human agency	Justifications that the work is still the student's own, through own voice, own input, prompting as a skill, retained ability, and responsibility for the result.	Own voice; input and ideas are mine; prompting as intellectual skill; retains ability to work unaided; responsibility for the final output.
5. Functions of use (what students do)	Concrete descriptions of what students use GenAI for.	Explanation; feedback and evaluation; writing and language; coding and technical work; summarizing and retrieval; self-quizzing; task initiation; formatting; content generation.
6. Drivers and pressures of use	Circumstances that push students toward GenAI, locating the decision partly in the educational environment.	Workload and deadlines; insufficient feedback or support; peer norms and competitive pressure; efficiency incentives; difficulty starting tasks.
7. Learning and overreliance concerns	Worries that GenAI harms critical thinking, idea generation, foundations, or independence, and that it enables performance without real learning.	Critical thinking; independent idea generation; foundational knowledge; dependency and deskilling; accuracy and hallucinations; performance without understanding.
8. Policy and institutions	Comments on institutional rules, teaching, assessment, staff, adoption, and access.	Rule clarity and consistency; AI-literacy education; assessment redesign; lecturer competence; institutional adoption; approved tools; equity and access (minority case).
9. Normalization and inevitability	GenAI use seen as normal, inevitable, or a workplace skill, including the calculator analogy.	Everyone uses it; workplace relevance; calculator analogy; universities should adapt.
10. Survey and research feedback (non-substantive)	Comments on the instrument, suggestions for future research, and administrative remarks, kept out of the substantive analysis.	Instrument and methodological feedback; future-research suggestions; administrative or contextual remarks.

18.5 Distribution of themes

Table 59 reports how often each theme appears, both as the primary theme of a comment and across all coded mentions. Because comments were multi-coded, the total number of mentions (139) is larger than the number of comments (68), giving an average of about two themes per comment.

Table 59 – Distribution of themes across the 68 comments

Theme	As primary theme	Total mentions
1. Disclosure: ambiguity, burden and opposition	8	15
2. Risk, stigma, credibility, and false accusation	6	6
3. Responsible use and task-dependent boundaries	9	27
4. Authorship, ownership, and human agency	3	9
5. Functions of use	12	19
6. Drivers and pressures of use	2	8
7. Learning and overreliance concerns	5	10
8. Policy and institutions	10	19
9. Normalization and inevitability	3	10
10. Survey and research feedback	10	16
Number of comments		68
Total theme mentions	68	139
Average themes per comment		2.04

Three themes stand out by volume. Responsible use and task-dependent boundaries is the most frequently coded idea. Comments within this theme commonly framed AI use as conditional rather than universally acceptable. Its legitimacy depended on the type of task and whether the AI-generated output was independently checked. Functions of use follows, capturing the specific ways students reported using AI, and Policy and institutions follows, focused on the role of universities and educators.

Two less frequent themes were, ‘Drivers and pressures of use’ and ‘Risk, stigma, credibility, and false accusation’. These are important nonetheless, and directly connect to the reason why a student might conceal use.

18.6 Thematic findings

The following sections present the themes in five clusters. Each theme is illustrated with excerpts drawn directly from the comments.

Making sense of disclosure (Themes 1 and 2)

A recurring difficulty is that students do not know what disclosure actually requires. Several respondents found it difficult to determine where ordinary assistance ends and reportable AI use begins. They also drew a clear distinction between minor uses of AI and more substantial uses, arguing that a single declaration would treat both in the same way. One respondent set out the problem directly:

“Do you mean that I leave out small things, or big things? For example, claiming I used no AI at all when it wrote all the text, versus saying I did not use AI when I once asked it for a link to an article I could not find myself? Those are quite big differences.” (Respondent 521, translated from Dutch)

Another respondent questioned how substantial an undisclosed use of AI must be before non-disclosure could reasonably be considered dishonest, arguing that the seriousness of the omission was often unclear (Respondent 74). A related concern is the effort of documenting. As one respondent put it, if every step had to be recorded, the declaration of AI use would often run to more pages than the assignment itself (Respondent 65, translated from Dutch). Some respondents went a step further and argued that disclosure is not always required. Where students see prompting as their own thinking, they may not regard use as reportable at all. One respondent explained that *“GenAI is based on my thought process of asking questions”*, leading them to view AI use as part of “my whole thought process” rather than as something requiring disclosure (Respondent 360). Finally, group work introduces a disclosure challenge that individual rules do not fully anticipate. One respondent noted that in shared documents they have *“no control or idea of whether anything the others put in there is AI generated or not,”* and argued that AI use has *“made group work more difficult with people I don't already know and trust”* (Respondent 72).

These uncertainties connect to a second cluster of comments about risk, stigma, and credibility. Here the concern is not what to disclose, but what disclosure costs. Several respondents felt that admitting even careful, minor use lowers the standing of their work:

“It is more difficult and seems worse to disclose actually responsible uses of AI. I find it does still reduce credibility even if it is a small contribution which is thoroughly checked afterwards.” (Respondent 89)

Others feared the opposite error, being wrongly accused of use they had not made. One respondent described a false accusation and having to prove their case, concluding that the current systems for assessing AI use are not trustworthy (Respondent 157).

Together these comments highlight a disclosure dilemma. On the one hand, disclosure can protect students from accusations of cheating. On the other hand, it may lead others to view their work less favorably or question its credibility. Non-disclosure avoids this risk but increases the possibility of being accused of misconduct. This tension is clearest in Respondent 50, who both feared being *“automatically discredited”* for disclosing and argued that responsible use should be de-stigmatized.

Responsibility and authorship (Themes 3 and 4)

The most common single idea across the comments is that GenAI is acceptable when it supports rather than replaces the student’s own work. The acceptability depends on the task. Many respondents judged AI use based on how it was used rather than whether it was used at all: *“the question is less about the amount AI is used and more about how we use it”* (Respondent 364). For many respondents, whether AI use was acceptable depended on the

task. One respondent was fairly strongly against using GenAI to generate text for reports, yet felt no discomfort using it to write statistical code (Respondent 303, translated from Dutch). Other respondents distinguished between using AI for core parts of an assignment and using it for supporting tasks, as well as between studying and graded work. A common view was that AI should support students rather than do the work for them. As one respondent put it, AI should play “*a supporting role, not a leading one*” (Respondent 238, translated from Dutch). Another described it as “*a sparring partner, not a slave*” (Respondent 166, translated from Dutch).

Closely related is a set of comments about authorship and ownership, in which students justify use by stressing that the work remains theirs. They describe keeping their own voice, verifying outputs, and still being able to work independently of GenAI. One respondent captured the ownership claim, as: “*all the input comes from me; I only let AI phrase it more clearly*” (Respondent 471, translated from Dutch). Another rejected any text that did not sound like them, because AI does not use their own voice and thoughts (Respondent 570). This theme matters for disclosure, because a student who sees GenAI as a tool through which they express their own thinking may feel there is little to declare.

What students use GenAI for, and what pushes them (Themes 5 and 6)

When respondents described their actual practice, the uses were varied and mostly specific: explaining difficult concepts, checking and generating code, summarizing long material, drafting and editing text, generating practice exercises, self-quizzing, and handling formatting or technical tasks. One respondent, a student of AI, valued being able to “*focus more on the big picture and architecture rather than small programming errors*” (Respondent 52, translated from Dutch). Another respondent restricted their GenAI use to peripheral tasks while writing the substance alone (Respondent 585, translated from Dutch). The value of these comments lies less in their individual details than in what they show collectively. Together, they demonstrate that GenAI use is not a single behavior but a range of different practices, each with different implications for whether students consider disclosure necessary.

A smaller but pointed set of comments describes the pressures that push students toward GenAI in the first place. Unlike comments that simply describe AI use, these responses highlight how students’ surroundings and circumstances shape their disclosure decisions. Respondents referred to workload and deadlines, little or late feedback from lecturers, and competition with peers. One began using AI tools only after noticing that heavily dependent classmates “*were achieving better marks than myself*” (Respondent 467). Another described a growing sense of compulsion: “*there are disadvantages or delays if I opt out of GenAI compared to peers that use GenAI*” (Respondent 462). Respondent 50 links use explicitly to shrinking staff feedback, turning to AI to assess and double-check their own work.

Read together, these comments suggest that non-disclosure may sometimes be less a moral choice than a response to structural pressure.

Consequences and context (Themes 7, 8, and 9)

A further group of comments show concern about the effects of GenAI on learning. Respondents worried about dependence, weakened critical thinking, and the loss of foundational skills. One reflected that “*people before GenAI could accomplish so much, and now we are dependent on AI for everything*” (Respondent 464). Another feared that overuse “*will degrade critical thinking ability*” (Respondent 162). Several distinguished students who use AI to learn from those who use it to avoid learning: “*get away with it, and do not learn anything*” (Respondent 73).

Many comments turn to the institution. Respondents repeatedly asked for clearer and more consistent rules, noting that expectations differ across courses, lecturers, and universities. One asked simply that every course should state its rules clearly, which would also reduce conflict in group work (Respondent 55, translated from Dutch). Some criticized staff for lagging behind, arguing that institutions risk becoming irrelevant if they resist rather than adapt (Respondent 70). Others proposed changing assessment rather than policing tools, for example by moving toward in-class or application-based tasks (Respondents 73 and 371). One respondent raised another important nuance, the risk that unequal access to AI could widen the divide between well-resourced and under-resourced institutions (Respondent 337).

A final group treats GenAI as already normal or inevitable. For some, use is unremarkable because “*everyone uses it*” (Respondent 159). Others framed it as a future skill, or drew an analogy to similar tools from the past: with human oversight, AI would come to resemble a calculator that students must learn to work alongside (Respondent 537). This sense of normalization is a relevant dimension of disclosure, since behavior that feels standard may feel less like something that needs to be declared.

Feedback on the study (Theme 10)

A number of comments addressed the survey itself rather than GenAI use. Some offered methodological observations, for example that the survey phrasing risked sounding judgmental (Respondent 171) or that the meaning of dishonest disclosure was ambiguous (Respondent 74). Others suggested directions for future work, such as examining the drivers of use or applying a behavioral model (Respondent 245). These comments were kept separate from the substantive themes.

18.7 Contribution to the thesis

Although exploratory, this analysis strengthens the study in three ways. First, it triangulates the quantitative model. Several themes map onto constructs in the conceptual model and give them a nuanced view. The comments on unclear and inconsistent rules speak directly to policy awareness. The comments on fear of sanctions, false accusation, and credibility loss illustrate the perceptions that deterrence and subjective norms are meant to capture. The comments on normalization show subjective norms operating in students’ own words. The pressures theme connects to academic self-efficacy and to workload, indicating conditions under which students lean on GenAI and may prefer to conceal that reliance.

Second, the analysis supports the study's central conceptual move, which is to treat non-disclosure as more than a binary act. Respondents repeatedly distinguished minor from substantial uses and questioned what a declaration should contain. Comments from Respondents 521 and 74 show why students consider full and partial non-disclosure to be meaningfully different. They also demonstrate how a simple declaration can misrepresent what actually happened.

Third, and most broadly, the comments provide insights that regression analysis alone cannot show. Taken together, the comments suggest that students consider several factors when evaluating GenAI use. They consider the type and importance of the task and how much of the work is their own. They also consider whether they still understand and control the work, what their institution expects, and the possible consequences of disclosure. Their decisions are further shaped by workload, competition, the availability of feedback, and the clarity of the rules. Non-disclosure, on this reading, is not simply a matter of individual honesty. It is a judgment made under specific task, social, and institutional conditions. This nuance complements the modeled relationships and helps to explain them.

18.8 Limitations

Several limitations should be kept in mind. The comments come from a select subset of respondents who chose to write something, so they may over-represent stronger opinions. The analysis was conducted by a single coder without a formal reliability check, which suits its supplementary purpose but limits the strength of the claims that can be drawn from it. Finally, the theme counts reflect emphasis within this set of comments and should not be read in the context of the wider student population. Within these limits, the analysis is offered as a transparent and illustrative complement to the main findings, not as an independent test of them.

Appendix 19 – AI Disclosure statement

In line with the requirements for transparency about the use of generative artificial intelligence (GenAI) in academic work, I set out below how, where, and for what purposes I used GenAI during the preparation of this thesis. The main tool I used was Anthropic's Claude, and within it I relied mainly on the Opus 4.8 model throughout the research process. I used GenAI to support my work rather than to produce it. In every case, I reviewed, verified, and edited the outputs myself, and I take full responsibility for the final content.

Literature search and identification of sources.

I used GenAI to make my literature search more efficient. It helped me generate area-specific terminology, synonyms, and related keywords, and it pointed me towards relevant articles and papers. This was particularly useful for finding sources that I would likely have missed otherwise, since it broadened my search terms and made the search more targeted. I did not use GenAI as a substitute for reading. All sources were carefully consulted in their original form, which were in turn saved locally for later consultation. I decided myself which sources were relevant and how I used them.

Language, grammar, and clarity

I used GenAI to check the grammar, spelling, and flow of text that I had written. It helped me spot errors, including some that I probably would not have found on my own, which I then corrected myself. I also used Grammarly as a writing-support tool to fix errors and to rephrase sentences for greater clarity and a more academic tone. I reviewed every suggested change and applied it at my own discretion.

Explanation of concepts and conceptual sparring

I used GenAI to obtain simplified explanations of complex concepts, which helped me build a clearer understanding of the underlying material more quickly. I also used it as a sparring partner when I reached a point of decision, since it helped me reason through my options and infer a sensible direction. This improved my efficiency, but it did not replace my own judgement, which guided every decision I made.

Survey development and translation

I used GenAI to support the development of my measurement instrument in Qualtrics. It helped me brainstorm, edit, and rephrase survey items to reduce bias, ambiguity, and double-barreled phrasing, and my items went through several iterations before I settled on the final versions. I also used it to translate survey items and open-response fields so that they stayed clear and unbiased across languages. I checked all translations myself.

Data analysis support

I used GenAI to clarify certain nuances of data analysis in SPSS. Because I had mainly worked in R in earlier projects, it helped me understand how to set up specific analyses in SPSS. It did not carry out the analyses, and I performed all analytical work and interpretation myself.

Declaration of Authorship

I declare that this thesis is my own work and that I am responsible for its content. Any ideas, findings, text, or other material derived from the work of others have been appropriately acknowledged, cited, and referenced. The use of generative AI and other digital writing tools has been disclosed above and does not replace my own academic judgement or responsibility.

Name: Finn Timmermans

Location: Delft, the Netherlands

Date: August 17, 2026

Signature:

A handwritten signature in black ink, appearing to be 'F. Timmermans', written in a cursive style.