PIER

PUEY UNGPHAKORN INSTITUTE FOR ECONOMIC RESEARCH

Efficiency at a Cost: How a Fiscal Rule on Disbursement Timelines Shifted Public Investment Toward Repairs

by

Athiphat Muthitacharoen, Phatarakorn Thongsathit, and Tapanat Paiboonsin

January 2026

Discussion Paper

No. 246

Efficiency at a Cost: How a Fiscal Rule on Disbursement Timelines

Shifted Public Investment Toward Repairs[*]

Athiphat Muthitacharoen   Phatarakorn Thongsathit

Tapanat Paiboonsin

January 2026

ABSTRACT

This paper examines how an efficiency-oriented fiscal rule influences the composition of public investment. We study Thailand's 2021 fiscal rule reform limiting the disbursement period for investment budgets to one year after initial approval, replacing a previously unrestricted timeline. We apply seeded LDA model to classify over 360,000 Thai-language government projects and to estimate the probability that each project reflects repair activity. We then implement a difference-in-differences framework. Our identification exploits variation in departments' pre-policy reliance on construction projects, which proxy exposure to tighter disbursement constraints. We find that departments more exposed to the reform significantly increased their reliance on repair-oriented projects, which are typically simpler and faster to execute. This adjustment emerges gradually, is more pronounced among larger departments, and coincides with higher post-reform disbursement rates, indicating that the policy achieved its primary objective of accelerating budget execution. At the same time, the results suggest a trade-off: efficiency-oriented fiscal rules can alter agencies' incentives in ways that shift investment portfolios toward projects that are easier to implement under tighter timelines. Overall, the study highlights how fiscal rule design can shape not only the pace but also the composition of public investment.

Keywords: Fiscal rules; Public investment composition; Budget execution; Administrative behavior; Impact evaluation

Statements and Declarations: The authors declare no conflict of interest. Muthitacharoen received financial support from the Chulalongkorn Economics Research Centre (CERC).
Ethics and Consent to Participate Declarations: This research was exempted from institutional ethics review because it does not involve human participants or the collection of identifiable personal data.
Acknowledgments and disclaimers: We thank comments and suggestions from participants at the PIER Research Exchange. The opinions expressed in this study are those of the authors and should not be attributed to Bank of Thailand.

[*] Muthitacharoen: Faculty of Economics, Chulalongkorn University, Thailand (athiphat.m@chula.ac.th); Thongsathit: Bank of Thailand (phatarat@bot.or.th); Paiboonsin: Bank of Thailand (tapanatp@bot.or.th).

INTRODUCTION

Public investment is central to long-term economic development, particularly in economies with infrastructure gaps and limited fiscal space. To safeguard public capital formation and strengthen budget discipline, governments increasingly rely on fiscal rules that constrain how public resources are allocated and executed (Dahan and Strawczynski, 2013; Ardanaz et al., 2021). While such rules are often designed to improve efficiency and credibility, they can also alter the incentives faced by implementing agencies, potentially reshaping not only the level but also the composition of public investment (Burret and Feld, 2018).

A large body of research has examined how fiscal rules influence the level of public investment (see, for example, Ardanaz et al., 2021; Delgado-Téllez et al., 2022; Vinturis, 2022). Much less attention, however, has been devoted to how these rules influence what types of investment projects governments choose to undertake. This distinction matters because different forms of capital spending, such as new construction versus repair and maintenance, differ substantially in complexity, implementation timelines, and long-run growth implications (Dabla-Norris et al., 2012; Rajaram et al., 2014). Understanding how fiscal and administrative constraints shape investment composition is therefore critical for evaluating the broader consequences of fiscal rule design.

This paper studies a fiscal rule reform in Thailand that provides a useful setting to examine this issue. Effective in 2021, Thailand introduced a binding disbursement deadline requiring public investment budgets to be fully disbursed within one year of approval. Under the previous system, investment funds could remain available across multiple fiscal years. The reform was intended to accelerate budget execution and reduce implementation delays.

Beyond altering the timing of budget execution, the reform changed the risk environment faced by implementing agencies. Under the new rule, failure to disburse funds within the deadline results in the forfeiture of remaining budgets and resubmission in a subsequent budget cycle. For departments managing large or technically complex construction projects, this constraint substantially increases the risk of failing to complete projects within the prescribed timeframe, thereby creating stronger incentives to adjust project selection and implementation strategies.

As a result, departments face a clear organizational trade-off. New construction projects may offer higher long-term returns but carry greater execution risk under a binding disbursement deadline. Repair projects, by contrast, are typically smaller in scale and faster to implement. Shifting investment portfolios toward repair-oriented projects therefore becomes a rational response to minimize the risk of budget forfeiture. This behavioral response does not require changes in formal priorities or political objectives; it arises mechanically from the interaction between project complexity and a time-based fiscal constraint. Accordingly, departments with greater pre-policy exposure to construction-intensive investments should exhibit larger post-reform shifts toward repair-oriented projects.

This paper examines whether such a response occurred. While the analysis focuses on Thailand, the underlying question is broadly applicable. Many governments impose time-based fiscal or administrative constraints to improve budget execution, particularly in developing and emerging economies. If these constraints systematically favor simpler, short-term projects, they may unintentionally discourage more complex infrastructure investments with higher long-run returns.

Addressing this question poses two empirical challenges. First, public investment data, especially in developing-country settings, rarely contain structured classifications distinguishing repair and maintenance from new construction. Second, even when such distinctions are implicit in project titles, extracting them at scale requires systematic text-based methods.

We address both challenges using a combination of semi-supervised text analysis and quasi-experimental evaluation. First, we apply a seeded Latent Dirichlet Allocation (LDA) model to classify more than 360,000 Thai-language public investment project descriptions and estimate the probability that each project reflects repair activity, guided by expert-defined seed terms. This approach yields a continuous, interpretable measure of repair intensity that can be aggregated and analyzed econometrically. We then exploit the 2021 disbursement rule reform as a quasi-experiment, implementing a difference-in-differences (DID) design that leverages variation in departments' pre-policy reliance on construction projects as a measure of exposure to the new constraint.

Our findings show that the disbursement rule change significantly altered the composition of public investment. Departments that were more reliant on construction projects prior to the reform shifted toward repair projects in the post-reform period. The magnitude is meaningful: in our baseline specification, a one–standard-deviation increase in pre-policy construction exposure raises the likelihood that a project is classified as repair by 4.4 percentage points (roughly 8% of the pre-policy mean). This pattern is robust across alternative exposure measures and outcome definitions, including simple keyword-based classifications. The adjustment emerged gradually and was more pronounced among larger departments, consistent with differential exposure to execution risk.

We also examine whether the reform achieved its core objective of improving disbursement performance. Departments with higher pre-policy construction exposure exhibit significantly faster disbursement rates after the rule change. A one–standard-deviation increase in exposure raises the disbursement rate by about 5% of the pre-policy mean.

Taken together, the findings point to a policy trade-off. While time-based fiscal rules can improve disbursement efficiency, they may also alter agencies' incentives in ways that favor simpler, faster-executing projects, such as repairs, over more complex investments. Importantly, this study does not assess whether repair investment is inherently inferior to new construction; rather, it highlights an underappreciated dimension of fiscal rule design—its influence on the strategic behavior of implementing agencies and, in turn, on the composition of public investment.

This paper makes three key contributions. First, we provide empirical evidence on a previously underexplored mechanism: how efficiency-oriented fiscal rules influence the composition, not just the level of public investment. Second, we demonstrate how semi-supervised text analysis can uncover latent policy responses in administrative data. Third, we offer a replicable framework for monitoring investment quality in settings where project information exists only as unstructured text.

The remainder of the paper is organized as follows. Section 2 reviews the literature on fiscal and budget rules. Section 3 describes the disbursement time-limit reform, summarizes public investment project data, and illustrates how we identify repair projects using a seeded LDA approach. Section 4 presents the main empirical findings based on a DID framework. Section 5 concludes.

RELATED STUDIES

This study contributes to two intersecting strands of research: 1) fiscal rules, regulatory constraints and public investment, and 2) the application of text analysis for policy-relevant classification.

The first strand of literature examines how fiscal rules and regulatory constraints influence public investment. Fiscal rules, by design, constrain government borrowing and spending, thereby shaping how resources are allocated across investment projects. Most studies focus on the aggregate level of investment, assessing whether fiscal rules promote or crowd out public capital formation. For instance, Vinturis (2023) finds that fiscal rule adoption significantly reduces government consumption while leaving investment largely unaffected.

The design of fiscal rules emerges as a key determinant of their impact. Flexible rules that include investment-friendly provisions—such as escape clauses or capital expenditure exemptions—tend to support sustained or even higher levels of investment. In contrast, rigid constraints can limit governments' ability to finance large or long-term projects (Bléssé, Dorn & Lay, 2023; Ardanaz et al., 2021). Similarly, Delgado-Téllez et al. (2022) show that overly stringent rules may unintentionally depress public investment by incentivizing governments to prioritize compliance with fiscal targets over socially optimal spending, often resulting in cuts to infrastructure programs.

Beyond total investment levels, a smaller but growing body of research emphasizes that fiscal rules can also shape the composition of government spending. Dahan and Strawczynski (2013) found a negative effect of fiscal rules on the share of social transfers to government consumption. Burret and Feld (2018) highlight that strict compliance requirements may lead to short-term budgetary adjustments at the expense of long-term investment goals, particularly when rules are narrowly defined or inflexible. Relatedly, Rajaram et al. (2014) argue that bureaucracies operating under tight budgetary controls tend to favor low-complexity, easy-to-execute projects to avoid breaching financial constraints. Further, Dabla-Norris et al. (2012) evaluates institutional environment underlying public investment management. It shows that weak project appraisal systems exacerbate misallocation, steering funds toward less productive or politically expedient investments—a challenge especially relevant in developing fiscal systems.

The second strand relates to employing text analysis techniques, particularly topic modeling, in economic research. Latent Dirichlet Allocation (LDA), introduced by Blei et al. (2003), is a foundational method for extracting latent themes from unstructured text. Subsequent innovations like seeded LDA allow experts to guide the model via curated keywords, enhancing applicability for policy classification tasks (Jagarlamudi et al., 2012; Watanabe & Baturo, 2024).

In economic policy research, topic modeling has assisted in interpreting central bank communications and forecasting monetary policy (Hansen et al., 2018; Luangaram & Wongwachara, 2017). In the fiscal policy literature, Latifi et al. (2024) use LDA to classify parliamentary debate tones, linking fiscal sentiment shifts to spending behavior. Topic modeling has also been employed in sustainability policy research; for example, Cruciani and Santagiustina (2023) use LDA to track ESG-related themes in corporate reports.

Despite advances in both strands, these literatures have rarely been brought together to study how fiscal rules shape the content of public investment at the project level. While prior studies have examined how fiscal rules affect overall investment levels, much less is known about how regulatory constraints, such as disbursement period limits, influence project design and selection. Existing work that touches on project composition tends to be policy-oriented rather than empirical, and no study to date has applied text-based classification methods to analyze how binding fiscal rules shape the strategic composition and quality of public investment.

This study helps fill that gap by examining how Thailand's disbursement time-limit rule affected the balance between repair and new construction investments. Specifically, it applies a seeded LDA approach to classify Thailand's public investment projects as either new construction or repair. In doing so, it directly engages with the concern raised by Rajaram et al. (2014) and Burret and Feld (2018) that governments facing stringent fiscal or procedural constraints may gravitate toward low-complexity projects. In addition, the study introduces a replicable analytical framework that can be applied in other developing-country contexts where project-level information is available primarily in unstructured text form.

INSTITUTIONAL SETTING, DATA AND MEASUREMENT

This section explains how we exploit a fiscal rule change as a quasi-experimental setting and describes the investment project data used in the study. It then outlines our approach to identifying repair-oriented investment using seeded LDA.

Quasi-experimental setting: Disbursement time-limit rule

We exploit the introduction of a binding disbursement time limit on public investment spending in Thailand as a quasi-experimental setting. The reform first applied to the fiscal year 2021 budget cycle. Under the new rule, government investment budgets must be fully disbursed by the end of the fiscal year following the year of initial budget approval (T+1). Any undisbursed balance is forfeited and projects must be resubmitted for approval in a subsequent Appropriations Act. Prior to the reform, investment budgets faced no formal disbursement deadline and could remain available across multiple fiscal years.

While the reform applies to all government investment projects, but its effects are unlikely to be uniform. New construction projects typically involve larger budgets, more complex procurement processes, and longer implementation timelines, making them less likely to be completed within a fixed disbursement window. Repair and maintenance projects, by contrast, are generally smaller in scale and faster to execute.

As a result, departments with greater pre-reform reliance on construction-intensive investment are more exposed to the new constraint. This differential exposure generates variation in the effective treatment intensity of the reform, which we exploit empirically to test whether departments facing tighter constraints adjust their investment portfolios toward repair-oriented projects following the rule change.

Investment Project Data

We use project-level public investment data for Thailand covering fiscal years 2018–2023, drawn from the Government Fiscal Management Information System (GFMIS). GFMIS is the administrative platform used by government agencies for budget planning, financial reporting, and fiscal oversight. The dataset provides detailed information on individual investment projects, including project titles, approved budgets, responsible agencies, and project locations.

Each project title, written in Thai, is relatively detailed and provides informative descriptions of the nature and scope of government investment. Table 1 presents examples of project titles to illustrate the level of detail in the data. Our analysis focuses on construction-type investment projects recorded in the GFMIS database, ensuring that we examine projects directly subject to the disbursement rule.

Table 1. Examples of project titles.

Original Project Title (in Thai)	English Translation
ปรับปรุงประสิทธิภาพสถานีวิทยุกระจายเสียงแห่งประเทศไทย ตำบลปากน้ำ อำเภอเมืองกระบี่ จังหวัดกระบี่ ระบบเอฟ.เอ็ม ขนาดกำลังส่ง 1 กิโลวัตต์	Improving the efficiency of the National Broadcasting Station of Thailand, Pak Nam Subdistrict, Mueang Krabi District, Krabi Province — FM system with 1-kilowatt transmission power
ซ่อมแซมหลังคาอาคารโรงเก็บแก๊ส สอต.ประจวบคีรีขันธ์	Repair of the roof of the gas storage building, Provincial Police Training Center, Prachuap Khiri Khan
ปรับปรุงซ่อมแซมอาคารเรียนอาคารประกอบและสิ่งก่อสร้างอื่นโรงเรียนบ้านโคกเสี่ยว ตำบลคูขาดอำเภอคง จังหวัดนครราชสีมา	Renovation and repair of classroom buildings, auxiliary structures, and other facilities at Ban Khok Siao School, Khukhat Subdistrict, Khong District, Nakhon Ratchasima Province
เขื่อนป้องกันตลิ่งริมแม่น้ำเจ้าพระยา หมู่ที่ 6 เทศบาลตำบลพยุหะ อำเภอพยุหะคีรี จังหวัดนครสวรรค์ ความยาว 517 เมตร	Construction of a riverbank protection embankment along the Chao Phraya River, Village No. 6, Phayuha Subdistrict Municipality, Phayuha Khiri District, Nakhon Sawan Province — 517 meters long
สร้างอาคารศูนย์ฝึกอบรมเทคโนโลยีสารสนเทศ กองกรรมวิธีข้อมูล ศลภ.คปอ.แขวงสายไหม เขตสายไหม กรุงเทพมหานคร	Construction of an Information Technology Training Center Building, Data Processing Division, Directorate of Logistics, Saimai Subdistrict, Saimai District, Bangkok

Table 2 provides summary statistics for budget of public investment projects. Overall, there are roughly 360,000 projects, with an average budget of 6.3 million baht and a median budget of 0.7 million baht.

Table 2. Summary statistics of public investment budget (in million baht).

Budget (in million baht)	N	Mean	Median	SD
All years	362,979	6.330	0.715	37.785
2018	66,609	5.726	0.737	35.947
2019	60,971	6.348	0.610	50.564
2020	63,854	6.197	0.615	36.438
2021	60,786	6.637	0.677	35.651
2022	52,075	7.830	0.868	38.574
2023	66,609	5.726	0.737	35.947

Identifying Repair Investment Using Seeded LDA

A central challenge in studying the composition of public investment is the lack of structured classifications identifying repair projects. In Thailand, as in many developing-country settings, administrative investment data consist primarily of project titles/descriptions written in a local language, without standardized indicators of project type. To address this challenge, we combine semi-supervised text analysis with administrative budget data to construct a continuous, interpretable measure of repair-oriented investment at the project level.

We apply a seeded Latent Dirichlet Allocation (LDA) model to classify investment projects based on their textual descriptions. Seeded LDA extends standard topic modeling by incorporating expert-defined seed words that anchor topics to substantively meaningful categories, while allowing the remaining word distributions to be learned from the data (Watanabe and Baturo, 2024). This approach is particularly useful in settings where prior domain knowledge can guide classification, but where project descriptions are too heterogeneous or incomplete for simple rule-based methods.

In our application, we focus on a single topic of interest, repair and maintenance investment, and guide the classification using a set of Thai seed words that are known ex ante to characterize such activities. These seed terms capture common administrative language associated with repair, rehabilitation, maintenance, and surface improvement in public investment projects. Seeded LDA incorporates this prior information by assigning greater weight to the predefined seed words, anchoring the repair topic in substantively meaningful terminology while allowing the remaining word–topic associations to be learned from patterns of word co-occurrence in the data. As a result, the model is able to classify projects that lack explicit repair keywords but share similar contextual language, producing a continuous probability measure that reflects the degree to which each project is repair-oriented.

This approach is well suited to the structure of Thai project titles. Compared with simple keyword matching, seeded LDA exploits word co-occurrence patterns to classify projects even when explicit repair-related terms are absent. At the same time, unlike large language models, seeded LDA offers transparent topic definitions and reproducible outputs, making it well-suited for policy-oriented analysis using administrative data.

The model yields, for each project, a probability that its description corresponds to a repair-oriented investment. We aggregate these probabilities to the department–year level to construct a budget-weighted measure of repair intensity, which serves as the primary outcome variable in the Difference-in-Differences analysis.

We assess the validity of the classification using several complementary checks. Table 3 reports the top words associated with each estimated topic, showing that repair-related terms (e.g., renovation, maintenance, improvement) load strongly on the repair topic, while construction-related terms (e.g., construction, building, expansion) dominate the background topic.

Table 3. Top 10 words in the Repair and Background topics.

Repair topic	**Background topic**
Improve	Construct
Maintenance	Building
Repair	School
Water	Water
Fix	Office
Surface	Temple
Upkeep	Dam
Dredge	Highway
Road	Road
Canal	Design

Note: The background topic is not seeded but emerges as a residual topic.

Table 4 presents examples of projects with the highest and lowest predicted repair probabilities. High-probability projects clearly correspond to repair and maintenance activities, while low-probability projects reflect new construction, supporting the substantive interpretability of the classification.

Table 4. Examples of projects with highest and lowest repair probability.

Project name	Repair probability
Top 10 projects	
Repair of the Water Delivery and Maintenance Office No. 2, Lopburi Irrigation Project	0.983
Repair of Water Conveyance System and Canal Embankments, Lateral Canal 2 (Left), Main Canal No. 1	0.981
Repair of Sluice Gate Winches in the Water Delivery System (Diameter 0.30–1.25 m), 53 units	0.980
Repair of Water Pipeline System and Associated Structures, Huai Khan Reservoir (11 locations), Mukdahan Irrigation Project	0.980
Maintenance of CCTV System in the Chao Phraya River Basin, Water Situation Monitoring Division	0.980
Repair of Irrigation Canal 1R–Chainat–Pasak, 300 meters, Wat Khok Subdistrict, Manorom District	0.980
Repair of Stone Protection Structures Near Canal Banks, 10 locations, Pase Mas Subdistrict	0.980
Repair of Water Supply System for Ban Mae Omki School, Mae Wa Luang Subdistrict, Tha Song Yang District	0.979
Repair of Water Control Structure at Tha Sung Canal, Including Associated Facilities	0.979
Repair of Downstream Drainage System, Huai Na Nuea Reservoir, Nong Waeng Subdistrict, Lahan Sai District	0.979
Bottom 10 projects	
Construction of Asphalt Concrete Road, Route NMT 321–05, Ban Thamnop Pattana to Ban Phrangam	0.016
Improvement of Building Facilities, Thung Song Hong Subdistrict, Lak Si District, Bangkok	0.016
Consulting Services for Construction Supervision and Infrastructure System Procurement	0.017
Construction Supervision for Efficiency Improvement and Capacity Expansion (Mae Tuen Hydropower Project)	0.017
Improvement of Backup Energy System for Solar Power Production Facility	0.017
Construction of 40×40 m Tree Nursery, Na Nong Thum Subdistrict, Chum Phae District, Khon Kaen	0.017
Construction of 40×40 m Tree Nursery, Mai Na Phiang Subdistrict, Waeng Noi District, Khon Kaen	0.017
Modular 20×20 m Tree Nursery, Ban Sahakon Subdistrict, Mae On District, Chiang Mai	0.017
Earth-Retaining Wall and Fence Construction, Forest Resource Management Office 4, Nakhon Sawan	0.017
Installation of Water Supply System for Chiang Mai Seedling Nursery Center	0.017

Figure 1 plots the distribution of repair probabilities across projects. The distribution is bimodal, with mass concentrated near zero and one, indicating that most projects are clearly classified rather than ambiguously assigned. For descriptive purposes, we classify projects with predicted probabilities above 0.5 as repair-oriented, though all regression analyses use the continuous probability measure.

Figure 1. Density distribution of repair probability across all government investment projects (2023).

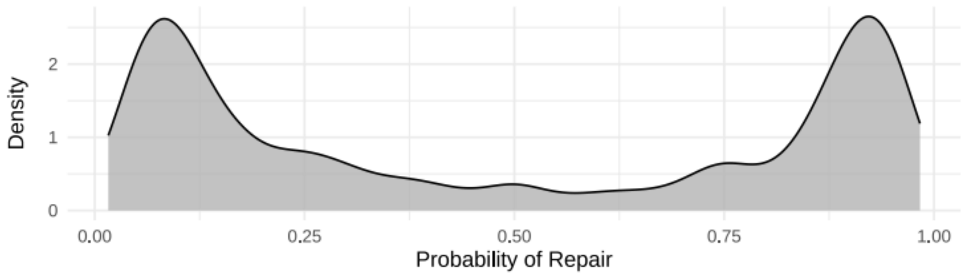

Note: This figure illustrates the density distribution of repair probabilities across government investment projects for fiscal year 2023.

Like all text-based classification methods, seeded LDA has limitations, including sensitivity to the choice of seed words and the absence of the deeper semantic representations available in large language models. To assess the robustness of our Difference-in-Differences results to these features, we conduct sensitivity analyses using alternative seed word lists and a simple keyword-based repair indicator. These checks are discussed in the next section.

An important policy question is the extent to which the government's construction budget is devoted to repair activities. Because our classification is probabilistic, this share cannot be observed directly. To provide an interpretable benchmark, we define repair investment as construction projects with an estimated repair probability greater than 0.5. This threshold-based classification is necessarily imperfect, as some projects may combine repair and new construction elements, and probabilistic scores near the cutoff may reflect classification uncertainty rather than clear project intent.

Figure 2 illustrates the resulting budget-weighted shares of construction investment classified as repair. Using this definition, we find that 46% of total construction investment in 2023 is categorized as repair. Moreover, the share of repair-oriented investment appears to increase over time (Figure 4), rising from 38% in the pre-policy period (2018–2020) to 47% in the post-policy period (2021–2023). These patterns motivate the difference-in-differences analysis that follows. Appendix B provides a detailed description of the seeded LDA methodology, including preprocessing steps, estimation details, and robustness to alternative seed definitions.

Figure 2. Scale of repair investment within Thailand's government construction budget.

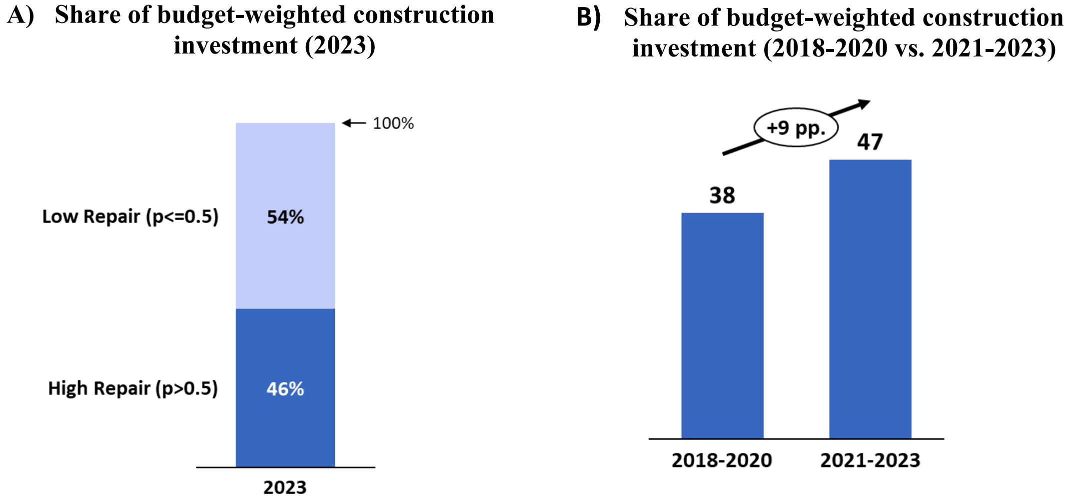

Notes: Panel A illustrates the share of budget-weighted construction investment (2023), while Panel B shows the share of budget-weighted construction investment for the 2018-2020 and the 2021-2023 periods. Repair investment is defined as construction projects with an estimated repair probability p>0.5. The repair probability p is a model-based measure derived from a seeded Latent Dirichlet Allocation (LDA) classification of project descriptions.

EMPIRICAL FINDINGS

This section presents our DID analysis and reports the main results on how the introduction of the disbursement time-limit reform affected the composition of public investment.

Empirical Strategy

We use a DID framework to estimate the impact of the reform on the likelihood that a department's investment portfolio is oriented toward repair projects. The analysis is conducted at the department-year level. Our main outcome variable is the budget-weighted probability that a department's investment projects are classified as repair-oriented.

To capture differential exposure to the policy, we use the pre-policy share of construction projects in total investment budget as a continuous treatment variable. This measure reflects each department's vulnerability to the new time constraint. Departments with higher construction shares are more likely to face larger adjustment pressures and, in turn, to shift toward faster-executing, repair-oriented projects in response to the reform.

Table 5 summarizes the descriptive statistics for the department-level dataset covering fiscal years 2018–2023, including key variables used in the DID analysis: the budget-weighted repair probability, the pre-policy construction share (our measure of policy exposure), and relevant control variables.

Table 5. Descriptive statistics.

Variables	N	Mean	Median	S.D.
Repair probability (Budget-weighted mean)	684	0.465	0.453	0.228
Disbursement rate (Budget-weighted mean)	682	0.633	0.642	0.231
Post	684	0.507	1.000	0.500
Exposure	684	0.226	0.083	0.276
Project numbers	684	446.6	32.0	1,551.2
Total_budget (million baht)	684	2,865.8	139.5	12,216.3
Repair probability (Simple mean)	684	0.472	0.455	0.214
Repair keywords-Having any repair project (Binary)	684	0.937	1.000	0.243
Repair keywords-Share of repair projects	684	0.432	0.430	0.281

Note: The table describes summary statistics of data used in the analyses.

We define the pre-policy period as fiscal years 2018–2020, and the post-policy period as 2021–2023. The estimation equation is specified as follows:

$$y_{it} = \alpha_0 + \alpha_1 Post_t + \alpha_2 Post_t \cdot Exposure_i + \alpha_3' X_{it} + DepartmentFE + YearFE + \varepsilon_{it}, \quad (1)$$

where y_{it} denotes outcome variable, $Post_t$ denotes a dummy variable that equals one for the post-policy years and equals zero otherwise, $Exposure_i$ denotes the department's pre-policy average share of construction projects in total investment budget (measured over 2018–2020), and control variables X_{it} denotes the set of control variables including total number of projects (log) and total budget (log). Department and year fixed effects are included. We use robust standard errors clustered at the department level.

Under the identification assumption that unobserved determinants of the outcome variable (ε_{it}) do not vary differentially across high- and low-exposure group around the time of the policy change, the coefficient α_2 captures the causal effect of the disbursement time-limit-reform policy change on the outcome.

We further examine heterogeneous treatment effects by department size and ministry type. In addition, we assess the impact of the reform on disbursement rates, a direct policy target intended to improve the efficiency of budget execution.

Finally, we conduct a series of sensitivity analyses. We first consider alternative outcome measures, including the unweighted mean repair probability and indicators based on counts of projects containing repair-related keywords. We also assess sensitivity to seed-word selection by estimating the seeded LDA model using an alternative set of keywords. In addition, we test the robustness of our exposure measure by applying a log transformation to the exposure variable.

To further support our identification strategy, we examine whether departments with higher pre-policy construction exposure exhibited differential trends in repair intensity prior to the reform. Specifically, we conduct a placebo test using only the pre-policy period (2018–2020), interacting construction exposure with a time trend. We find no evidence of such differential pre-trends: the interaction terms are small in magnitude and statistically insignificant throughout the pre-reform period (Appendix Table A1). This result indicates that departments with higher construction exposure were not already shifting toward repair-oriented projects before the policy change.

We next present an event-study analysis around the reform (Figure 3). For visual clarity and statistical precision, we bin all years prior to the policy into a single pre-policy period (Year ≤ -1) and use fiscal year 2020 (Year 0)—the last full year before implementation—as the reference category. The estimated coefficient for the pre-policy period is small and statistically insignificant, providing no evidence of differential pre-trends. The coefficient for the first post-policy year is also insignificant, while estimates for Years 2 and 3 are positive and statistically significant, indicating a gradual adjustment in project composition over time.

Taken together, the placebo test and event-study evidence strengthen the causal interpretation of our difference-in-differences estimates. The placebo test addresses concerns about selection on pre-policy construction exposure by showing no differential trends prior to the reform, while the event-study analysis supports the parallel trends assumption and documents the dynamic adjustment following policy implementation. The results suggest that departments adjusted their investment planning gradually rather than immediately, consistent with procurement and organizational constraints.

Figure 3. Event-study estimate of the effect of the disbursement policy on budget-weighted repair project probability.

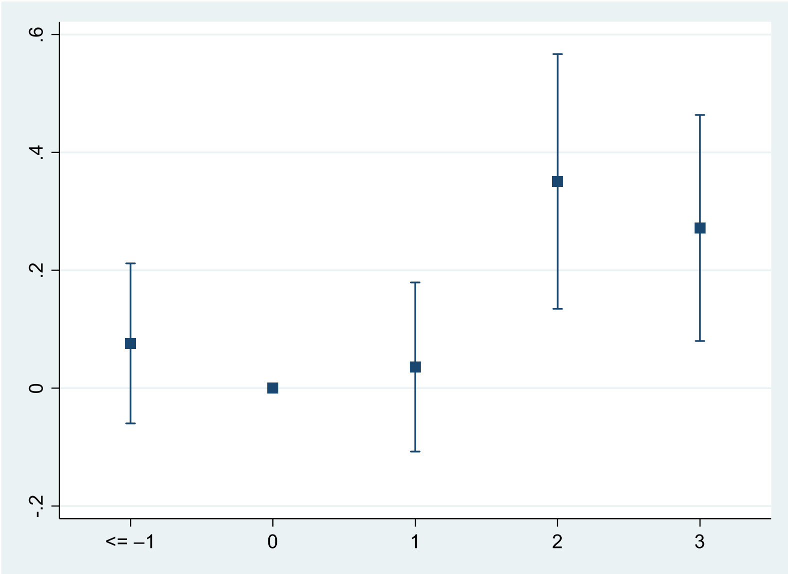

Notes: The figure plots event-study estimates of the effect of the disbursement policy on the budget-weighted probability that a department's investment project is classified as a repair project. The omitted reference category is Year 0 (fiscal year 2020). Error bars represent 95% confidence intervals.

Difference-in-Differences Findings

We begin by examining the overall effect of the disbursement time-limit reform using the budget-weighted mean probability that a department's investment projects are classified as repair-oriented as the outcome variable. Table 6 reports the results from the DID specification (equation 1), which exploits variation in departments' pre-policy exposure to construction projects.

Table 6. Effects of the disbursement policy on repair-project probability and disbursement rate.

	Repair-project probability				Disbursement rate
	(1)	(2)	(3)	(4) Baseline	(5)
$Post_t$	-0.191***	-0.358***	-0.356***	-0.359***	-0.058**
	(0.023)	(0.026)	(0.026)	(0.026)	(0.028)
$Post_t \cdot Exposure_i$	0.163***	0.166***	0.162***	0.158**	0.111**
	(0.062)	(0.062)	(0.061)	(0.064)	(0.050)
$Project\ numbers_{it}$				-0.009	-0.006
				(0.018)	(0.016)
$Total\ Budget_{it}$				0.001	0.023
				(0.012)	(0.017)
$Exposure_i$	-0.116**	-0.060			
	(0.052)	(0.057)			
$Constant$	0.570***	0.643***	0.659***	0.689***	0.581***
	(0.017)	(0.044)	(0.022)	(0.064)	(0.090)
Observations	684	684	684	684	682
Number of Departments	142	142	142	142	142
R-squared	0.125	0.293	0.317	0.318	0.061
Dept. FE	No	No	Yes	Yes	Yes
Year FE	No	Yes	Yes	Yes	Yes
Ministry FE	No	Yes	n/a	n/a	n/a
Control	No	No	No	Yes	Yes

Notes: This table reports the estimated effects of the disbursement time-limit policy on repair-project probability and disbursement rate. For Columns 1-4, the dependent variable is the budget-weighted mean probability that a department's project is classified as a repair project. For Column 5, the dependent variable is the budget-weighted mean of fiscal-year-end disbursement rate. $Post_t$ is a dummy variable that equals one for 2021-2023, and zero for 2018-2020. $Exposure_i$ is the pre-policy average share of construction projects in total budget (measured over 2018–2020). $Post_t \cdot Exposure_i$ is the interaction variable between $Post_t$ and $Exposure_i$ and captures heterogeneous effects based on policy exposure. Columns 1–4 incrementally include year fixed effects, department fixed effects, and control variables: the log number of projects and log total budget. Column (4) is the baseline specification. Standard errors are heteroscedasticity-robust and clustered at the department level. Robust standard errors are reported in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Across all specifications, we find a positive and statistically significant interaction between the post-policy period and construction exposure. In our baseline specification (Column 4 of Table 6), the interaction coefficient is 0.158. This indicates that departments with higher pre-policy reliance on construction projects significantly increased their share of repair-type investments following the reform. Interpreting the magnitude, a one–standard-deviation increase in pre-policy construction exposure (0.28) increases the repair probability by 4.4 percentage points, which is roughly 8% of the pre-policy mean repair probability (0.54). This represents a meaningful reallocation toward repair-type projects, consistent with departments responding to the tighter disbursement timeline by favoring simpler, faster-executing investments.

It is also important to assess whether the policy achieved its primary objective of accelerating investment disbursement. To test this, we examine the effect on the budget-weighted average disbursement rate at the end of the fiscal year (Column 5 of Table 2). We find that departments with higher pre-policy reliance on construction projects experienced a significantly faster disbursement rate following the reform. A one–standard-deviation increase in pre-policy exposure raises the disbursement rate by 3.1 percentage points, equivalent to about a 5% increase relative to the pre-policy mean (0.63). This finding suggests that while the policy succeeded in improving disbursement efficiency, it likely did so at the unintended cost of steering departments toward repair-oriented projects, which are relatively easier to execute.

In addition to the placebo test and event-study estimates discussed earlier, we conduct a series of sensitivity tests to assess the robustness of our baseline findings. First, we replace the budget-weighted outcome with the unweighted (simple) mean probability that a project is classified as repair. As shown in Column (1) of Table 7, the interaction between the post-policy period and construction exposure remains positive and statistically significant.

Table 7. Sensitivity tests: Alternative outcome variables and policy exposure measure.

	Alternative outcome var.				Alternative exposure var.
	(1)	(2)	(3)	(4)	(5)
	Simple mean of repair probability	Repair keywords-Having any repair project (Binary)	Repair keywords-Share of repair projects	Alternative Seeded LDA keywords	Log(Policy Exposure)
$Post_t$	-0.369***	-0.014	0.043	-0.410***	-0.277***
	(0.022)	(0.021)	(0.026)	(0.033)	(0.036)
$Post_t \cdot Exposure_i$	0.190***	0.087**	0.027	0.139**	
	(0.059)	(0.041)	(0.035)	(0.053)	
$Post_t \cdot Log(Exposure_i)$					0.019**
					(0.009)
$Project\ numbers_{it}$	-0.002	0.045*	0.037*	-0.021	-0.012
	(0.016)	(0.027)	(0.021)	(0.018)	(0.018)
$Total\ Budget_{it}$	-0.005	-0.015	-0.070***	0.012	0.004
	(0.010)	(0.012)	(0.016)	(0.012)	(0.012)
$Constant$	0.729***	0.858***	0.618***	0.769***	0.689***
	(0.059)	(0.057)	(0.073)	(0.071)	(0.065)
Observations	684	684	684	684	684
Number of Departments	142	142	142	142	142
R-squared	0.377	0.035	0.145	0.290	0.312
Dept. FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Control	Yes	Yes	Yes	Yes	Yes

Notes: This table reports the sensitivity tests with respect to alternative outcome variables policy exposure measure. Column (1) uses the unweighted (simple) mean repair probability as the outcome. Column (2) uses a binary indicator for whether a department has any project title containing repair or (improvement. Column (3) uses the share of projects containing these keywords. Column (4) uses an alternative set of Seeded LDA keywords. Column (5) uses the log of the construction share as an alternative policy exposure measure. $Post_t$ is a dummy variable that equals one for 2021-2023, and zero for 2018-2020. $Exposure_i$ is the pre-policy average share of construction projects in total budget (measured over 2018–2020). $Post_t \cdot Exposure_i$ is the interaction variable between $Post_t$ and $Exposure_i$. $Log(Exposure_i)$ is the log of pre-policy average share of construction projects in total budget (measured over 2018–2020). (measured over 2018–2020). $Post_t \cdot Log(Exposure_i)$ is the interaction variable between $Post_t$ and $Log(Exposure_i)$. All specifications include department and year fixed effects and control for the log number of projects and the log of total investment budget. Standard errors are heteroscedasticity-robust and clustered at the department level. Robust standard errors are reported in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Second, we test alternative outcome definitions based on keyword matching. Specifically, we use simple dictionary-based indicators that flag projects with the Thai terms "Somsam" (repair) or "Prabprung" (renovation/improvement) in the project title. We construct two alternative outcomes: 1) a binary variable indicating whether the department has at least one such project in a given year (Column 2 of Table 7), and 2) the share of projects containing those keywords (Column 3). While this approach is highly interpretable, it likely undercaptures the full range of repair-type projects, particularly those using alternative phrasing or embedded in more complex project descriptions.

As shown in Table 7, the estimate in Column (2) is positive and statistically significant, while the estimate in Column (3) is positive but imprecisely estimated. As expected, the keyword-based measures yield smaller and less precise effects than the seeded LDA specification. Nonetheless, the direction of the estimates is consistent across approaches, supporting the robustness of our topic-model-based outcome measure.

In Column (4), we assess sensitivity to the choice of seeded LDA keywords by using a narrower set—som ("repair"), bamrung-raksa ("maintenance"), and prap-prung ("improve")—in place of the baseline list. The interaction term remains positive and statistically significant, indicating that our results are robust to seed-word selection.

We also test whether the results are sensitive to the functional form of the exposure variable. In Column (4), we use the log of the pre-policy construction share as an alternative exposure measure. The interaction remains positive and statistically significant. This suggests that our main results are not driven by the linearity assumption in exposure.

Finally, we explore heterogeneity in policy response with respect to department size, and ministry type. First, we divide departments into two groups, above and below the median of pre-policy average investment budget, to proxy for size (Columns 1–2 of Table 8). The estimated effect among smaller departments is positive but statistically insignificant, whereas the effect among larger departments is both larger in magnitude and statistically significant. This pattern is consistent with the interpretation that departments managing larger and more complex project portfolios face greater exposure to execution risk under the reform and therefore exhibit stronger adjustment responses.

Next, we examine heterogeneity by ministry type, grouping departments into three broad categories: (1) Infrastructure & Social Services, (2) Industry & Resources, and (3) Others. While the estimated effects are positive across all three groups, they are less precisely estimated, with the largest effect observed among departments under Infrastructure & Social Services ministries.

However, interpreting heterogeneity by ministry type requires caution. Ministries often oversee departments with diverse operational roles, which may not align neatly with broad functional categories. For example, within the Ministry of Agriculture and Cooperatives, the Royal Irrigation Department is responsible for large-scale infrastructure projects such as dam construction and canal systems, whereas the Department of Agricultural Extension primarily provides training, support, and subsidies to farmers, with far fewer capital-intensive projects. As such, ministry affiliation alone may not consistently reflect the type of investment activities undertaken, making this source of heterogeneity less informative than department size.

Table 8. Heterogeneity analyses: Size and ministry type.

	Size		Ministry Type		
	(1)	(2)	(3)	(4)	(5)
	Small	Large	Infra&Social	Industry& Resources	Others
$Post_t$	-0.357***	-0.362***	-0.338***	-0.411***	-0.353***
	(0.037)	(0.038)	(0.042)	(0.060)	(0.042)
$Post_t \cdot Exposure_i$	0.104	0.191**	0.210*	0.064	0.098
	(0.079)	(0.089)	(0.117)	(0.131)	(0.079)
$Project\ numbers_{it}$	0.012	-0.035	0.016	-0.047*	-0.016
	(0.019)	(0.036)	(0.027)	(0.027)	(0.033)
$Total\ Budget_{it}$	-0.037**	0.025	-0.024	-0.006	0.021
	(0.016)	(0.016)	(0.027)	(0.019)	(0.017)
$Constant$	0.793***	0.596***	0.691***	0.897***	0.576***
	(0.077)	(0.096)	(0.107)	(0.088)	(0.096)
Observations	327	353	263	127	294
Number of Departments	67	73	55	27	63
R-squared	0.437	0.285	0.342	0.532	0.294
Dept. FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Control	Yes	Yes	Yes	Yes	Yes

Notes: This table presents heterogeneity analyses of the effects of the disbursement time-limit policy, based on department size and pre-policy disbursement rate. Columns (1) and (2) split departments into small and large based on the median of their pre-policy average total investment budget (measured over 2018–2020). Columns (3)-(5) split departments based on ministry type. The dependent variable is the budget-weighted probability that a department's investment project is classified as a repair project. $Post_t$ is a dummy variable that equals one for 2021-2023, and zero for 2018-2020. $Exposure_i$ is the pre-policy average share of construction projects in total budget (measured over 2018–2020). $Post_t \cdot Exposure_i$ is the interaction variable between $Post_t$ and $Exposure_i$. All specifications include department and year fixed effects and control for the log number of projects and the log of total investment budget. Standard errors are heteroscedasticity-robust and clustered at the department level. Robust standard errors are reported in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

CONCLUSION

This paper examines the unintended consequences of a fiscal rule aimed at improving investment disbursement efficiency. Focusing on Thailand's 2021 reform, which introduced a binding one-year disbursement deadline for public investment budgets, we combine seeded LDA with a DID design to study how government departments adjusted their investment portfolios in response to the new constraint.

We find that departments with greater pre-policy reliance on construction projects significantly increased their share of repair-oriented investments following the reform. Event-study estimates show that this adjustment unfolded gradually, consistent with budget planning and procurement cycles. The response is particularly pronounced among departments with larger investment budgets, which face greater exposure to execution risk when managing complex, multi-year projects.

While the reform appears to have achieved its goal of accelerating disbursement, it did so at the cost of shifting investment away from new construction toward simpler, short-term repairs. This pattern points to a fundamental trade-off in fiscal rule design between promoting timely budget execution and shaping the composition of public investment.

More broadly, the study demonstrates how text-based methods can be used to extract policy-relevant measures from administrative data. This is particularly useful for monitoring investment composition in settings where structured project classifications are limited. Importantly, our analysis does not assess whether repair-oriented investment is inherently inferior to new construction; repair and maintenance play a critical role in preserving public capital. Rather, our contribution is to show that the reform altered the propensity of departments to shift toward repair-oriented projects—a response that was not an explicit objective of the policy. These findings underscore that even wellintentioned fiscal rules can influence investment decisions in subtle but consequential ways, highlighting the importance of aligning such rules with long-term development priorities.

REFERENCES

Ardanaz, M., Cavallo, E., Izquierdo, A., & Puig, J. (2021). Growth-friendly fiscal rules? Safeguarding public investment from budget cuts through fiscal rule design. *Journal of International Money and Finance*, *111*, 102319.

Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent dirichlet allocation. *Journal of machine Learning research*, *3*(Jan), 993-1022.

Burret, H. T., & Feld, L. P. (2018). (Un-) intended effects of fiscal rules. *European Journal of Political Economy*, *52*, 166-191.

Cruciani, C., & Santagiustina, C. R. (2023). The present and future of sustainability disclosure in equity investment funds' pre-contractual documents: Mapping ESG discourse through STM. *Finance Research Letters*, *58*, 104033.

Dabla-Norris, E., Brumby, J., Kyobe, A., Mills, Z., & Papageorgiou, C. (2012). Investing in public investment: an index of public investment efficiency. *Journal of Economic Growth*, *17*(3), 235-266.

Dahan, M., & Strawczynski, M. (2013). Fiscal rules and the composition of government expenditures in OECD countries. *Journal of Policy Analysis and Management*, *32*(3), 484-504.

Delgado-Téllez, M., Gordo, E., Kataryniuk, I., & Pérez, J. J. (2022). The decline in public investment:``social dominance''or too-rigid fiscal rules?. *Applied Economics*, *54*(10), 1123-1136.

Hansen, S., McMahon, M., & Prat, A. (2018). Transparency and deliberation within the FOMC: A computational linguistics approach. *The Quarterly Journal of Economics*, *133*(2), 801-870.

Jagarlamudi, J., Daumé III, H., & Udupa, R. (2012, April). Incorporating lexical priors into topic models. In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics* (pp. 204-213).

Latifi, A., Naboka-Krell, V., Tillmann, P., & Winker, P. (2024). Fiscal policy in the Bundestag: Textual analysis and macroeconomic effects. *European Economic Review*, *168*, 104827.

Luangaram, P., & Wongwachara, W. (2017). More than words: a textual analysis of monetary policy communication. *PIER Discussion Papers*, *54*, 1-42.

Rajaram, A., Kaiser, K., Le, T.M., Kim, J.H., Frank, J. (2014) *The Power of Public Investment Management: Transforming Resources into Assets for Growth*. Washington, D.C.: World Bank.

Vinturis, C. (2023). How do fiscal rules shape governments' spending behavior?. *Economic Inquiry*, *61*(2), 322-341.

Watanabe, K., & Baturo, A. (2024). Seeded sequential LDA: A semi-supervised algorithm for topic-specific analysis of sentences. *Social Science Computer Review*, *42*(1), 224-248.

APPENDIX A: SUPPLEMENTARY ROBUSTNESS CHECKS

Table A1. Placebo test: Pre-policy trends in repair project probability (2018–2020).

	Repair-project probability	
	(1)	(2)
$Trend_t$	-0.041***	-0.044***
	(0.015)	(0.015)
$Trend_t \cdot Exposure_i$	0.031	0.032
	(0.038)	(0.038)
Constant	0.636***	0.799***
	(0.012)	(0.100)
Observations	337	337
R-squared	0.317	0.325
Dept FE	Yes	Yes
Year FE	Yes	Yes
Control	No	Yes
Number of Departments	140	140

Notes: This table reports results from a placebo test conducted using only the pre-policy period (2018-2020). The dependent variable is the budget-weighted mean repair project probability, defined as the average probability that a department's investment projects are classified as repair projects. Trend denotes a linear time trend over the pre-policy period. Exposure is the department's pre-policy average share of construction projects in total investment budget (measured over 2018–2020). The interaction term Trend × Exposure tests whether departments with higher construction exposure exhibited differential trends in repair intensity prior to the reform. Standard errors are heteroscedasticity-robust and clustered at the department level. Robust standard errors are reported in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

APPENDIX B: TEXT CLASSIFICATION METHODOLOGY

Text preprocessing

The analysis begins with a preprocessing phase to prepare the Thai-language text for modeling. Project titles are written in Thai and consist of short, standardized descriptions. This phase involves tokenizing the text (breaking it into individual words) and removing common words, stop words (frequent but uninformative terms such as “lae” (and) or “samrap” (for), as well as punctuation and address-related terms. These steps ensure that only the most meaningful words remain for analysis.

Seeded LDA framework

We employ a seeded Latent Dirichlet Allocation (LDA) model, which builds on the standard LDA framework by incorporating prior information through topic-specific seed words (Watanabe & Baturo, 2024). As in classical LDA, each document is modeled as a mixture of latent topics, and each topic is represented as a distribution over words. The seeded variant modifies the prior distributions to assign greater probability mass to predefined seed terms within designated topics, anchoring topic formation in substantively meaningful categories while preserving flexibility in estimation.

Formally, seeded LDA retains the same generative assumptions as classical LDA: each document d is represented as a multinomial distribution θ_d over topics, drawn from a Dirichlet prior Dir(α), and each topic z is characterized by a multinomial distribution $\emptyset_z$ over words, drawn from Dir(β). However, in seeded LDA, the prior on $\emptyset_z$ is modified such that seed words receive higher pseudo-counts in their designated topics, effectively encoding prior beliefs into the estimation process.

The model is estimated using a variational expectation–maximization (VEM) algorithm, which approximates the posterior distributions of document-level topic proportions and topic-level word distributions.

Seed word selection and implementation

In our application, we define a single topic of interest corresponding to repair and maintenance activities and specify a set of Thai seed words that are known ex ante to characterize such projects. The full set of seed terms includes: som (repair), som-saem (repair), bamrung (maintain), raksa (maintain), burana (restore), kaekhai (fix), fuen-fuu (rehabilitate), bamrung-raksa (maintenance), prap-prung (improve), prap-plian (modify), som-bamrung (maintenance), khut (dredge or excavate), lok (dredge), phiw-jarajon (road surface), and phiw (surface). These terms reflect common administrative language used in Thai public investment projects related to repair, rehabilitation, maintenance, and surface improvement.

The seeded LDA model is estimated using a variational expectation–maximization (VEM) algorithm, which approximates the posterior distributions of topic proportions for each document and word distributions for each topic. Prior information enters the model through modified Dirichlet priors that place greater weight on the predefined seed words, anchoring the repair topic while allowing the remaining word–topic associations to be learned from the data.

Seed terms are selected based on expert knowledge of public investment terminology in Thailand. The model is estimated on the full corpus of project titles, yielding posterior probabilities that each project corresponds to a repair-oriented investment. These probabilities are subsequently aggregated to the department–year level using budget weights to construct the repair intensity measures employed in the empirical analysis.

Generalization beyond keyword matching

A key advantage of seeded LDA is its ability to generalize beyond explicit keyword presence. While documents containing seed words are more likely to be assigned to the repair topic, the model also infers topic probabilities for documents that lack seed terms by learning patterns of word co-occurrence across the corpus. For example, if terms such as thom-din (earth fill), lat yang (asphalt paving), or phiu thang (pavement surface) frequently appear in documents that also contain repair-related language, the model learns to allocate positive repair-topic probability to those terms as well.

At the same time, seeded LDA avoids naive keyword tagging. Some projects that contain seed words such as som-saem (repair) receive low repair-topic probabilities when their broader context aligns more closely with new construction or infrastructure expansion. This outcome reflects the model's ability to balance local lexical evidence with global contextual information, rather than relying solely on the presence of individual keywords.

Comparison with large language models and limitations

Unlike large language models (LLMs), which rely on opaque neural representations and may yield inconsistent classifications, seeded LDA incorporates expert-defined seed terms while still learning contextual patterns from the corpus. This balance is essential in our setting, where project titles are formulaic and specific to Thai public-sector construction terminology—features that often cause LLMs to over-generalize or misinterpret administrative language.

Nevertheless, the seeded LDA approach has limitations, including sensitivity to the choice of seed words and the absence of the deeper semantic representations available in LLM-based models. We conduct sensitivity analyses using alternative seed word lists and a simple keyword-based classification. These robustness checks are discussed in the empirical finding section.