Beyond the Incentive Threshold: Academic Support Use and Persistence After a First-Year Participation Incentive

Heriberto Espino-Montelongo
Universidad de las Américas Puebla

Abstract

Universities increasingly use structured first-year participation requirements and academic-support services to encourage student engagement, yet much less is known about whether engagement with support services persists after a particular incentive-linked context has passed. We study formal mathematics-support use at the Centro de Aprendizaje de Matemáticas (CMAT) of Universidad de las Américas Puebla. In the institutional setting examined here, participation in CMAT during first-year Matemáticas Universitarias (MU) could contribute to a broader first-year participation requirement, while later Calculus I use is treated as occurring outside that same PPA1-linked incentive context. We follow 3,241 students who passed their first observed MU attempt, progressed to Calculus I in the next regular academic term, had CMAT coverage in both periods, and had valid classroom-relative performance measures. The probability of any later CMAT use in Calculus increased from 17.25% among students with no MU visits to 43.43% for 1–2 visits, 54.37% for exactly three visits, and 66.84% for four or more visits. The association was strong overall ($\chi^2(3) = 392.47$, $p < .001$, Cramér's $V = .348$). Students with four or more MU visits were 12.47 percentage points more likely to use CMAT later than students with exactly three visits (95% CI [0.92, 23.90] percentage points). Adjusted logistic models controlling for prior classroom-relative performance, degree programme, and term preserved the gradient in later use. These findings do not identify a causal effect of the participation incentive or of CMAT attendance. Instead, they show substantial behavioural persistence in formal academic help-seeking across a change in incentive context and illustrate why evaluations of voluntary support services should distinguish initial exposure, threshold-linked participation, and later self-selected continuation.

Keywords: academic help-seeking; student engagement; mathematics support; higher education; incentives; persistence; first-year transition

1 Introduction

Formal academic-support services are a common part of higher-education infrastructure. Mathematics and statistics support centres, tutoring programmes, supplemental instruction, and related services are designed to provide students with opportunities to seek help outside the standard classroom environment (Lawson et al. 2020; Mullen et al. 2024). Yet evaluating these services is difficult because use is typically voluntary. Students decide whether to attend, how often to return, and when to stop. As a result, observed attendance combines the institutional availability of support with student-level need, preferences, prior experiences, scheduling constraints, and other forms of selection.

This distinction matters because engagement is not a single construct. Higher-education research has long separated institutional conditions and student characteristics from behavioural manifestations of engagement and from later consequences (Kahu 2013; Kahu and Nelson 2018). Formal help-seeking is itself associated with academic achievement, but the direction and meaning of that relationship depend on the type of help-seeking and on the process through which students enter support (Fong et al. 2023). Within mathematics support specifically, prior research documents both potential benefits and persistent non-engagement, while also emphasizing methodological difficulties created by self-selection (Jacob and Ní Fhloinn 2019; Mac an Bhaird et al. 2013; Mullen et al. 2024).

Institutional incentives add another layer. Interventions can change participation directly, but the behaviour observed while an incentive is active need not reveal what students will do once that incentive is removed or becomes irrelevant. Economic and behavioural research therefore distinguishes the immediate response to incentives from behaviour after the incentivized context ends (Gneezy et al. 2011). In education, randomized and quasi-experimental studies show that incentives, encouragements, and nudges can alter service use or intermediate behaviours, although effects on academic outcomes are often heterogeneous or limited (Angrist et al. 2009;

Blondeel et al. 2024; Damgaard and Nielsen 2018; Paloyo et al. 2016; Pugatch and Wilson 2018). Evidence on behaviour after incentive removal is less common, but it is particularly informative when the substantive question is whether participation is confined to the incentivized context or persists beyond it (Agnew et al. 2021).

This paper uses a longitudinal administrative setting in which these issues can be examined directly, although not causally. At Universidad de las Américas Puebla (UDLAP), the Centro de Aprendizaje de Matemáticas (CMAT) provides mathematics and statistics support to undergraduate students. During the first-year context considered here, CMAT participation could contribute to the institutional Programa de Primer Año (PPA1), a broader participation requirement. The operational threshold relevant to this study is three recorded CMAT visits during the MU period. Later Calculus I use is treated as occurring after that same PPA1-linked CMAT incentive is generally no longer operative.

The design asks a simple longitudinal question: *does formal support use observed in an incentive-linked first-year context persist into later use when that same incentive context is no longer assumed to apply?* We do not observe why an individual student attended CMAT, whether a particular visit was motivated by PPA, or the exact date at which every student completed PPA1. We therefore do not estimate a causal effect of PPA. Instead, we compare observable patterns of formal support use over time and keep three explanations visible throughout: threshold-limited participation, persistence or familiarization with formal support, and stable student selection into help-seeking.

The paper makes three contributions. First, it shifts the outcome of interest from contemporaneous grades to *later support use*, treating persistence in formal help-seeking as substantively important in its own right. Second, it exploits the institutional threshold to distinguish students who stop exactly at three recorded visits from those who continue beyond the threshold, without treating that threshold as exogenous. Third, it connects mathematics-support research to broader higher-education debates on engagement and institutional participation by separating incentive-linked participation from later self-selected continuation. This broader framing is important for a setting-specific study to contribute beyond the local service itself, particularly for journals that expect empirical work to speak to general higher-education processes (Blondeel et al. 2024; Bowden et al. 2021; Felby et al. 2026; Herpen et al. 2020; Kahu 2013).

2 Conceptual framing

2.1 Institutional conditions, engagement, and formal help-seeking

A useful starting point is to distinguish institutional antecedents, student characteristics, behavioural engagement, and outcomes. In Kahu's framework, engagement emerges from the interaction between institutional and student factors rather than being reducible to a single student trait (Kahu 2013). The educational-interface formulation similarly treats engagement as situated within an interaction between students and their learning environment (Kahu and Nelson 2018). This distinction is important for administrative attendance data: a CMAT visit is an observed behaviour, not a direct measure of motivation, belonging, self-regulation, or need.

Formal academic help-seeking can be adaptive, but its empirical relationship with achievement varies across forms of help-seeking and across students (Fong et al. 2023). Mathematics-support research likewise shows that centre use is neither randomly distributed nor universal (Lawson et al. 2020; Mac an Bhaird et al. 2013; Mullen et al. 2024). Students who use support may differ from non-users before the first recorded visit. Consequently, persistence in CMAT use can reflect both experiences acquired through prior contact with the service and stable propensities that make some students more likely to seek help repeatedly.

2.2 Incentive-linked participation and post-incentive behaviour

The literature on incentives provides a second distinction: behaviour while an incentive is active and behaviour after its removal are different outcomes. Incentives can alter the immediate cost-benefit structure of participation, while later behaviour may reflect learning about the service, reduced search costs, familiarity, crowding effects, or simply persistent differences between individuals (Gneezy et al. 2011). Educational interventions illustrate the same separation between treatment assignment, actual uptake, and later outcomes. For example, encouragement designs can increase use of academic support without making attendance itself exogenous (Paloyo et al. 2016; Pugatch and Wilson 2018), and randomized nudges can show weak average treatment effects even when voluntary intensity of engagement is associated with performance (Blondeel et al. 2024).

For the present study, three observable patterns are therefore conceptually useful. A *threshold-limited pattern* occurs when a student reaches exactly the operational threshold in MU and does not use CMAT later. A *persistent-use pattern* occurs when a student uses CMAT during MU and later returns during Calculus. A *late-adoption pattern* occurs when a student has no recorded MU use but begins using CMAT in Calculus. These are behavioural descriptions, not diagnoses of motivation.

The comparison between exactly three and four or more MU visits is especially informative descriptively. If students who stop exactly at three were much less likely to return later than students who continue beyond three, this would be compatible with heterogeneity between threshold-limited and more persistent use. However, because students choose their own visit counts, the comparison is not a regression-discontinuity design and cannot identify a treatment effect of crossing the threshold.

3 Institutional setting and data

CMAT is an institutional mathematics-learning support centre at UDLAP. It serves undergraduate students across mathematics and statistics subjects. Advising encounters are recorded through a Google Forms process, with one row representing one recorded visit. The form includes a student identifier and contemporaneous information about the advisory encounter. Multiple visits on the same day can be legitimate because faculty advisers operate in time blocks and students may continue working with another adviser; same-day concentration is therefore not treated as evidence of manipulation.

The academic-record data contain course attempts, final academic outcomes, professor, academic period, and official degree programme. A key administrative complication is that academic records can contain revalidation or programme-change rows that do not represent a new course attempt. The longitudinal pipeline therefore distinguishes likely real attempts from later administrative replications of already passed courses. For classroom-relative outcomes, a classroom is defined as professor × same course × academic period.

PPA1 is a first-year institutional participation requirement that can be completed through different university activities, including CMAT-related activities. PPA1 is not part of the MU course grade. The study uses an operational threshold of three recorded CMAT visits during the MU period; it does not equate one recorded visit with one PPA point. The design assumes that MU typically occurs in the PPA1-linked first-semester context and treats later Calculus I as a context in which the same PPA1-linked CMAT incentive is generally no longer operative. Individual motivation and exact PPA completion timing are not observed.

4 Study design

4.1 Longitudinal cohort

The primary analysis uses a strict MU-to-Calculus progression cohort. Students must have: (1) a first real observed attempt at Matemáticas Universitarias; (2) a numeric passing grade of at least 7.5 in that attempt; (3) a later first real Calculus I attempt; (4) a numeric Calculus I grade; (5) CMAT coverage in both periods; (6) valid classroom-relative performance measures in both courses; and (7) Calculus I in the next regular spring/fall academic period after MU.

Figure 1 summarizes the cohort construction. The resulting primary cohort contains 3,241 students.

The 3,389 students who satisfy all conditions except the next-regular-term restriction form a delayed-progression sensitivity cohort in the historical longitudinal analysis. A broader 4,211 future-Calculus progressor cohort is also documented in the methodology record, but it is not yet fully reconciled with the stricter longitudinal pipeline and is therefore not substituted silently for the primary estimand in this draft.

4.2 Exposure and outcomes

The primary predictor is the number of recorded CMAT visits during the student's MU period, grouped as shown in Equation (1).

$$0, \quad 1\text{–}2, \quad 3, \quad 4+. \tag{1}$$

The exactly-three group corresponds to the operational institutional threshold and the four-or-more group captures recorded use beyond that threshold.

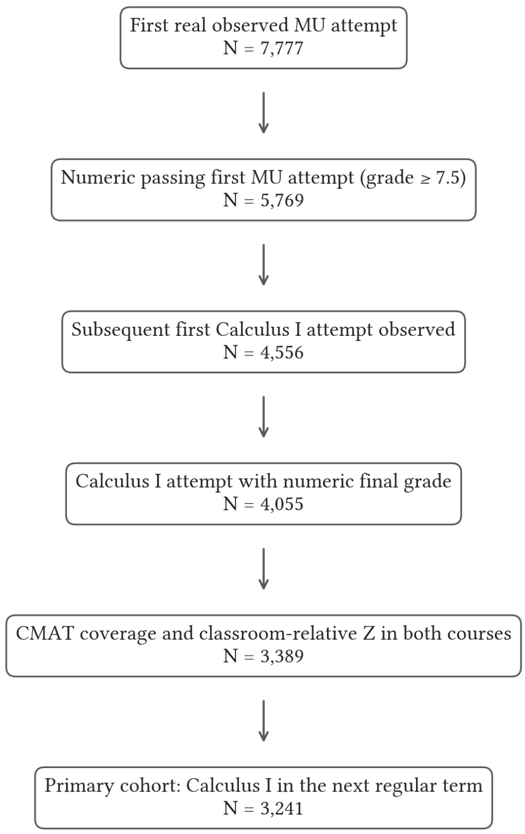

Figure 1: Construction of the primary longitudinal cohort. Counts are reproduced from the retained aggregate cohort-flow table; the final analytic sample contains students who passed their first real observed MU attempt, subsequently had a numeric Calculus I grade, had CMAT coverage and valid classroom-relative performance measures in both courses, and progressed to Calculus I in the next regular term.

The primary outcome is any recorded CMAT use during the subsequent eligible Calculus I period. This outcome is binary and measures later formal support use, not academic achievement.

The secondary academic outcome is classroom-relative Calculus performance, defined in Equation (2). For student i in classroom c,

$$Z_{ic} = \frac{Y_{ic} - \bar{Y}_c}{s_c}, \tag{2}$$

where c is defined by professor, course, and academic period. Classroom reference distributions use numeric grades from eligible real attempts rather than only the final longitudinal subsample.

4.3 Statistical analysis

We first estimate the probability of any later CMAT use for each MU visit group and report Wilson 95% confidence intervals. A Pearson chi-square test summarizes the global association between the four MU groups and later use, with Cramér's V as an effect-size measure.

The pre-specified threshold-neighbour comparison contrasts students with exactly three MU visits against students with four or more. We report the risk difference with a Newcombe confidence interval, the risk ratio, odds ratio, Pearson test, and two-sided Fisher exact test. This is an observational risk comparison.

We then estimate logistic regressions for later Calculus use. The adjusted model includes MU visit group, prior classroom-relative MU performance (Z_{MU}), official degree programme at the MU attempt, and MU academic period. Standard errors are clustered by MU classroom.

As an exploratory description of the shape of the relationship, a piecewise visit-count model allows one slope through three visits and a second slope after three. The model is descriptive and is not interpreted as regression discontinuity.

For the secondary academic outcome, we estimate linear models of Z_{Calc} including Z_{MU} and MU visit group, followed by a specification that additionally adjusts for official degree programme in Calculus and Calculus period. Standard errors are clustered by Calculus classroom.

5 Results

5.1 Persistence of formal support use

Later CMAT use differs sharply by prior MU engagement (Table 1). Among students with no recorded MU visits, 426 of 2,470 (17.25%) used CMAT during Calculus. The corresponding proportions were 43.43% for students with 1–2 MU visits, 54.37% for exactly three visits, and 66.84% for four or more.

Table 1: Later CMAT use in Calculus I by prior MU visit group

MU visits	N	Later users	Probability	95% CI
0	2,470	426	17.25%	[15.81%, 18.79%]
1–2	472	205	43.43%	[39.03%, 47.94%]
3	103	56	54.37%	[44.77%, 63.66%]
4+	196	131	66.84%	[59.98%, 73.05%]

The global association is strong: $\chi^2(3) = 392.47$, $p = 9.47 \times 10^{-85}$, with Cramér's $V = .348$. The minimum expected cell count is approximately 26, so sparse expected cells do not drive the Pearson result.

5.2 Exactly three versus four or more visits

The comparison most directly tied to the institutional threshold contrasts exactly three visits with four or more. Later CMAT use was 54.37% in the exactly-three group and 66.84% in the four-or-more group. The risk difference was 12.47 percentage points (95% CI [0.92, 23.90]). The risk ratio was 1.229 (95% CI [1.004, 1.505]) and the odds ratio was 1.691 (95% CI [1.038, 2.757]). The Pearson test gave $p = .0343$ and Fisher's exact test gave $p = .0440$.

The comparison is consistent with behavioural heterogeneity around the operational threshold: students who continued beyond three MU visits were more likely to use CMAT again later than students who stopped

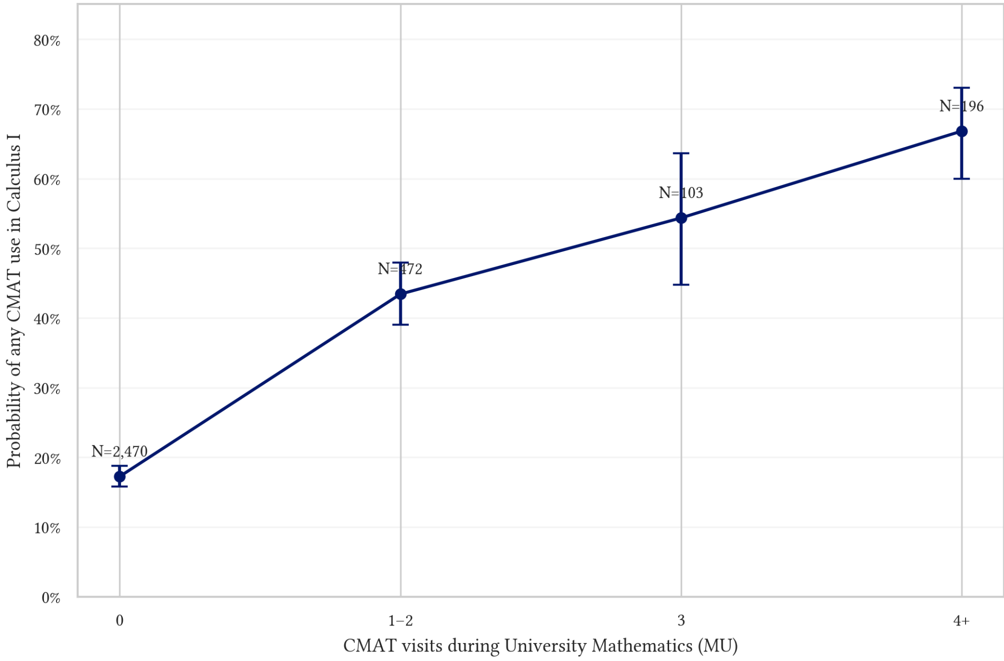

Figure 2: Probability of any CMAT use during Calculus I by prior University Mathematics (MU) visit group in the primary longitudinal cohort ($N = 3{,}241$). Points show observed group proportions and error bars show Wilson 95% confidence intervals.

exactly at three. It does not establish that crossing three visits caused persistence, because the number of visits is student-controlled.

5.3 Visit-count shape around the threshold

A descriptive piecewise model shows a much steeper association between visit count and later use through three visits than after three. Up to the threshold, the estimated odds ratio per additional visit is 1.994 (95% CI [1.807, 2.200], $p < .001$). Beyond three visits, the corresponding odds ratio is 1.063 (95% CI [0.982, 1.150], $p = .131$). Figure 3 summarizes this contrast. The flattening is descriptive; it does not identify a discontinuity caused by the institutional threshold.

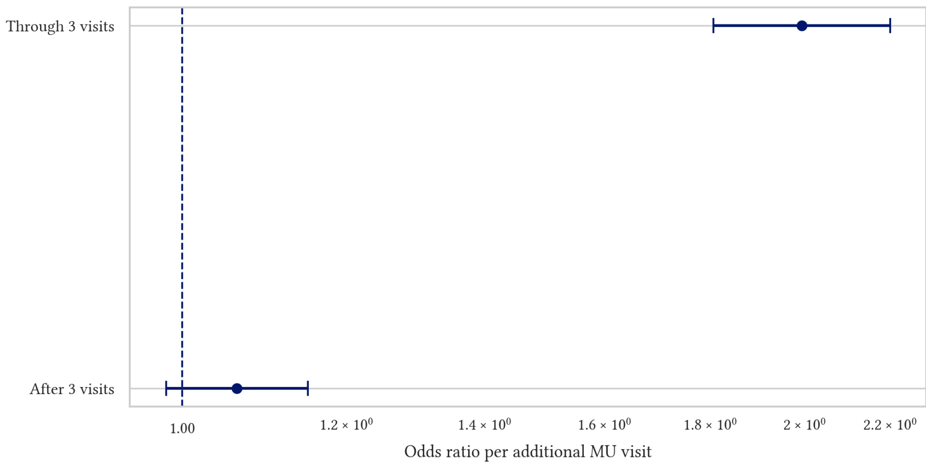

Figure 3: Estimated per-visit odds ratios from the descriptive piecewise logistic model of later CMAT use, with cluster-robust 95% confidence intervals. The first slope applies through three MU visits and the second applies after three visits. The dashed reference line denotes an odds ratio of 1. Visit count is student-controlled, so the figure is not a regression-discontinuity estimate.

5.4 Adjusted persistence models

The persistence gradient remains after adjustment for prior classroom-relative MU performance, official degree programme, and MU period (Figure 4). Relative to students with no MU visits, adjusted odds of later CMAT use were 3.53 times higher for 1–2 visits, 4.84 times higher for exactly three visits, and 9.82 times higher for four or more visits. All three group coefficients were statistically distinguishable from zero in the adjusted model.

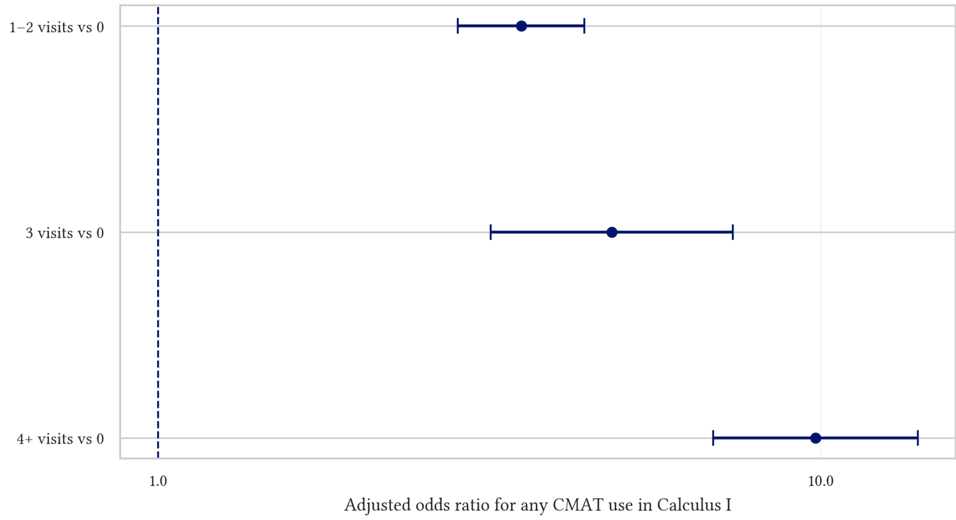

Figure 4: Adjusted odds ratios for any CMAT use during Calculus I by prior MU visit group. The reference group is zero MU visits. The model adjusts for prior classroom-relative MU performance, official degree programme at the MU attempt, and MU academic period; error bars show cluster-robust 95% confidence intervals with clustering by MU classroom.

The adjusted model also shows that higher prior relative performance predicts later CMAT use ($OR = 1.59$ per one-unit increase in Z_{MU}, 95% CI [1.27, 1.98]). This is another indication that later use cannot be reduced to a simple remediation story in which only low-performing students seek support.

5.5 Secondary association with later academic performance

Prior CMAT-use groups also remain positively associated with classroom-relative Calculus performance after controlling for Z_{MU}, official Calculus degree programme, and Calculus period. Relative to students with no MU visits, coefficients were $0.122Z$ for 1–2 visits (95% CI [0.051, 0.194]), $0.153Z$ for exactly three visits (95% CI [0.037, 0.269]), and $0.172Z$ for four or more visits (95% CI [0.070, 0.274]). Prior Z_{MU} remained the dominant academic predictor ($\hat{\beta} = 0.578$, 95% CI [0.495, 0.661]).

These estimates are secondary and observational. They do not imply that prior CMAT attendance caused later academic gains. Their main role in this paper is to show that the observed persistence gradient is not confined to a population with uniformly weak prior academic standing and that later performance remains strongly related to earlier classroom-relative performance.

6 Discussion

The central result is straightforward: formal mathematics-support use shows substantial persistence across the MU-to-Calculus transition. Only 17% of students with no recorded MU use subsequently used CMAT in Calculus, compared with 43% of students with 1–2 visits, 54% of those with exactly three visits, and 67% of those with four or more. This gradient remains large after adjustment for prior relative performance, degree programme, and academic period.

The findings are relevant to higher-education research for three reasons. First, they show the value of separating an institutional participation condition from the later behaviour that follows it. An observed visit during an incentive-linked period is not the same outcome as later voluntary continuation. Studies of incentives and nudges often distinguish offer, exposure, uptake, and downstream behaviour (Angrist et al. 2009; Blondeel et al. 2024; Paloyo et al. 2016; Pugatch and Wilson 2018); the same distinction is useful when studying academic-support services.

Second, the results suggest that support-centre use is not well represented as a homogeneous dose. The largest conceptual distinction is between no previous use and any previous use, but behaviour also differs within positive-use groups. Students who continued beyond the operational threshold were more likely to return later than students who stopped exactly at the threshold. At the same time, the piecewise count model flattens after three visits, which cautions against interpreting each additional visit as an equivalent increment of engagement.

Third, the results connect formal mathematics support to broader theories of student engagement. Institutional conditions can shape opportunities and incentives for behavioural engagement, but continued engagement reflects an interaction between the environment and student characteristics (Kahu 2013; Kahu and Nelson 2018). The observed persistence pattern is therefore compatible with multiple mechanisms. Early contact with CMAT may reduce future search costs, increase familiarity with formal help-seeking, or make the service more salient. Alternatively, students who are already more willing or able to seek formal support may use CMAT repeatedly across courses. The present data cannot causally distinguish these mechanisms.

The exact-three versus four-or-more comparison is particularly useful for interpretation. Students who stop at three and do not return later display the observable pattern most compatible with threshold-limited use. Students who exceed three and later return display the clearest persistence pattern. But neither group reveals internal motivation. The same visit sequence could arise from academic need, schedule constraints, instructor referrals, perceived usefulness, prior preferences, or other unobserved factors. We therefore avoid labeling students as "incentive-motivated" or "intrinsically motivated."

The secondary performance results should be read similarly. Positive adjusted associations with later Z_{Calc} are consistent with the possibility that persistent support users remain academically engaged, but they are also compatible with selection. Recent mathematics-support research has increasingly emphasized stronger causal and quasi-causal designs where possible (Büchele and Schürmann 2024), while reviews of the field continue to note wide variation in evaluation quality (Mullen et al. 2024). The present paper contributes a different object: longitudinal persistence in formal help-seeking, not a causal estimate of tutoring effectiveness.

6.1 Implications for institutional support design

For institutions, the practical implication is not that participation incentives should automatically be expanded. Rather, the results suggest that evaluations of first-year participation structures should measure what happens after the incentivized window. If a support service is intended partly to make help-seeking more accessible or familiar, later voluntary use is an informative outcome. Conversely, if participation disappears entirely when an incentive ends, high use during the incentivized period may overstate durable engagement.

This distinction also suggests a useful design principle for future institutional evaluation. Administrative systems should separately record incentive eligibility, incentive completion, service exposure, service use, and later service use. Doing so would allow researchers to distinguish whether an intervention primarily increases initial contact, repeated use during the active incentive, or continued use after the incentive ends.

7 Limitations

The study has several important limitations. First, CMAT use is observational and student-selected. Adjusted regression reduces measured imbalance but does not eliminate unobserved confounding. The results should therefore be interpreted as associations in behavioural trajectories, not causal treatment effects.

Second, the institutional assumption about PPA1 timing is not observed at the individual level. We know the operational first-year context and the three-visit threshold used in the study, but we do not directly observe which visit, if any, a student undertook specifically to satisfy a PPA requirement. We also do not observe internal motivation.

Third, the primary cohort conditions on passing MU and progressing to Calculus with valid linked records. This defines a substantively meaningful risk set for later Calculus use but also selects on academic progression. Results therefore generalize most directly to students who follow that pathway, not to all first-year students.

Fourth, visit records measure recorded encounters rather than duration, intensity, content quality, or learning processes. Multiple same-day visits can represent legitimate continued support across adviser blocks, and the data do not support treating them as manipulation.

Fifth, the study is based on one institution and one support-centre context. The contribution therefore depends on connecting the local institutional design to broader higher-education processes rather than assuming direct generalizability. Cross-institution replication would be valuable.

Finally, the strict $N = 3{,}241$ longitudinal analysis was produced in a later development stage than the currently reproducible methodology-report pipeline. The underlying analysis and retained outputs are documented, but the longitudinal logic still needs to be fully reintegrated into a single canonical executable pipeline before submission.

8 Conclusion

This study shows that formal mathematics-support use can persist well beyond an incentive-linked first-year context. Among students progressing from MU to Calculus, later CMAT use increased sharply with prior use, from 17% among non-users to 67% among students with four or more MU visits. Students who continued beyond the operational three-visit threshold were also more likely to return later than students who stopped exactly at three. These patterns survive adjustment for prior relative performance, degree programme, and term.

The evidence does not show that PPA caused persistence or that CMAT attendance caused academic improvement. Its contribution is instead to separate three objects that are often conflated in evaluations of voluntary support: institutional conditions that encourage initial participation, student-selected use while those conditions are active, and later continuation when the same incentive context no longer applies. Treating these as distinct outcomes provides a more informative way to study academic support, engagement, and persistence in higher education.

References

Agnew, S., J. Kerr, and R. Watt (2021). "The effect on student behaviour and achievement of removing incentives to complete online formative assessments". In: *Australasian Journal of Educational Technology* 37.4, pp. 173–185. DOI: 10.14742/ajet.6203.

Angrist, J., D. Lang, and P. Oreopoulos (2009). "Incentives and Services for College Achievement: Evidence from a Randomized Trial". In: *American Economic Journal: Applied Economics* 1.1, pp. 136–163. DOI: 10.1257/app.1.1.136.

Blondeel, E., P. Everaert, and E. Opdecam (2024). "A little push in the back: nudging students to improve procrastination, class attendance and preparation". In: *Studies in Higher Education* 49.11, pp. 2016–2035. DOI: 10.1080/03075079.2023.2288170.

Bowden, J. L.-H., L. Tickle, and K. Naumann (2021). "The four pillars of tertiary student engagement and success: a holistic measurement approach". In: *Studies in Higher Education* 46.6, pp. 1207–1224. DOI: 10.1080/03075079.2019.1672647.

Büchele, S. and M. Schürmann (2024). "Causal evidence of the effect of a course specific mathematical learning and support center on economics students' performance and affective variables". In: *Studies in Higher Education* 49.9, pp. 1569–1587. DOI: 10.1080/03075079.2023.2271029.

Damgaard, M. T. and H. S. Nielsen (2018). "Nudging in education". In: *Economics of Education Review* 64, pp. 313–342. DOI: 10.1016/j.econedurev.2018.03.008.

Felby, L. C., R. T. Jensen, and M. Fibiger (2026). "The stars, compass, and anchor: how relationships shape first-year students' perception of their access to disciplinary content". In: *Studies in Higher Education*. DOI: 10.1080/03075079.2026.2718800.

Fong, C. J., C. Gonzales, C. Hill-Troglin Cox, and H. B. Shinn (2023). "Academic Help-Seeking and Achievement of Postsecondary Students: A Meta-Analytic Investigation". In: *Journal of Educational Psychology* 115.1, pp. 1–21. DOI: 10.1037/edu0000725.

Gneezy, U., S. Meier, and P. Rey-Biel (2011). "When and Why Incentives (Don't) Work to Modify Behavior". In: *Journal of Economic Perspectives* 25.4, pp. 191–210. DOI: 10.1257/jep.25.4.191.

Herpen, S. G. A. van, M. Meeuwisse, W. H. A. Hofman, and S. E. Severiens (2020). "A head start in higher education: the effect of a transition intervention on interaction, sense of belonging, and academic performance". In: *Studies in Higher Education* 45.4, pp. 862–877. DOI: 10.1080/03075079.2019.1572088.

Jacob, M. and E. Ní Fhloinn (2019). “A quantitative, longitudinal analysis of the impact of mathematics support in an Irish university”. In: *Teaching Mathematics and its Applications* 38.4, pp. 216–229. DOI: 10.1093/teamat/hry012.

Kahu, E. R. (2013). “Framing student engagement in higher education”. In: *Studies in Higher Education* 38.5, pp. 758–773. DOI: 10.1080/03075079.2011.598505.

Kahu, E. R. and K. Nelson (2018). “Student engagement in the educational interface: understanding the mechanisms of student success”. In: *Higher Education Research & Development* 37.1, pp. 58–71. DOI: 10.1080/07294360.2017.1344197.

Lawson, D., M. Grove, and T. Croft (2020). “The evolution of mathematics support: a literature review”. In: *International Journal of Mathematical Education in Science and Technology* 51.8, pp. 1224–1254. DOI: 10.1080/0020739X.2019.1662120.

Mac an Bhaird, C., O. Fitzmaurice, E. Ní Fhloinn, and C. O’Sullivan (2013). “Student non-engagement with mathematics learning supports”. In: *Teaching Mathematics and its Applications* 32, pp. 191–205. DOI: 10.1093/teamat/hrt018.

Mullen, C., E. Howard, and A. Cronin (2024). “A scoping literature review of the impact and evaluation of mathematics and statistics support in higher education”. In: *Educational Studies in Mathematics* 117, pp. 1–22. DOI: 10.1007/s10649-024-10332-6.

Paloyo, A. R., S. Rogan, and P. Siminski (2016). “The effect of supplemental instruction on academic performance: An encouragement design experiment”. In: *Economics of Education Review* 55, pp. 57–69. DOI: 10.1016/j.econedurev.2016.08.005.

Pugatch, T. and N. Wilson (2018). “Nudging study habits: A field experiment on peer tutoring in higher education”. In: *Economics of Education Review* 62, pp. 151–161. DOI: 10.1016/j.econedurev.2017.11.003.