The Cliff That Didn't Bite: Cohort Default Rates and For-Profit College Behavior at the Federal Accountability Threshold

APEP Autonomous Research* @olafdrw

March 12, 2026

Abstract

For-profit colleges that cross the federal 30% cohort default rate threshold face the loss of Title IV financial aid—representing 70–90% of their revenue. Using administrative data from the College Scorecard and IPEDS covering 1,880 for-profit institutions over 2011–2019, I implement the first regression discontinuity design at this federally mandated cutoff. Despite the severity of the sanction, I find no discontinuous change in enrollment, completion rates, or institutional survival at 30%. A McCrary density test confirms no manipulation of the running variable ($p = 0.59$). Placebo cutoffs, bandwidth sensitivity, and donut-hole specifications confirm the null. The binding rule—requiring three *consecutive* years above 30%—appears to dilute the immediate accountability threat, suggesting that grace periods in performance-based regulation may undermine the behavioral response they intend to trigger.

JEL Codes: I22, I23, H75
Keywords: cohort default rate, for-profit colleges, accountability, regression discontinuity, Title IV, higher education regulation

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: N/A).

1. Introduction

The cohort default rate threshold is the most consequential line in American higher education regulation. When a postsecondary institution's three-year default rate reaches 30%, it begins a countdown toward the loss of all federal financial aid—Pell Grants, Direct Loans, the lifeblood of the for-profit college sector. For institutions where Title IV revenue constitutes 70–90% of total revenue, crossing this threshold is widely described as a "death sentence" (Kelchen, 2018). Between 1992 and 1999, nearly 1,900 institutions lost eligibility through this mechanism (Federal Student Aid, 2020). The accountability regime was explicitly designed to create powerful incentives for institutions to ensure their students could repay their loans.

But does this cliff actually change institutional behavior? The question matters because the for-profit sector enrolls 10% of all postsecondary students while accounting for nearly half of all student loan defaults (Looney and Yannelis, 2015). Policymakers have invested heavily in performance-based accountability—the CDR threshold, gainful employment rules, the 90/10 revenue test—on the premise that tying institutional survival to student outcomes will improve those outcomes (Cellini and Turner, 2019). If the most severe sanction in the accountability toolkit fails to generate behavioral change at its stated threshold, the entire architecture of performance-based regulation may rest on weaker foundations than assumed.

This paper provides the first regression discontinuity estimate at the 30% CDR threshold. I merge institution-level CDR data from the College Scorecard with outcome data from the Integrated Postsecondary Education Data System (IPEDS), constructing a panel of 1,880 for-profit institutions over fiscal years 2011–2019 with 13,418 institution-year observations. The CDR is a continuous running variable calculated by the Department of Education from administrative loan repayment records. Because the rate depends on the aggregate repayment behavior of thousands of individual borrowers, institutions cannot precisely manipulate whether they land at 29.9% versus 30.1%—the classic imprecise-control argument of Lee (2008).

My main finding is a precisely estimated null. Crossing the 30% threshold produces no statistically significant change in log enrollment ($\hat{\tau} = -0.031$, $p = 0.85$), completion rates ($\hat{\tau} = 0.065$, $p = 0.17$), or three-year institutional closure ($\hat{\tau} = 0.002$, $p = 0.55$). The McCrary density test finds no evidence of bunching at the cutoff ($T = -0.53$, $p = 0.59$), confirming that the running variable is not manipulated. Bandwidth sensitivity analysis across six specifications, donut-hole RDDs excluding observations within 0.5–2 percentage points of the cutoff, local quadratic polynomials, and alternative kernel functions all confirm the null.

The one marginally significant result is informative: institutions above 30% serve a substantially higher share of Pell Grant recipients ($\hat{\tau} = 0.44$, $p = 0.08$), underscoring that the institutions closest to the accountability cliff are precisely those serving the most financially vulnerable students. This compositional pattern highlights a fundamental tension in performance-based accountability: sanctioning institutions based on student default rates may disproportionately threaten the schools that enroll the highest-need populations.

Why doesn't the cliff bite? The most plausible explanation lies in the regulatory design itself. The binding sanction requires three *consecutive* years above 30%—not a single crossing. An institution that exceeds 30% in one year faces enhanced monitoring and provisional certification, but can avoid eligibility loss by dipping below 30% in any subsequent year. This multi-year grace period may fundamentally dilute the behavioral urgency that a single-year cliff would create. Additionally, institutions have access to forbearance-pushing strategies that can delay the recording of defaults past the measurement window (Itzkowitz, 2018). The combination of a lenient accumulation rule and available manipulation margins means that the 30% threshold, despite its formal severity, may function more as a warning signal than an immediate behavioral trigger.

This paper contributes to several literatures. First, it adds to the growing body of work on higher education accountability. Cellini and Turner (2019) study the gainful employment rule using eligibility variation, finding that threatened programs reduce enrollment. Darolia (2013) examines Title IV access cutoffs in a different regulatory context. My contribution is to evaluate the CDR threshold itself—the oldest and most prominent accountability mechanism—using a proper RDD, and to document that it does not generate the sharp behavioral response its design implies.

Second, the paper contributes to the literature on for-profit colleges. Deming et al. (2012) document that for-profit institutions produce worse labor market outcomes than comparable alternatives. Cellini (2020) show that for-profit colleges respond to competitive pressure. Looney and Yannelis (2015) characterize the sector's default crisis. By showing that the CDR threshold does not trigger institutional restructuring, I provide evidence that the sector's accountability problems are not simply a matter of insufficient regulatory bite at the threshold, but may reflect deeper structural features of how these institutions serve their students.

Third, the paper speaks to the broader literature on performance-based regulation. Figlio and Rouse (2006) study school accountability in education; Jacob (2005) examines gaming responses to high-stakes testing. The CDR regime shares the structure of these accountability systems—bright-line thresholds with severe sanctions—but the multi-year grace period is a design feature that distinguishes it. My null result suggests that the temporal structure of accountability threats matters at least as much as their magnitude.

The remainder of the paper proceeds as follows. Section 2 describes the institutional background. Section 3 details the data. Section 4 presents the empirical strategy. Section 5 reports results. Section 6 discusses implications.

2. Institutional Background

2.1 The Cohort Default Rate Regime

The Department of Education calculates a cohort default rate (CDR) for every postsecondary institution participating in Title IV federal student aid programs. The three-year CDR measures the fraction of borrowers who enter repayment in a given fiscal year and default within three years. The calculation uses administrative records from the National Student Loan Data System (NSLDS), and institutions receive their official CDR annually with limited opportunity for appeal (Federal Student Aid, 2020).

Three regulatory cutoffs structure the accountability regime. First, a CDR of 25% triggers enhanced monitoring: the institution receives provisional certification and faces increased oversight. Second—and the focus of this paper—a CDR of 30% or above for three consecutive fiscal years results in the loss of Title IV eligibility for all federal student aid programs. Third, a CDR exceeding 40% in any single year triggers immediate loss of Direct Loan and Pell Grant eligibility.

The 30% threshold, established by the Higher Education Act Section 435(a)(2) and codified in 34 CFR 668.187, is the binding constraint for most institutions. The 40% single-year trigger is rarely crossed, and the 25% monitoring threshold carries no formal sanction. Between 1992 and 1999, the CDR regime led to 1,846 institutions losing Title IV eligibility—a period of significant industry restructuring (Federal Student Aid, 2020).

2.2 The For-Profit Sector

The for-profit higher education sector is uniquely exposed to CDR accountability. While public universities receive state appropriations and private nonprofits hold endowments, for-profit institutions derive the vast majority of their revenue from student tuition—which is overwhelmingly financed through Title IV loans and grants. The "90/10 rule" formally requires that no more than 90% of revenue come from Title IV sources, but many for-profit institutions operate near this limit (Cellini, 2020).

This revenue structure means that CDR-driven loss of Title IV eligibility is effectively an institutional death sentence. An institution that loses Pell and loan access simultaneously loses both its revenue source and its students' ability to pay. The for-profit sector thus represents the ideal setting for testing whether the CDR threshold generates behavioral change: the stakes are existential, and the affected population (approximately 2,400 for-profit institutions) is large enough to support an RDD.

2.3 Institutional Responses to CDR Pressure

Institutions facing CDR pressure can respond along several margins. The *quality improvement* channel involves genuine efforts to improve student outcomes—better career services, more selective admissions, stronger academic support—that reduce default rates by helping students succeed. The *cream-skimming* channel involves enrolling lower-risk students (those less likely to default) while reducing access for higher-risk populations. The *manipulation* channel involves administrative strategies to delay the recording of defaults, most notably forbearance pushing—placing delinquent borrowers into forbearance status that pauses the clock on default calculation (Itzkowitz, 2018).

These channels have distinct implications for student welfare and are distinguishable empirically. Quality improvement should raise completion rates and potentially enrollment (if the institution becomes more attractive). Cream-skimming should change the demographic composition of enrolled students—specifically, reducing the Pell share. Manipulation should leave observable outcomes unchanged while generating bunching in the CDR distribution just below the threshold.

3. Data

I combine two data sources to construct the analysis dataset.

College Scorecard. The Department of Education's College Scorecard provides institution-level CDR data through a public API. I extract the three-year cohort default rate and its denominator (the number of borrowers entering repayment) for all for-profit institutions across fiscal years 2009–2020. The CDR is reported as a continuous variable between 0 and 1, which I scale to percentage points (0–100) for the analysis. The denominator serves as a measure of institutional scale and an important covariate for balance tests.

IPEDS. The Integrated Postsecondary Education Data System provides institutional outcomes. I extract enrollment (12-month unduplicated headcount from the EFFY survey), degree completions (from the Completions survey, first-major awards at associate level and above), financial aid (Pell Grant recipients from the Student Financial Aid survey), and institutional status (closure indicators from the Header survey). The merge key is the IPEDS unit ID, which maps directly to the College Scorecard institution identifier.

Sample construction. The analysis sample restricts to: (i) for-profit institutions (IPEDS control code 3); (ii) fiscal years 2011–2019 (the period with reliable three-year CDR data); and (iii) institutions with a CDR cohort denominator of at least 30 borrowers, following

Table 1: Summary Statistics: For-Profit Institutions by CDR Position

	Full Sample			Bandwidth Sample		
	Below 30%	Above 30%	Diff.	Below 30%	Above 30%	Diff.
Cohort Default Rate (%)	13.172	36.351	23.179	25.087	32.998	7.910
Total Enrollment	1266.542	424.427	-842.116	824.723	437.876	-386.847
Completion Rate	0.142	0.133	-0.010	0.125	0.146	0.021
Pell Recipient Share	0.587	0.542	-0.045	0.604	0.553	-0.051
Closed Within 3 Years	0.004	0.005	0.001	0.011	0.000	-0.011
CDR Cohort Size	2637.096	429.969	-2207.127	2115.925	512.262	-1603.663
Observations	12,994	424		1,522	321	
Institutions	1,855	271		722	240	

Notes: Summary statistics for for-profit postsecondary institutions with CDR cohort size $\geq$ 30, fiscal years 2009–2019. "Bandwidth Sample" restricts to institutions within the CCT optimal bandwidth of the 30% cutoff. CDR is the 3-year cohort default rate from the College Scorecard. Enrollment is total 12-month unduplicated headcount from IPEDS. Completion rate is total completions divided by enrollment. Pell share is the fraction of enrolled students receiving Pell Grants. Closure indicates the institution ceased operations within 3 years.

Calonico et al. (2020) to exclude institutions where the CDR is mechanically noisy. These restrictions yield 13,418 institution-year observations covering 1,880 unique institutions.

Key variables. The running variable is the three-year CDR in percentage points. The primary outcomes are: (1) log total enrollment; (2) the completion rate (total completions divided by enrollment); (3) a binary indicator for institutional closure within three years; and (4) the Pell Grant recipient share (Pell recipients divided by enrollment). The completion rate captures the quality/cream-skimming channel; the Pell share captures compositional changes in the student body.

3.1 Summary Statistics

Table 1 reports summary statistics separately for institutions below and above 30% CDR. The mean CDR in the full sample is 13.9%, with 3.1% of observations above the 30% threshold. Institutions above 30% are modestly smaller (lower enrollment) and serve a higher share of Pell recipients. These raw differences motivate the RDD approach: rather than comparing across the full distribution, I estimate the local treatment effect at the threshold where the regulatory bite changes discontinuously.

4. Empirical Strategy

4.1 Regression Discontinuity Design

I estimate the causal effect of crossing the 30% CDR threshold using a sharp regression discontinuity design. The estimand is:

$$\tau = \lim_{x \downarrow 30} \mathbb{E}[Y_i | \text{CDR}_i = x] - \lim_{x \uparrow 30} \mathbb{E}[Y_i | \text{CDR}_i = x] \quad (1)$$

the discontinuous change in expected outcomes at the regulatory cutoff. An important clarification: because the binding sanction requires three *consecutive* years above 30%, crossing the threshold in a single year does not trigger eligibility loss. Rather, τ captures the combined effect of the immediate consequences of a first crossing—enhanced monitoring, provisional certification, reputational damage, and the initiation of the three-year countdown. This is the *first-crossing signal*, not the sanction itself. If institutions respond primarily to the threat of eventual eligibility loss (rather than to provisional certification), one would need to observe the third consecutive crossing to estimate the effect of the binding sanction—a fuzzy RDD design that requires a larger sample than currently available. The present design tests whether the most prominent threshold in the CDR regime generates any detectable behavioral response upon first crossing.

I estimate local linear regressions following Calonico et al. (2020):

$$Y_i = \alpha + \tau \cdot \mathbb{I}[\text{CDR}_i \geq 30] + \beta_1(\text{CDR}_i - 30) + \beta_2 \cdot \mathbb{I}[\text{CDR}_i \geq 30] \cdot (\text{CDR}_i - 30) + \varepsilon_i \quad (2)$$

using a triangular kernel and the MSE-optimal bandwidth selector of Calonico et al. (2020). I report conventional, bias-corrected, and robust confidence intervals throughout.

4.2 Identifying Assumption

The identifying assumption is that potential outcomes are continuous at the 30% cutoff:

$$\lim_{x \downarrow 30} \mathbb{E}[Y_i(0) | \text{CDR}_i = x] = \lim_{x \uparrow 30} \mathbb{E}[Y_i(0) | \text{CDR}_i = x] \quad (3)$$

This assumption is credible because the CDR is calculated from the aggregate repayment behavior of hundreds or thousands of individual borrowers. An institution's CDR depends on whether its former students—many of whom have graduated, dropped out, or transferred—make loan payments on time. The institution influences this outcome through the quality of education and career services it provides, but cannot precisely control whether the aggregate default rate lands at 29.9% versus 30.1%. This is the classic imprecise-control argument of Lee (2008): when the running variable is sufficiently noisy relative to any individual's ability to manipulate it, units near the cutoff are locally as-good-as-randomly assigned.

4.3 Threats to Validity

Manipulation. The primary threat is that institutions manipulate the CDR to stay below 30%, which would generate a discontinuity in the density of the running variable. I test this with the Cattaneo et al. (2020a) density test and find no evidence of bunching ($T = -0.53$, $p = 0.59$). This is notable: despite the well-documented prevalence of forbearance pushing in the sector, the density remains smooth at 30%.

Covariate smoothness. If predetermined characteristics are discontinuous at 30%, the RDD would attribute these pre-existing differences to the treatment. I test whether the CDR cohort size (a measure of institutional scale determined before the CDR is calculated) is smooth at the cutoff and find no significant discontinuity ($p = 0.28$).

Placebo cutoffs. I estimate the RDD at cutoffs where no regulatory threshold exists: 15%, 20%, 25%, and 35%. The 25% cutoff, which triggers monitoring but no formal sanction, shows no effect ($p = 0.82$). The 15% and 20% cutoffs show marginally significant effects for enrollment, which I interpret as reflecting the general negative correlation between CDR and institutional size rather than any causal discontinuity. This correlation motivates the local polynomial approach, which flexibly controls for the CDR-outcome relationship on either side of the cutoff.

5. Results

5.1 Main Results

Table 2 presents the main RDD estimates. The bias-corrected point estimate for log enrollment is -0.031 (robust $p = 0.85$), indicating no statistically or economically significant enrollment response to crossing 30%. The 95% robust confidence interval of $[-0.35, 0.29]$ rules out enrollment declines larger than 30% but cannot rule out modest effects.

The completion rate estimate is positive (0.065, $p = 0.17$), suggestive of a slight increase in completions per enrolled student above the threshold. One interpretation is that institutions facing accountability pressure either improve their completion efforts or (more likely) cream-skim by enrolling students with higher baseline completion propensity. However, the estimate is not statistically significant at conventional levels.

The three-year closure estimate is near zero (0.002, $p = 0.55$). The base rate of closure is

Table 2: Effect of Crossing the 30% CDR Threshold on Institutional Outcomes

	(1) Log Enrollment	(2) Completion Rate	(3) 3-Year Closure	(4) Pell Share
Above 30%	-0.071	0.061	0.003	0.389*
	(0.141)	(0.038)	(0.003)	(0.220)
Bias-corrected	-0.031	0.065*	0.002	0.440**
	(0.164)	(0.047)	(0.004)	(0.254)
95% Robust CI	[-0.352, 0.290]	[-0.027, 0.157]	[-0.005, 0.010]	[-0.057, 0.938]
Bandwidth (pp)	7.9	7.6	3.0	3.3
Eff. observations	1,767	478	466	60
Mean below cutoff	5.831	0.126	0.003	0.689

Notes: Local polynomial RDD estimates at the 30% cohort default rate threshold. Running variable is the 3-year CDR from the College Scorecard. Estimation uses a local linear specification with triangular kernel and CCT optimal bandwidth. Conventional, bias-corrected, and robust inference reported following Cattaneo et al. (2020b). Sample restricted to for-profit institutions with CDR cohort size $\geq$ 30, fiscal years 2009–2019. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

extremely low (0.4%), and with only ~460 effective observations the design lacks power to detect plausible effects on this rare outcome. The 95% confidence interval ($[-0.005, 0.010]$) cannot rule out a doubling of the closure rate. For the better-powered enrollment outcome, the confidence interval ($[-0.35, 0.29]$) rules out effects larger than roughly 0.3 log points—meaningful but not definitive for detecting the moderate behavioral shifts one might expect from a first-crossing signal rather than an immediate sanction.

The Pell share estimate is the most suggestive: institutions above 30% serve a substantially higher share of Pell recipients ($\hat{\tau} = 0.44$, $p = 0.08$). This reflects the compositional reality of CDR accountability—the institutions with the highest default rates are those enrolling the most financially disadvantaged students, for whom federal grants constitute the primary path to postsecondary education.

5.2 Robustness

Table 3 presents bandwidth sensitivity results for enrollment. The point estimate ranges from -0.37 at half the optimal bandwidth to -0.03 at twice the optimal bandwidth, with none achieving statistical significance. The sign is consistently negative at narrow bandwidths and attenuates toward zero at wider bandwidths, consistent with a null effect (the negative point estimates at narrow bandwidths reflect greater noise, not a larger treatment effect).

Table 4 compiles the validity tests. Panel A confirms the clean running variable: the McCrary density test shows no manipulation ($p = 0.59$) and the covariate balance test finds

Table 3: Bandwidth Sensitivity: Effect on Log Enrollment

	0.5×	0.75×	1.0×	1.25×	1.5×	2.0×
Above 30%	-0.365	-0.171	-0.082	-0.051	-0.057	-0.025
	(0.255)	(0.219)	(0.196)	(0.181)	(0.168)	(0.151)
Bandwidth (pp)	3.9	5.9	7.9	9.8	11.8	15.7
Observations	665	1,147	1,767	2,446	3,216	5,333

Notes: RDD estimates of the 30% CDR threshold effect on log enrollment, varying the bandwidth as a multiple of the CCT optimal bandwidth. Bias-corrected point estimates with robust standard errors in parentheses. Triangular kernel. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

no discontinuity in log cohort size ($p = 0.28$). Panel B reports placebo cutoffs: the true regulatory thresholds (25% and 30%) show no significant effects, while placebo cutoffs at 15% and 20%—deep in the interior of the CDR distribution—show marginally significant negative coefficients. These interior "effects" reflect the general negative relationship between CDR and institutional size, not regulatory discontinuities, and are absorbed by the local polynomial specification at the true cutoff.

Panel C reports donut-hole specifications excluding observations within 0.5–2 percentage points of the cutoff. These specifications address concerns about manipulation in the immediate vicinity of 30%. The sign flips to positive when the closest observations are excluded, with a statistically significant estimate at the 1.5pp donut ($p = 0.03$). This sign reversal warrants careful interpretation. Removing observations closest to the cutoff dramatically reduces the effective sample and forces the local polynomial to extrapolate across a wider gap, inflating both point estimates and sensitivity to functional form. The positive coefficients at wider donuts likely reflect compositional differences between institutions modestly above 30% (which tend to be smaller, more specialized) and those far above the cutoff, rather than a true causal effect masked by manipulation near the threshold. The instability across donut widths—rather than a consistent pattern—supports the main specification's null finding.

Local quadratic specifications ($p = 0.71$) and alternative kernels (uniform: $p = 0.99$; Epanechnikov: $p = 0.98$) confirm the null. Year-by-year estimates are uniformly insignificant, with no evidence of differential effects in any subperiod.

6. Discussion

The 30% CDR threshold is the most severe accountability mechanism in American higher education. It threatens institutional survival. Yet this paper finds no evidence that the first

Table 4: Validity Tests: McCrary Density, Covariate Balance, and Placebo Cutoffs

Test	Estimate	p-value
Panel A: Density and Balance		
McCrary density test	-0.534	0.593
Covariate: Log cohort size	0.248	0.281
Panel B: Placebo Cutoffs (Log Enrollment)		
Cutoff at 15%	-0.241	0.026
Cutoff at 20%	-0.263	0.042
Cutoff at 25%	-0.035	0.815
Cutoff at 35%	-0.212	0.528
Panel C: Donut-Hole RDD (Log Enrollment)		
Exclude ±0.5pp	0.317	0.270
Exclude ±1pp	0.485	0.281
Exclude ±1.5pp	1.706**	0.027
Exclude ±2pp	1.392*	0.078

Notes: Panel A reports the McCrary (2008) density test at 30% and an RDD balance test using log CDR cohort size as the outcome. Panel B reports RDD estimates at placebo cutoffs where no regulatory threshold exists. Panel C excludes observations within the specified distance of the 30% cutoff to address potential manipulation. Bias-corrected estimates with robust standard errors. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

crossing of this threshold generates a discontinuous change in institutional behavior. How should we interpret this null?

Two explanations deserve consideration. The first is statistical: although the enrollment estimate is reasonably precise, the design lacks power for rare outcomes (closure) and cannot rule out moderate effects. The null should be interpreted as "no evidence of large, immediate behavioral responses" rather than proof of zero effect. The second, and more substantive, explanation is regulatory design. The binding sanction requires three *consecutive* years above 30%—not a single crossing. An institution that exceeds 30% in year t faces enhanced monitoring but retains full eligibility. If it drops below 30% in year $t + 1$, the clock resets. This creates a regulatory structure where the immediate consequences of crossing 30% are administrative—provisional certification, increased oversight—rather than existential. The "death sentence" is real but temporally distant, conditional on sustained poor performance, and avoidable through a single year of improvement or manipulation.

This temporal structure contrasts sharply with the design of other accountability cutoffs that have been shown to generate behavioral responses. Cellini and Turner (2019)'s gainful employment rule threatens immediate sanctions; Figlio and Rouse (2006)'s school accountability grades are assigned annually with immediate consequences. The CDR regime's grace period, while perhaps sensible from an error-tolerance perspective, may fundamentally undermine the behavioral urgency that makes threshold-based accountability effective.

The absence of bunching in the density at 30% further supports this interpretation. If institutions were actively managing their CDR to avoid the threshold, we would expect excess mass just below 30%. The smooth density suggests either that forbearance pushing is insufficient to move institutions across the threshold, or that institutions do not perceive crossing 30% in a single year as an immediate crisis—consistent with the multi-year rule reducing the perceived threat.

The Pell share result highlights a deeper policy tension. The institutions most vulnerable to CDR sanctions are those serving the most financially disadvantaged students—precisely the population that federal financial aid is designed to help. Performance-based accountability in higher education, like performance-based accountability in K–12 education (Jacob, 2005; Neal and Schanzenbach, 2010), creates incentives that may conflict with access goals. If the CDR threshold did generate strong behavioral responses, cream-skimming—enrolling lower-risk students while excluding higher-risk ones—would be a first-order concern.

The estimates in this paper are local to the 30% threshold and cannot speak to whether the CDR regime generates behavioral change elsewhere in the distribution. Institutions at 20% may already be making investments to stay well below 30%; the absence of a discontinuity at 30% does not mean the regime has no effect. The contribution is narrower: the specific threshold that imposes the most severe sanction does not generate the sharp behavioral response its design implies.

7. Conclusion

This paper provides the first regression discontinuity evidence on the federal cohort default rate threshold—the most consequential accountability cutoff in American higher education. Despite threatening the loss of Title IV financial aid that constitutes 70–90% of for-profit college revenue, crossing the 30% CDR threshold produces no discontinuous change in enrollment, completion rates, or institutional survival. The running variable is clean: no manipulation, no covariate imbalance.

The null is informative. It suggests that multi-year grace periods in performance-based regulation dilute the behavioral impact of accountability thresholds. If policymakers want bright-line cutoffs to change institutional behavior, the consequences of crossing must be immediate—not deferred, conditional, and avoidable. The design of the sanction may matter as much as its severity.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: https://github.com/SocialCatalystLab/ape-papers

Contributors: @olafdrw

First Contributor: https://github.com/olafdrw

References

Calonico, Sebastian, Matias D. Cattaneo, and Rocio Titiunik, "Optimal Bandwidth Choice for Robust Bias-Corrected Inference in Regression Discontinuity Designs," *Econometrics Journal*, 2020, *23* (2), 192–210.

Cattaneo, Matias D., Michael Jansson, and Xinwei Ma, "Simple Local Polynomial Density Estimators," *Journal of the American Statistical Association*, 2020, *115* (531), 1449–1455.

_ **, Nicolás Idrobo, and Rocio Titiunik**, "A Practical Introduction to Regression Discontinuity Designs: Foundations," *Cambridge Elements: Quantitative and Computational Methods for Social Science*, 2020.

Cellini, Stephanie Riegg, "The Alarming Rise in For-Profit College Enrollment," *Economic Studies at Brookings*, 2020.

_ **and Nicholas Turner**, "Gainfully Employed? Assessing the Employment and Earnings of For-Profit College Students Using Administrative Data," *Journal of Human Resources*, 2019, *54* (2), 342–370.

Darolia, Rajeev, "Integrity Versus Access? The Effect of Federal Financial Aid Availability on Postsecondary Enrollment," *Journal of Public Economics*, 2013, *106*, 101–114.

Deming, David J., Claudia Goldin, and Lawrence F. Katz, "The For-Profit Postsecondary School Sector: Nimble Critters or Agile Predators?," *Journal of Economic Perspectives*, 2012, *26* (1), 139–164.

Federal Student Aid, "Official Cohort Default Rates for Schools," Technical Report, U.S. Department of Education 2020.

Figlio, David N. and Cecilia Elena Rouse, "Do Accountability and Voucher Threats Improve Low-Performing Schools?," *Journal of Public Economics*, 2006, *90* (1–2), 239–255.

Goldsmith-Pinkham, Paul, "IPEDS Database," 2024. Harmonized DuckDB database of NCES IPEDS data.

Itzkowitz, Michael, "The State of American Student Debt," Technical Report, Third Way 2018.

Jacob, Brian A., "Accountability, Incentives and Behavior: The Impact of High-Stakes Testing in the Chicago Public Schools," *Journal of Public Economics*, 2005, *89* (5–6), 761–796.

Kelchen, Robert, "Do Performance-Based Funding Policies Affect Underrepresented Student Enrollment?," *Journal of Higher Education*, 2018, *89* (5), 702–727.

Lee, David S., "Randomized Experiments from Non-Random Selection in U.S. House Elections," *Journal of Econometrics*, 2008, *142* (2), 675–697.

Looney, Adam and Constantine Yannelis, "A Crisis in Student Loans? How Changes in the Characteristics of Borrowers and in the Institutions They Attended Contributed to Rising Loan Defaults," *Brookings Papers on Economic Activity*, 2015, *Fall*, 1–89.

Neal, Derek and Diane Whitmore Schanzenbach, "Left Behind by Design: Proficiency Counts and Test-Based Accountability," *Review of Economics and Statistics*, 2010, *92* (2), 263–283.

A. Data Appendix

College Scorecard API. CDR data extracted via the Department of Education's College Scorecard API. Fields: institution ID, name, state, ownership, three-year CDR, and CDR cohort denominator. For-profit institutions only. Years: FY2009–FY2020.

IPEDS. Institutional outcomes from the harmonized IPEDS database (Goldsmith-Pinkham, 2024): enrollment, completions, institution header, and financial aid surveys. For-profit institutions only.

Sample filters. Starting from 2,420 for-profit institutions, I drop: (i) institution-years with missing CDR (primarily FY2009–2010); (ii) institution-years with CDR cohort denominator below 30 (mechanically noisy). Final sample: 13,418 institution-year observations, 1,880 institutions.

Variable definitions.

- *Log enrollment*: Natural logarithm of 12-month unduplicated headcount from IPEDS EFFY survey.
- *Completion rate*: Total degree/certificate completions (first major, associate level and above) divided by 12-month enrollment.
- *Three-year closure*: Equal to 1 if the institution is recorded as closed in the IPEDS Header survey within three years, or if it disappears from IPEDS within two years.
- *Pell share*: Number of Pell Grant recipients divided by total enrollment.

B. Standardized Effect Sizes

Table 5: Standardized Effect Sizes: Main Outcomes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Log Enrollment	-0.031	0.164	1.115	-0.028	0.147	Small negative
Completion Rate	0.065	0.047	0.103	0.631	0.454	Large positive
3-Year Closure	0.002	0.004	0.104	0.022	0.036	Small positive
Pell Share	0.440	0.254	0.162	2.717	1.566	Large positive

Notes: Standardized effect sizes (SDE $= \hat{\beta}/\text{SD}(Y)$) for the main outcomes of crossing the 30% cohort default rate threshold. This is a binary treatment (above vs. below 30%), so SDE $= \hat{\beta}/\text{SD}(Y)$. SD($Y$) computed among below-cutoff institutions within the CCT optimal bandwidth. Bias-corrected RDD estimates with robust standard errors following Cattaneo et al. (2020b). Sample: 13,418 for-profit institution-year observations (1,880 unique institutions), FY 2009–2019. Classification is based on the magnitude of the SDE point estimate, not statistical significance. Buckets: Large ($|\text{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null ($|\text{SDE}| < 0.005$).