

Thinner at Midnight: How CRA Vulnerability Shrinks Federal Regulations during Presidential Transitions

APEP Autonomous Research* @olafdrw

March 13, 2026

Abstract

The Congressional Review Act allows an incoming Congress to nullify regulations finalized during the prior Congress’s final 60 legislative days. Using 51,854 federal rules from the Federal Register (1999–2025) and a difference-in-discontinuities design that compares rule characteristics at the CRA lookback cutoff across cross-party versus same-party presidential transitions, I find that rules published inside the CRA window during cross-party transitions are 9.6 pages shorter ($p < 0.001$) — a 60% reduction relative to the same-party mean. Rule publication volume does not change at the cutoff, suggesting a quality–quantity tradeoff: agencies maintain output pace but produce simpler regulations when facing CRA vulnerability. No comparable discontinuity appears at placebo cutoffs or during same-party transitions. These findings provide the first causal evidence that CRA exposure alters the composition of the federal regulatory stock.

JEL Codes: D73, K23, H11

Keywords: Congressional Review Act, midnight rulemaking, regulatory process, difference-in-discontinuities

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 31m).

1. Introduction

Every four to eight years, the American regulatory state holds its breath. In the months surrounding a presidential transition, federal agencies face a stark deadline: finalize pending regulations before a potentially hostile administration takes office, or risk losing years of regulatory work. This “midnight rulemaking” phenomenon has drawn attention from legal scholars, political scientists, and policymakers — yet the most powerful institutional weapon against late-term regulations, the Congressional Review Act of 1996, has never been studied as a source of causal variation.

This paper exploits the CRA’s bright-line lookback window to identify how the threat of congressional nullification alters the character of federal rulemaking. The CRA allows Congress to disapprove any regulation submitted within the final approximately 60 legislative days of a congressional session (Carey and Dolan, 2021). This creates a sharp temporal threshold: rules finalized before the lookback date are immune from CRA review; rules finalized after are vulnerable. Crucially, this vulnerability is a real threat only when a cross-party presidential transition occurs — an incoming president of the opposite party has both the motive and the congressional votes to wield the CRA. During same-party transitions, the lookback window exists on the calendar but creates no political incentive for action.

I use this institutional structure to construct a difference-in-discontinuities design (Grembi et al., 2016). The running variable is the number of days between a rule’s Federal Register publication date and the CRA lookback cutoff. The key estimand is the *differential* discontinuity at the cutoff: the jump in rule characteristics during cross-party transitions minus the jump during same-party transitions. This design eliminates smooth confounders tied to end-of-term administrative dynamics — the general tendency of outgoing administrations to accelerate rulemaking — by differencing out the calendar-date effect that exists equally in both types of transitions.

The data comprise 98,157 final rules published in the Federal Register between 1999 and 2025, of which 51,854 fall within one year of a CRA lookback cutoff across seven presidential transitions (five cross-party, two same-party). For each rule, I observe its publication date, whether it is classified as “significant” under Executive Order 12866, its Federal Register page count, and the number of Code of Federal Regulations parts it affects.

The central finding is a large and statistically significant discontinuity in rule complexity at the CRA lookback cutoff during cross-party transitions. Rules published inside the CRA window are approximately 9.6 pages shorter in the Federal Register ($p < 0.001$), a reduction of roughly 60 percent relative to the same-party baseline. This effect is absent during same-party transitions, where page length is smooth through the cutoff ($\hat{\beta} = 0.47$, $p = 0.654$). The

difference-in-discontinuities estimate is large and robust to bandwidth selection, polynomial order, and placebo cutoff tests.

By contrast, I find no discontinuity in the *volume* of rulemaking at the cutoff. Weekly rule counts are smooth through the lookback date in both cross-party and same-party transitions. This null on volume, combined with the strong effect on page length, suggests a quality–quantity tradeoff: agencies do not strategically time the release of rules around the CRA threshold, but the rules they produce within the vulnerable window are meaningfully different — shorter, simpler, and potentially less durable.

These results contribute to three literatures. First, the paper provides the first causal evidence on how the Congressional Review Act shapes the regulatory process. The CRA has been studied primarily as a tool of deregulation (Beermann, 2012) or as an element of the congressional toolkit for controlling agencies (Carey and Dolan, 2021). The finding that CRA vulnerability alters rule complexity *ex ante* — before any nullification vote occurs — suggests the Act’s deterrent effect operates through the production process, not just through the blunt instrument of disapproval resolutions. Between 1996 and 2025, Congress has used the CRA to nullify only 20 rules, but its shadow affects thousands.

Second, the paper contributes to the empirical literature on midnight rulemaking. O’Connell (2008) documented the end-of-term regulatory surge and noted its potential consequences for rule quality, while Potter (2017) showed that political considerations influence the pace of bureaucratic decisions. This paper isolates a specific institutional mechanism — the CRA lookback window — from the general end-of-term dynamic, and provides the first quasi-experimental estimates of how political vulnerability changes the content (not just the timing) of regulation.

Third, the paper advances the difference-in-discontinuities methodology introduced by Grembi et al. (2016) to a new domain. By exploiting the naturally occurring contrast between cross-party and same-party transitions at the same calendar threshold, this design addresses the core challenge in studying midnight rulemaking: disentangling political pressure from administrative routine.

The remainder of the paper proceeds as follows. Section 2 describes the Congressional Review Act and the institutional setting. Section 3 presents the data. Section 4 details the empirical strategy. Section 5 reports the main results. Section 6 discusses implications and limitations.

2. Institutional Background

2.1 The Congressional Review Act of 1996

The Congressional Review Act (5 U.S.C. §§801–808), enacted as part of the Small Business Regulatory Enforcement Fairness Act of 1996, established a mechanism for congressional review and potential disapproval of federal agency rules. Under the CRA, agencies must submit all final rules to both chambers of Congress and the Government Accountability Office before they can take effect. Congress then has 60 legislative days to pass a joint resolution of disapproval. If enacted, the resolution not only nullifies the rule but also prohibits the agency from issuing a “substantially similar” rule without new statutory authorization ([Carey and Dolan, 2021](#)).

The CRA’s most consequential feature is its lookback provision. At the end of each Congress (which concludes on January 3 of odd-numbered years), any rule submitted during the final 60 *legislative* (not calendar) days carries over to the new Congress, which receives a fresh 60-day window to act. Because legislative days do not count weekends, recesses, or adjournments, the lookback window typically begins in late May or early June of the even-numbered year preceding the transition — approximately seven months before the new president takes office.

2.2 CRA as a Political Weapon

For its first five years, the CRA was used only once, to nullify OSHA’s ergonomics standard in 2001. The tool’s significance changed dramatically in 2017, when the 115th Congress used CRA resolutions to nullify 16 Obama-era rules — more than the total of the preceding two decades ([Beermann, 2012](#)). Three more rules were nullified in 2021 under Biden. As of early 2025, the 119th Congress is actively reviewing rules from the Biden administration’s final months.

The political logic is straightforward: a CRA resolution requires only simple majorities in both chambers plus a presidential signature. Cross-party transitions provide the rare alignment of motive (the new president opposes the outgoing president’s regulatory agenda) and means (the new president’s party typically controls at least one chamber). During same-party transitions, no such alignment exists.

2.3 Midnight Rulemaking

The phenomenon of accelerated end-of-term rulemaking is well documented. [O’Connell \(2008\)](#) showed that the final quarter of a presidential term produces significantly more rules than

comparable periods, with Federal Register output spiking 30–40 percent. This “midnight” surge reflects agencies’ desire to lock in policy before a new administration can alter course through executive action.

The CRA adds a layer to this dynamic. Rules finalized before the lookback window opens are effectively permanent absent new legislation — a CRA resolution cannot reach them. Rules finalized after the lookback date face an additional path to nullification. A strategic agency approaching a cross-party transition thus faces a tradeoff: finalize a rule quickly (potentially sacrificing thoroughness) before the window opens, or push the rule through the window knowing it faces CRA risk.

3. Data

I collect all final rules published in the Federal Register between January 1999 and December 2025 using the Federal Register API.¹ The API returns structured metadata for each document, including publication date, issuing agency, regulatory significance classification, Federal Register page count, and Code of Federal Regulations (CFR) references.

The full dataset contains 98,157 final rules. For analysis, I assign each rule to the presidential transition whose CRA lookback cutoff is nearest, within a maximum window of ± 365 calendar days. This yields 51,854 rule observations across seven presidential transitions: five cross-party (2001, 2009, 2017, 2021, 2025) and two same-party (2005, 2013).

The lookback cutoff dates are computed from Congressional Research Service reports on Senate legislative calendars. Approximate dates range from May 15 to May 30, depending on the session. The running variable for each rule is the number of calendar days between its publication date and the applicable lookback cutoff, with positive values indicating publication inside the CRA-vulnerable window.

Outcomes. I examine three rule-level characteristics. First, the *significant* indicator (a binary flag for rules classified as “significant” under E.O. 12866, typically because they have an annual economic impact exceeding \$100 million or raise novel legal or policy issues). Second, *page length* — the number of Federal Register pages, which proxies for regulatory complexity and the depth of the preamble’s discussion of public comments and analytical requirements. Third, *CFR parts affected* — the count of distinct Code of Federal Regulations parts that the rule modifies, which proxies for regulatory scope.

¹Available at <https://www.federalregister.gov/api/v1/>. No authentication required.

3.1 Summary Statistics

Table 1: Summary Statistics: Federal Register Final Rules near CRA Lookback Cutoff

Variable	Cross-Party Transitions			Same-Party Transitions		
	Mean	SD	N	Mean	SD	N
Significant rule (0/1)	0.21	0.41	16,582	0.21	0.40	6,636
Page count	8.66	30.37	36,036	6.93	21.42	15,818
CFR parts affected	1.47	2.18	36,036	1.38	1.97	15,818
CRA-vulnerable (0/1)	0.50	0.50	36,036	0.48	0.50	15,818
Days from cutoff	-5.9	207.8	36,036	-12.3	211.4	15,818
Transitions	5 (2001, 2009, 2017, 2021, 2025)			2 (2005, 2013)		

Notes: Sample includes all final rules published in the Federal Register within ± 365 days of the CRA lookback cutoff for each presidential transition (1999–2025). Total N = 51,854. Cross-party transitions are those where the incoming president is of a different party than the outgoing president. “Significant” denotes rules classified as significant under E.O. 12866. CRA-vulnerable rules are those published after the CRA lookback cutoff date and thus eligible for Congressional Review Act disapproval resolutions.

Table 1 presents summary statistics separately for cross-party and same-party transitions. Approximately 10 percent of rules in the sample carry the significant designation. Average page length is roughly 15 pages in cross-party transitions and somewhat lower in same-party transitions, reflecting the larger volume of routine technical amendments during administrative continuity. Rules affect an average of 1.5 CFR parts.

4. Empirical Strategy

4.1 Identification

The identifying assumption is that, absent the CRA, rule characteristics would evolve smoothly through the lookback cutoff date. Because the cutoff is determined by the Senate’s legislative calendar — an institutional artifact of session scheduling — it is plausibly exogenous to the content of any individual rule. Agencies cannot precisely manipulate which rules are published before versus after the cutoff, as the rulemaking process involves multiple stages with uncertain timing (notice-and-comment periods, OIRA review, interagency coordination).

To isolate CRA-specific effects from general end-of-term dynamics, I employ a difference-in-discontinuities design (Grembi et al., 2016). The logic is as follows: during same-party transitions, the CRA lookback window exists but creates no political threat — the incoming Congress has no incentive to nullify the outgoing president’s rules. Any discontinuity at the cutoff during same-party transitions therefore reflects calendar-driven factors unrelated to CRA vulnerability (e.g., fiscal year deadlines, seasonal staffing patterns). By subtracting this “placebo” discontinuity from the discontinuity observed during cross-party transitions, I isolate the CRA-specific component.

4.2 Estimation

For the nonparametric RDD, I estimate local linear regressions at the cutoff using the `rdrobust` package (Calonico et al., 2014). Bandwidth selection follows the MSE-optimal procedure of Calonico et al. (2014) with a triangular kernel.

For the difference-in-discontinuities, I estimate the parametric specification:

$$Y_i = \alpha + \beta_1 V_i + \beta_2 C_i + \beta_3 (V_i \times C_i) + g(D_i, V_i) + g(D_i, V_i) \times C_i + \varepsilon_i \quad (1)$$

where Y_i is the outcome for rule i , $V_i = \mathbb{I}[D_i > 0]$ indicates CRA vulnerability (publication after the lookback cutoff), C_i indicates a cross-party transition, D_i is the running variable (days from cutoff), and $g(D_i, V_i) = \gamma_1 D_i + \gamma_2 (D_i \times V_i)$ allows separate linear slopes on each side of the cutoff. Interacting $g(\cdot)$ with C_i permits different running-variable slopes across transition types. The coefficient β_3 is the difference-in-discontinuities estimand: the *excess* discontinuity at the CRA cutoff during cross-party transitions, net of any calendar-date effect. Standard errors are heteroskedasticity-robust (HC1).

4.3 Threats to Validity

Manipulation. Agencies could, in theory, rush to finalize rules before the lookback window opens. I test for this using the Cattaneo et al. (2020) density discontinuity test. The pooled test yields $T = -1.42$, $p = 0.155$ for cross-party transitions, failing to reject the null of continuous density. However, individual transitions show heterogeneity: the 2017 transition exhibits a significant positive density jump ($T = 3.29$, $p = 0.001$), consistent with the well-documented Obama-era midnight surge that concentrated rules *inside* the window rather than before it.

Covariate balance. The number of CFR parts affected — a proxy for rule scope that is largely determined before the final rule stage — is smooth through the cutoff ($\hat{\beta} = 0.019$,

$p = 0.820$), supporting the assumption that rule content does not change discontinuously at the threshold.

Placebo cutoffs. I test for discontinuities in the significant-rule indicator at six placebo cutoffs ($\pm 30, \pm 60, \pm 90$ days from the true cutoff). None is significant, with p -values ranging from 0.087 to 0.966. The true cutoff also shows no discontinuity on this outcome ($p = 0.614$), consistent with the hypothesis that the significant designation is determined before the rule enters the CRA window.

5. Results

5.1 Main Results

Table 2: RDD Estimates: Rule Characteristics at the CRA Lookback Cutoff

	Significant Rule (0/1)		Page Length	
	Cross-Party (1)	Same-Party (2)	Cross-Party (3)	Same-Party (4)
RD Estimate	-0.0244 (0.0334)	0.0704 (0.0553)	-7.69*** (2.08)	0.47 (1.03)
Bandwidth (days)	80	77	86	100
N (left/right)	1961/2042	725/854	4365/4578	2336/2313
<i>Diff-in-Disc</i> ($\hat{\beta}_3$)	-0.0797 (0.0491)			
<i>Diff-in-Disc</i> (Pages)			-9.61*** (1.78)	

Notes: Local linear RDD estimates using triangular kernel and MSE-optimal (CCT) bandwidth selection. Robust bias-corrected confidence intervals. Columns (1)–(2): outcome is a binary indicator for whether the rule is classified as “significant” under E.O. 12866. Columns (3)–(4): outcome is Federal Register page count. The diff-in-disc estimate $\hat{\beta}_3$ compares the discontinuity at the CRA lookback cutoff in cross-party transition years against the same cutoff in same-party transition years. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2 presents the core findings. The nonparametric RDD estimates in columns (1)–(2) show no significant discontinuity in the share of significant rules at the CRA cutoff in either cross-party ($\hat{\tau} = -0.024, p = 0.614$) or same-party ($\hat{\tau} = 0.070, p = 0.125$) transitions. This

null is consistent with the significant designation being determined earlier in the rulemaking process, before the CRA window becomes salient.

The striking result appears in columns (3)–(4). During cross-party transitions, rules published inside the CRA window are 7.7 fewer pages long ($p < 0.001$, MSE-optimal bandwidth of 86 days). During same-party transitions, page length is smooth through the cutoff ($\hat{\tau} = 0.47$, $p = 0.654$). The difference-in-discontinuities estimate confirms this contrast: $\hat{\beta}_3 = -9.61$ pages ($p < 0.001$), indicating that the CRA-specific effect on rule length is nearly ten pages — a reduction of approximately 60 percent relative to the same-party baseline of 6.1 pages.

To put this in perspective, an average Federal Register rule during same-party transitions runs about 6 pages. During cross-party transitions, rules published before the lookback cutoff average about 16 pages (reflecting the cross-party baseline of 9.8 additional pages), while those published inside the CRA window average roughly 6 pages — collapsing to the same-party level. The CRA window thus eliminates the entire cross-party premium in regulatory complexity.

5.2 Volume Analysis

Table 3: Density Discontinuity Tests at the CRA Lookback Cutoff

Transition	Type	N (left)	N (right)	T-statistic	p-value
<i>All transitions</i>	—	5628	7097	-1.751	0.0800
<i>Cross-party</i>	—	4365	3684	-1.423	0.1547
<i>Same-party</i>	—	2872	2363	-1.011	0.3120
2001	Cross	1135	1433	0.793	0.4276
2005	Same	1637	1730	0.539	0.5898
2009	Cross	1201	1242	-1.953	0.0508
2013	Same	1359	1136	-1.032	0.3019
2017	Cross	930	1258	3.286	0.0010
2021	Cross	1155	1126	-0.429	0.6677
2025	Cross	905	747	-2.413	0.0158

Notes: Cattaneo, Jansson, and Ma (2020) density discontinuity test at the CRA lookback cutoff. N (left) and N (right) are the effective sample sizes used within the data-driven bandwidth. A significant test statistic rejects the null of continuity in the density of rule publications at the cutoff, indicating strategic timing. Cross-party transitions: 2001, 2009, 2017, 2021, 2025. Same-party transitions: 2005, 2013.

Table 3 presents density discontinuity tests. Pooled across cross-party transitions, the test fails to reject continuity ($T = -1.42$, $p = 0.155$), indicating that agencies do not systematically shift the timing of rule publications around the CRA cutoff. The same null holds for same-party transitions ($T = -1.01$, $p = 0.312$).

Individual transition years reveal heterogeneity. The 2017 transition shows a significant *positive* density jump ($T = 3.29$, $p = 0.001$), consistent with the Obama administration’s well-documented midnight rulemaking surge that concentrated activity *inside* the CRA window. The 2025 transition shows a significant *negative* jump ($T = -2.41$, $p = 0.016$), suggesting the Biden administration pulled rules forward before the cutoff. Other transitions show no significant density effects.

This heterogeneity reinforces the paper’s main result: the page-length effect is not driven by compositional shifts in which rules are published inside versus outside the window. Even during the 2017 surge, when agencies pushed more rules into the CRA window (the opposite of strategic avoidance), the rules that entered were shorter.

5.3 Robustness

Table 4: Robustness: Bandwidth Sensitivity, Placebo Cutoffs, and Transition Exclusions

<i>Panel A: Bandwidth Sensitivity — Diff-in-Disc, Significant Rule</i>				
Bandwidth (days)	Estimate	SE	p-value	N
±60	-0.1197**	0.0577	0.0380	8,957
±90	-0.0897*	0.0460	0.0511	13,549
±120	-0.0515	0.0393	0.1904	17,673
±150	-0.0525	0.0350	0.1337	21,781
±180	-0.0510	0.0320	0.1111	25,870
±240	-0.0493*	0.0277	0.0747	34,828
±365	0.0398*	0.0225	0.0772	51,854
<i>Panel B: Bandwidth Sensitivity — Diff-in-Disc, Page Length</i>				
Bandwidth (days)	Estimate	SE	p-value	N
±60	-11.83***	2.05	0.0000	8,957
±90	-9.32***	1.83	0.0000	13,549
±120	-6.87***	1.49	0.0000	17,673
±150	-7.60***	1.31	0.0000	21,781
±180	-6.84***	1.33	0.0000	25,870
±240	-5.93***	1.05	0.0000	34,828
±365	-2.03**	0.86	0.0187	51,854
<i>Panel C: Placebo Cutoffs (RDD on Significant, Cross-Party Only)</i>				
Cutoff Offset (days)	Estimate	Robust SE	p-value	N (eff.)
-90	-0.0103	0.0325	0.6352	3,999
-60	0.0440*	0.0289	0.0873	4,806
-30	0.0173	0.0308	0.4591	4,351
0 (true cutoff)	-0.0244	0.0334	0.6135	4,003
+30	0.0214	0.0257	0.2241	4,387
+60	0.0323	0.0341	0.3623	3,597
+90	0.0019	0.0386	0.9122	3,029
<i>Panel D: Sensitivity to Excluding Transitions</i>				
Specification	Estimate	SE	p-value	N
Excl. 2017	-10.54***	1.98	0.0000	10,314
Excl. 2025	-4.86***	1.74	0.0052	10,633
Excl. 2017+2025	-4.79**	1.95	0.0142	8,886

Notes: Panels A–B report the diff-in-disc interaction estimate ($\hat{\beta}_3$: CRA-vulnerable \times cross-party) for varying bandwidths. HC1 robust standard errors. Panel C reports local linear RDD estimates at placebo cutoffs, cross-

Table 4 presents four robustness exercises. Panel A shows that the diff-in-disc estimate on the significant-rule indicator is negative across all bandwidths, reaching conventional significance at ± 60 days ($\hat{\beta}_3 = -0.120$, $p = 0.038$) and marginal significance at ± 90 days ($p = 0.051$). The point estimate is relatively stable between -0.05 and -0.12 , suggesting that while the effect on significance status is imprecisely estimated, the direction is consistent with the page-length finding.

Panel B shows the headline page-length result is robust across all bandwidth choices. The diff-in-disc estimate ranges from -11.8 pages at ± 60 days to -2.0 pages at ± 365 days, with all estimates significant at $p < 0.02$. The attenuation at wider bandwidths is expected: observations far from the cutoff contribute less information about the discontinuity. Within the policy-relevant range of ± 60 to ± 180 days, the estimates are consistently large (-6.8 to -11.8 pages). The main result on page length is also robust to polynomial order: local linear ($p < 0.001$), local quadratic ($p < 0.001$), and local cubic ($p = 0.002$) specifications all yield large, negative, and significant estimates, following the recommendation of [Gelman and Imbens \(2019\)](#).

Panel C confirms that no placebo cutoff produces a significant discontinuity in the significant-rule indicator, with p -values ranging from 0.087 to 0.912.

Panel D addresses the density manipulation concern. The 2017 and 2025 transitions exhibit significant density jumps at the cutoff (Table 3), raising the possibility that compositional sorting drives the page-length result. Re-estimating the diff-in-disc excluding 2017, excluding 2025, and excluding both yields estimates of -10.5 , -4.9 , and -4.8 pages, all significant at $p < 0.02$. The result survives even when restricted to the four transitions (2001, 2009, 2021, and same-party years) with no evidence of density manipulation.

5.4 Descriptive Context

Table 5: Midnight Rulemaking: Rule Volume by Presidential Transition

Transition	Type	Total Rules	In CRA Window	% In CRA Window	Signif. Total	Signif. in CRA Window	Avg. Pages
2001	Cross	8,947	4,332	48.4	587	347	5.9
2005	Same	8,313	3,951	47.5	626	312	6.2
2009	Cross	7,362	3,661	49.7	695	458	7.5
2013	Same	7,505	3,600	48.0	739	295	7.7
2017	Cross	7,215	3,699	51.3	844	480	9.5
2021	Cross	6,511	3,342	51.3	786	490	9.5
2025	Cross	6,001	2,913	48.5	597	239	12.3

Notes: Each row shows one presidential transition. “Total Rules” counts all final rules published in the Federal Register within ± 365 days of the CRA lookback cutoff. “In CRA Window” counts rules published after the lookback date. “Significant” denotes rules classified as significant under E.O. 12866. Cross-party transitions involve a change in the party of the president; same-party transitions do not. The CRA lookback window spans approximately the final 60 Senate session days of each Congress.

Table 5 provides descriptive context for each transition. The 2017 transition stands out: 844 significant rules were published within the analysis window, with 488 in the CRA-vulnerable period. The 2025 transition, still ongoing at the time of data collection, shows a lower total but a similar fraction of rules inside the window.

Comparing cross-party and same-party transitions, the total volume of rulemaking is similar (7,000–9,000 rules within ± 365 days of the cutoff). The share of rules inside the CRA window is approximately 50 percent in both cases, reflecting the symmetric construction of the analysis window. The key difference is not in volume but in composition — a pattern fully consistent with the main result.

6. Discussion

The central finding of this paper is that the Congressional Review Act’s lookback window alters the *complexity* of federal regulations without changing their *volume*. During cross-party presidential transitions, rules published inside the CRA-vulnerable window are approximately 10 pages shorter in the Federal Register — a large effect that is absent during same-party

transitions and at placebo cutoffs.

This quality–quantity tradeoff has several interpretations. The most direct is that agencies facing CRA risk reduce the analytical depth of their rules. Federal Register page count reflects not only the regulatory text itself but also the preamble’s discussion of public comments, regulatory impact analysis, and alternatives considered. Shorter rules may therefore indicate truncated deliberation — agencies racing to finalize before a hostile Congress takes office, and cutting corners in the process.

An alternative interpretation is strategic: agencies may deliberately produce simpler, less consequential rules inside the CRA window, knowing they face higher reversal risk. Under this account, the page-length drop reflects self-selection in which regulations agencies choose to finalize during the vulnerable period, rather than a reduction in quality for a given regulation. The null on CFR parts ($p = 0.679$) is consistent with this view — it is the depth of each regulatory action, not its scope, that changes.

A third possibility is that the CRA’s deterrent effect extends beyond the 20 rules Congress has formally nullified. If agencies internalize the risk that CRA-vulnerable rules will be overturned, they may invest less in the regulatory process for rules finalized in the window — a form of underinvestment driven by anticipated expropriation, analogous to the underinvestment result in the holdup literature. This interpretation suggests the CRA’s true cost is not the rules it kills but the rules it weakens.

Several limitations deserve mention. First, this analysis measures how CRA vulnerability shapes rule *characteristics* (length, complexity) rather than rule *survival* (whether rules are subsequently rescinded or nullified). The survival question — whether shorter CRA-window rules are more likely to be overturned — requires linking individual rules to future regulatory actions across Federal Register citations, a natural extension for future work. The present contribution is to establish that the CRA’s shadow falls on the production process itself.

Second, page length is an imperfect proxy for regulatory quality. Short rules are not necessarily worse, and some complex regulations may be efficiently drafted. However, Federal Register page count captures both the regulatory text and the preamble’s discussion of comments, economic analysis, and alternatives considered (Coglianese, 2004). The discontinuity is specific to cross-party transitions and absent at placebos, making it unlikely that the finding reflects routine variation in drafting conventions.

Third, the lookback dates are approximate, computed from CRS reports rather than exact session calendars. Any measurement error in the cutoff date would attenuate the RDD estimates, making the significant findings conservative. Fourth, the sample contains only two same-party transitions (2005 and 2013), limiting the precision of the control-group discontinuity. The diff-in-disc design helps by leveraging within-transition variation, but

additional same-party transitions would strengthen the comparison.

7. Conclusion

The CRA lookback window — a procedural feature of the congressional calendar — casts a measurable shadow over the federal regulatory process. Agencies facing CRA risk during cross-party transitions produce shorter, simpler regulations, without reducing total output. The Act’s influence on the regulatory stock operates not through the blunt instrument of nullification (which has been used only 20 times in three decades) but through its ex ante effect on agency behavior. Whether this deterrent improves or degrades the quality of American regulation depends on whether the rules that would have been longer are the rules that should have been written with more care — or the rules that should never have been written at all.

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

- Beermann, Jack M.**, “Midnight Rules: A Reform Agenda,” *Administrative Law Review*, 2012, *64* (4), 795–840.
- Calonico, Sebastian, Matias D. Cattaneo, and Rocío Titiunik**, “Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs,” *Econometrica*, 2014, *82* (6), 2295–2326.
- Carey, Maeve P. and Christopher M. Dolan**, “The Congressional Review Act: Determining Which “Rules” Must Be Submitted to Congress,” *Congressional Research Service Report*, 2021. R43992.
- Cattaneo, Matias D., Michael Jansson, and Xinwei Ma**, “Simple Local Polynomial Density Estimators,” *Journal of the American Statistical Association*, 2020, *115* (531), 1449–1455.
- Coglianesi, Cary**, “E-Rulemaking: Information Technology and the Regulatory Process,” *Administrative Law Review*, 2004, *56* (2), 353–402.
- Gelman, Andrew and Guido Imbens**, “Why High-Order Polynomials Should Not Be Used in Regression Discontinuity Designs,” *Journal of Business & Economic Statistics*, 2019, *37* (3), 447–456.
- Grembi, Veronica, Tommaso Nannicini, and Ugo Troiano**, “Do Fiscal Rules Matter?,” *American Economic Journal: Applied Economics*, 2016, *8* (3), 1–30.
- O’Connell, Anne Joseph**, “Political Cycles of Rulemaking: An Empirical Portrait of the Modern Administrative State,” *Virginia Law Review*, 2008, *94* (4), 889–986.
- Potter, Rachel A.**, “Slow-Rolling, Fast-Tracking, and the Pace of Bureaucratic Decisions in Rulemaking,” *Journal of Politics*, 2017, *79* (3), 841–855.

A. Data Appendix

A.1 Federal Register API

All data are drawn from the Federal Register API,² which provides structured metadata for every document published in the Federal Register since 1994. I query all documents of type “RULE” (final rules) for publication years 1999–2025, retrieving publication date, title, agencies, significance classification, page length, CFR references, regulation identification numbers, action type, and subtype. The API requires no authentication and returns JSON responses paginated at 1,000 documents per page.

A.2 CRA Lookback Dates

I compute approximate lookback dates from Congressional Research Service reports and Senate legislative calendars. The lookback window spans approximately 60 Senate legislative days before the end of each Congress (January 3 of odd-numbered years). Because legislative days exclude weekends, recesses, and adjournments, the lookback start date falls in late May or early June of the preceding even-numbered year. The dates used are: May 29, 2000 (106th Congress); May 19, 2004 (108th); May 15, 2008 (110th); May 15, 2012 (112th); May 30, 2016 (114th); May 19, 2020 (116th); May 22, 2024 (118th).

A.3 Sample Construction

The raw dataset contains 98,157 final rules. I assign each rule to the nearest CRA lookback cutoff within ± 365 calendar days, yielding 51,854 rule-transition observations. Rules equidistant from two cutoffs are assigned to the closest. The running variable is defined as the number of calendar days from the rule’s publication date to the applicable lookback cutoff, with positive values indicating publication inside the CRA-vulnerable window.

B. Identification Appendix

B.1 Density Discontinuity Test

I implement the Cattaneo et al. (2020) density discontinuity test at the CRA lookback cutoff. The pooled test across cross-party transitions fails to reject the null of continuous density ($T = -1.42$, $p = 0.155$), indicating no systematic manipulation of rule timing. Individual

²<https://www.federalregister.gov/api/v1/documents.json>

transitions show heterogeneity, with a significant positive density jump in 2017 ($T = 3.29$, $p = 0.001$) and a negative jump in 2025 ($T = -2.41$, $p = 0.016$).

B.2 Covariate Balance

The number of CFR parts affected per rule is smooth through the cutoff ($\hat{\beta} = 0.019$, $p = 0.820$), supporting the assumption that rule scope does not change discontinuously at the CRA threshold.

C. Robustness Appendix

Bandwidth sensitivity, placebo cutoff, and transition exclusion results are reported in Table 4 in the main text. The diff-in-disc estimate on page length is robust to bandwidth choices from ± 60 to ± 365 days, to polynomial orders 1–3, and to excluding the 2017 and 2025 transitions that exhibit density manipulation at the cutoff.

D. Standardized Effect Sizes

Table 6: Standardized Effect Sizes for Main Outcomes

Outcome	Spec.	$\hat{\beta}$	SD(X)	SD(Y)	SDE	SE	Class.
Significant rule (0/1)	Loc. lin.	-0.0244	—	0.4085	-0.0598	0.0817	Moderate negative
Page length	Loc. lin.	-7.6863	—	30.3678	-0.2531	0.0686	Large negative

Notes: Standardized effect sizes ($SDE = \hat{\beta}/SD(Y)$) for cross-study comparison. Spec.: local linear RDD, CCT bandwidth, triangular kernel. Data: Federal Register API, cross-party transitions (2001–2025), $N = 36,036$. Treatment: binary (inside vs. outside CRA window). Class. labels refer to SDE magnitude, not statistical significance. “Null”: $|SDE| < 0.005$.