Reserve-Price Disclosure in Public Procurement Auctions: Evidence from Japan

David Lander* & Kosei Tanaka†

April 23, 2026

Abstract

This paper studies the causal effect of reserve-price disclosure in public procurement auctions, in the context of Japan's price-cap system. We exploit a policy change under which the national government encouraged local governments to replace ex-ante disclosure with ex-post disclosure, i.e., to keep the price cap hidden until after the auction. Using a municipality-level panel on procurement rules and average winning bid rates for construction procurement auctions from 2008 to 2023, we identify municipalities that switched disclosure regimes and estimate the effect of the reform using staggered difference-in-differences methods robust to treatment-effect heterogeneity. We find that adopting ex-post disclosure increased the average winning bid rate by about 1.2–1.6 percentage points across specifications. Thus, hiding the price cap raised the average normalized winning bid across auctions conducted within a municipality-year. The results provide causal evidence that reserve-price disclosure can materially affect procurement outcomes in a high-stakes public-sector setting.

JEL codes: D44, D47, H54, H57, L74.

Keywords: reverse auction, secret reserve, price cap, price ceiling, Japanese economy.

*Affiliation: Peking University HSBC Business School. Correspondence: lander@phbs.pku.edu.cn.

†Affiliation: Peking University HSBC Business School.

1 Introduction

Public procurement auctions often impose a reserve price or price cap, but an important design question remains unresolved: should that cap be disclosed to bidders before bidding begins? The answer is not obvious. Ex-ante disclosure may facilitate coordination or focal-point bidding near the cap, thereby raising winning bids. But hiding the cap may either raise winning bids, by increasing bidders' uncertainty about the valid bidding range or discouraging participation, or lower them, by making coordination more difficult. Because these channels point in different directions, the effect of disclosure timing on procurement outcomes is ultimately an empirical question.

This paper studies that question in the context of Japan's public procurement system, where construction tenders are subject to a reserve price that functions as a maximum acceptable bid. Over time, Japanese municipalities have differed in whether they disclose this price cap before the auction or only after it concludes. We exploit a policy-driven shift in disclosure timing: following a campaign promoted by the national government, a number of municipalities switched from ex-ante to ex-post disclosure. This staggered adoption creates quasi-experimental variation that allows us to estimate the effect of hiding the price cap on procurement outcomes.

Our analysis uses a municipality-level panel constructed from annual surveys conducted by Japanese national ministries. The data record each municipality's disclosure regime and, for each fiscal year, the arithmetic average of the winning bid rate across its construction procurement auctions. We use this municipality-year average winning bid rate (AWBR) as our main outcome and estimate treatment effects using staggered difference-in-differences (DID) methods robust to treatment-effect heterogeneity. Because the available data are aggregated at the municipality level, our estimates identify reduced-form effects on a municipality-level summary measure of bidding outcomes rather than auction-level effects for a fixed set of comparable projects.

We find that municipalities switching from ex-ante to ex-post disclosure experienced a positive and economically meaningful increase in the AWBR. Our main estimate is about 1.6 percentage points; however, reweighting treated and control municipalities on baseline observables attenuates this estimate to around 1.2 percentage points. Thus, across all specifications, hiding the price cap increased the average normalized winning bid across auctions conducted within a municipality-year. We then examine whether this increase was accompanied by higher aggregate municipal civil-engineering expenditure. Using the same staggered-DID framework, we do not detect a post-treatment increase in aggregate expenditure. This pattern is consistent with rigid local government budgets, but it may also reflect adjustment along the quantity margin or offsetting changes in project composition.

The paper contributes to the literature on reserve-price disclosure in three ways. First, to the best of our knowledge, it provides the first quasi-experimental evidence on reserve-price disclosure and the first such evidence in public procurement. Existing empirical work has relied mainly on structural counterfactuals or experimental variation in seller-auction and laboratory settings. Most of these settings have considerably lower stakes than public procurement auctions. Second, we study a real public procurement environment in which disclosure timing is an institutional rule rather than a seller's discretionary design choice. Third, the paper provides reduced-form evidence from a high-stakes public-sector setting, where the policy relevance of disclosure timing is direct even though the available data do not allow us to distinguish among the mechanisms behind the observed effect.

The rest of the paper proceeds as follows. Section 2 reviews the related literature and situates our contribution within it. Section 3 provides institutional background on Japan's procurement system and the debate over disclosure timing, while Section 4 describes the data. Section 5 outlines the empirical strategy, and Section 6 presents the main results together with a range of robustness checks. Section 7 concludes.

2 Related literature

This paper relates to two broad literatures on reserve-price secrecy and disclosure: theory, which studies when keeping the reserve secret may be desirable, and empirical work, which has examined the question mainly through structural counterfactuals and experimental variation. Across these literatures, the common question is whether secrecy can improve seller or auctioneer outcomes, and if so, under what conditions.

2.1 Theory

The theoretical literature begins from a benchmark in which secrecy is difficult to rationalize. In the standard independent-private-values, risk-neutral first-price framework, secret reserves do not generally outperform public reserves for the seller (Elyakime et al., 1994). This benchmark creates a "puzzle": secret reserves are frequently observed in practice, yet canonical models do not provide a simple general argument for their use.[1]

Subsequent theory identifies specific environments in which secrecy may matter. Secrecy may affect participation and information transmission in common-value or interdependent-value settings (Vincent, 1995), and may interact with resale or signaling incentives in ways

[1] Similar difficulties arise in canonical second-price environments and related extensions with entry costs and seller risk aversion (Lovo and Spaenjers, 2017; Moreno and Wooders, 2017; Lander and Li, 2026).

that public reserves cannot (Horstmann and LaCasse, 1997). Other papers rationalize secrecy through bidder risk aversion (Li and Tan, 2017), imperfect beliefs about past auction data (Jehiel and Lamy, 2015), reference-dependent preferences (Rosenkranz and Schmitz, 2007), and expectations-based loss aversion (Balzer and Rosato, 2025). Secrecy may also be valuable when the seller faces informational constraints: Rosar (2014) studies a setting in which a secret reserve helps the seller respond to information learned over time about the seller's own value, while Andreyanov and Caoui (2022) show how secrecy can allow the seller to learn from bids about auction-specific heterogeneity.

Taken together, this literature does not deliver a single general prediction in favor of either public or secret reserves. Rather, the desirability of secrecy depends on the auction environment, bidders' beliefs and preferences, and the seller's information. That theoretical ambiguity is important for our paper, because it implies that the effect of disclosure timing is ultimately an empirical question.

2.2 Prior empirical evidence

Existing empirical evidence is limited and falls mainly into two groups. The first uses structural models to estimate primitives of the auction environment and compare secret and public reserve regimes through counterfactual simulation. Foundational structural evidence comes mainly from seller auctions. Elyakime et al. (1994), using first-price timber auctions, reinforce the benchmark conclusion that public reserves dominate secrecy in standard environments. Li and Perrigne (2003), using French timber auctions with random reserve prices, and Eklöf and Lunander (2003), using Swedish apartment auctions with secret reserves, likewise find that reserve-price announcement improves seller outcomes in their applications. By contrast, the closest structural procurement paper, Ji and Li (2008), studies multi-round first-price sealed-bid procurement with secret engineer's estimates and finds that secrecy can reduce expected government payment in that setting. Andreyanov and Caoui (2022) provide a more recent structural contribution in which secrecy is useful because it allows the seller to learn from bids within the current auction. The strength of this literature is that it permits richer mechanism analysis and policy counterfactuals; its weakness is that conclusions depend on maintained assumptions about information, equilibrium behavior, and functional form.

The second group uses more design-based methods, especially experiments, to compare public and secret reserves under exogenous disclosure variation. The clearest field-experimental benchmark is Katkar and Reiley (2007), who randomize matched eBay auctions between a secret reserve and an economically equivalent public minimum bid and find that secret reserves reduce seller outcomes through lower sale probabilities and fewer serious bidders. Grether et al. (2016), combining field and laboratory evidence from salvage-vehicle auctions, find little difference in revenue between public and secret reserve disclosure.[2] Evidence closest to procurement is much thinner. Brisset et al. (2015) study a reverse-auction procurement environment in the laboratory and show that the relative performance of public and secret reserves depends on suppliers' risk aversion. Overall, the broader lesson from both structural and design-based evidence is that the effects of secrecy are highly institution-specific, and that credible empirical evidence from real public procurement remains scarce.

2.3 Contribution of this paper

Our paper contributes to this literature by exploiting a policy-driven change in reserve-price disclosure timing in a real public procurement environment. In contrast to private seller markets, where secrecy is typically a discretionary design choice of the seller, disclosure timing in our setting changed because municipalities responded to a nationally promoted reform agenda. This generates quasi-experimental variation on a question that the literature has studied mainly through theory, structural counterfactuals, or experiments.

Relative to the structural literature, our approach is less suited to decomposing precise mechanisms or solving for the seller's optimal reserve policy under hypothetical environments. But that is also its value: we provide comparatively assumption-light reduced-form evidence without imposing a structural bidding model. Relative to the experimental literature, our setting offers greater external validity and much higher economic stakes because it concerns bidding in government procurement rather than low-value consumer auctions or stylized laboratory environments. To the best of our knowledge, this makes our paper the first quasi-experimental evidence on reserve-price disclosure and the first such evidence in public procurement.

3 Institutional background

Public procurement commonly uses reverse auctions. The government announces a project, suppliers decide whether to participate and submit bids, and the lowest valid bid typically wins. Some tenders use a comprehensive-evaluation procedure in which price and non-price criteria are combined into a score, with the highest score winning. Many procurement systems also impose a "price floor," or minimum acceptable bid, to discourage abnormally low offers.[3]

[2] Related design-based evidence includes Walley and Fortin (2005) for online auctions and Kimbrough et al. (2024) on preferences over reserve-price publicity in the laboratory.

[3] Recent work by Chassang and Ortner (2019) suggests that a price floor may also inhibit bid rigging, by limiting the range of possible "punishments", i.e., low bids, after observed deviations from collusive bidding strategies.

3.1 Price cap in Japan's public procurement auction system

While the auction system used by Japanese government agencies is fundamentally similar to the general structure described above, it incorporates a distinctive feature: the use of a reserve price (or "price cap").[4] When the government plans to hold a tender, officials begin by making an estimate for the price of the project. To do so, they examine market prices for comparable goods or services, refer to similar past contracts, and use this information to set a reasonable benchmark, including considerations for materials, labor, transportation, and necessary equipment. For large-scale or specialized procurements, external consultants may also be involved to ensure that the price is appropriate. This estimated price becomes the price cap of the project and serves as the maximum acceptable price.

The price floor is set after the price cap has been fixed. It is calculated based on a formula predetermined by the auctioneer. In the case of construction-related procurement, it generally falls within 75–92% of the price cap. Together, the cap and floor define a band representing the acceptable range of bids: bids within this range are considered valid, while those outside are automatically disqualified. If no valid bids are received, the auctioneer may either rerun the auction, encouraging bidders to adjust their pricing, or cancel the auction and revise the price cap estimation before reopening the tender.

While the price floor aims to prevent unrealistically low bids, the price cap is intended to prevent the contract from being awarded at an excessively high price. For example, if only one supplier is capable of providing the required goods or services at a given time, that supplier may submit a bid that significantly exceeds the price anticipated by officials. Such a situation could hinder budget control and distort the procurement process of public agencies. In this case, the price cap serves as the upper limit on the winning bid and helps ensure that the project cost stays within the range estimated by the officials.

The use of a price cap in Japanese public procurement auctions gives rise to the concept of the "winning bid rate," which measures the ratio of the winning bid to the cap. For example, if the price cap for a project is 1,000,000 JPY and a bidder wins with a bid of 875,000 JPY, the winning bid rate would be 87.5%. Since this represents the normalized size of the winning bid relative to the maximum acceptable price for each project, it can be considered an indicator of the level of winning bid prices. This measure is widely used as an outcome variable in studies on Japanese public procurement auctions, since it enables consistent comparisons.[5]

[4] In the context of procurement, it has also been referred to as a pre-determined price, maximum estimated price, and engineer's estimate. In Japanese, it is known as *yotei kakaku*; literally, target price (Ohno and Harada, 2006; Kusunoki, 2007). Henceforth, we use the term price cap to refer to it.

[5] See, e.g., Ohashi (2009), Chassang and Ortner (2019), Kawai and Nakabayashi (2022).

3.2 Disclosure timing of the price cap

Throughout Japan, the timing of price-cap disclosure varies across local governments. While some governments keep the price cap confidential and disclose it only after the auction (ex-post disclosure), as with the price floor, others include the price cap in the tender notice when publishing the project (ex-ante disclosure). In addition, some governments adopt a mixed approach: for example, they set a price threshold, such that if the price cap for a project exceeds that threshold, it is withheld until the auction concludes.

According to a survey conducted by Japan's national ministries in 2023, among the 47 prefectural governments, 18 (38.3%) follow ex-post disclosure, 13 (27.7%) follow ex-ante disclosure, and 16 (34.0%) use both methods. Among the 1,741 municipal governments, 659 (37.9%) follow ex-post, 647 (37.2%) follow ex-ante, 323 (18.6%) use both, and the remaining 112 (6.4%) do not disclose the price cap, most of which are small towns and villages.[6]

One might assume that these differences reflect an urban-rural divide, but this is not necessarily the case. For example, among the ten most populous prefectures in Japan, four, including Tokyo and Osaka, use both methods. Of the remaining six, four adopt ex-post disclosure, while two, including Aichi, use ex-ante. This pattern is also observed at the municipal level. Figure 1 illustrates the location of each municipality, colored according to its disclosure status as of 2023: red, blue, yellow, and white correspond to ex-post, ex-ante, both, and non-disclosure, respectively. While some clustering of municipalities with the same disclosure method is observed, these patterns do not appear to be systematically associated with urbanization.

3.3 Policy debate over disclosure timing

The timing of price-cap disclosure has long been debated in Japan by national ministries, advisory councils, and local governments. The core issue is whether disclosing the price cap before bidding improves procurement outcomes or instead facilitates undesirable bidding behavior. The arguments run in both directions, and they imply different predictions for the winning bid rate.

Arguments in favor of ex-post disclosure emphasize auction fairness. A commonly stated concern is that ex-ante disclosure may facilitate bid rigging by revealing the maximum acceptable bid in advance.[7] Even absent explicit collusion, some policymakers argue that a publicly

[6]Specifically, among the 112 municipalities, 74 do not disclose the price cap for any tenders, while 38 follow a general policy of non-disclosure but may disclose it in exceptional cases. However, the survey does not provide information on the specific rules or conditions under which disclosure occurs for the latter group.

[7]For recent studies on collusion in Japanese procurement auctions, see Chassang et al. (2022), Kawai and Nakabayashi (2022), and Kawai et al. (2023).

known cap may serve as a focal point and thereby push bids upward. Others note that because the price floor is formulaically related to the cap, ex-ante disclosure may also make it easier for bidders to infer the floor, leading to clustering near the lower bound and more lotteries among identical lowest bids.

Figure 1: Map of Japanese municipalities by price-cap disclosure timing

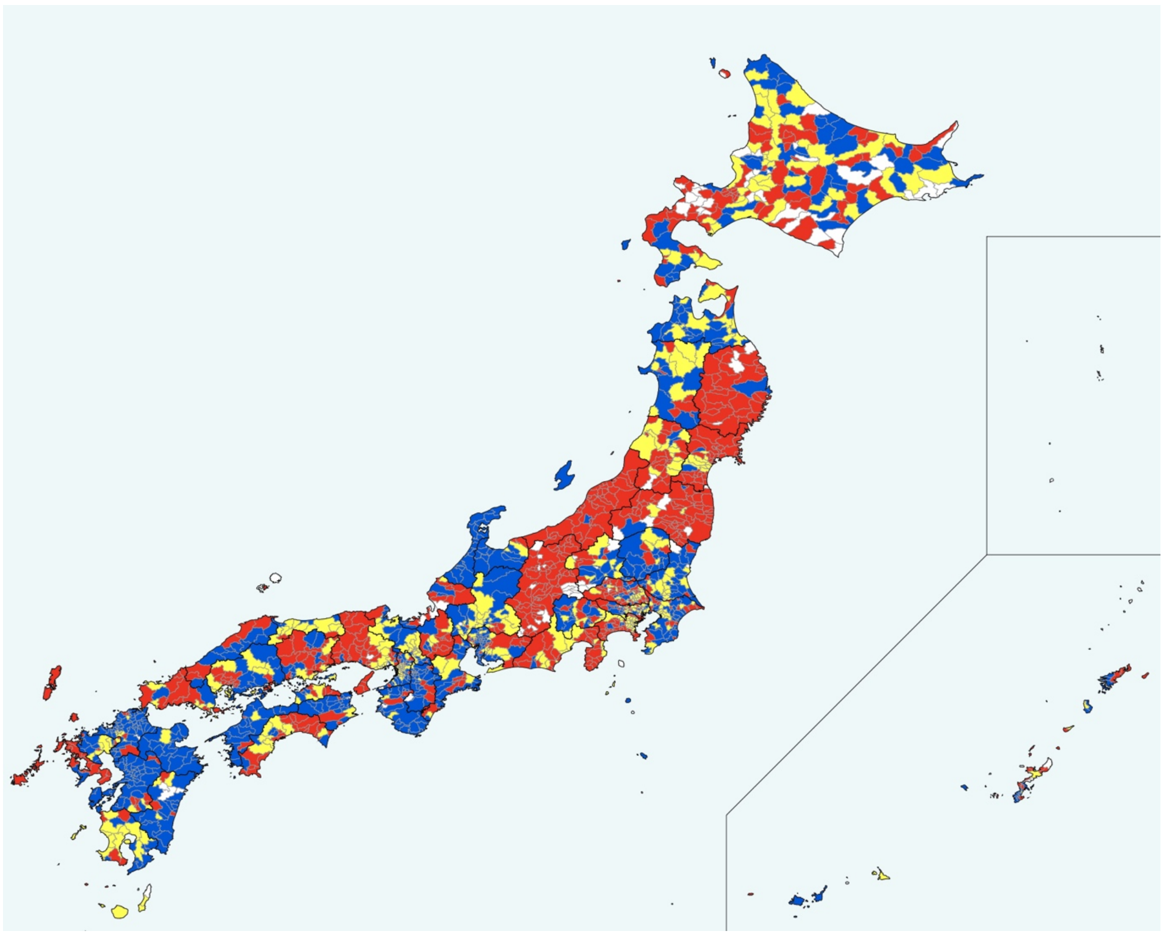

Note: Each municipality is colored according to its disclosure status (as of 2023): red, blue, yellow, and white correspond to ex-post, ex-ante, both, and non-disclosure, respectively. See Appendix Figure A1 for data sources.

Arguments in favor of ex-ante disclosure emphasize bidder uncertainty and participation. If the cap is hidden, suppliers must estimate the valid bidding range themselves, which may induce less aggressive (i.e., higher) bids, discourage participation, or increase the probability of unsuccessful procurement due to invalid bids. Some policymakers also argue that ex-post disclosure may disproportionately disadvantage smaller firms, which may have fewer resources to estimate the bidding range, and may create incentives to obtain cap information through informal or illicit means. More generally, some observers contend that disclosure timing may matter little if bidders can infer hidden caps reasonably well from past procurement records.

For the purposes of this paper, the main implication of this policy debate is that the sign of the effect on the winning bid rate is ambiguous *a priori*. Some arguments imply that ex-ante disclosure raises winning bids, while others imply the opposite. Thus, the Japanese policy debate itself does not resolve whether switching from ex-ante to ex-post disclosure should increase or decrease normalized winning bids.

Finally, several municipalities have compared AWBRs under the two regimes, but these exercises use non-random project assignment, small samples, and little formal inference. Their mixed findings are therefore descriptive rather than causal and motivate our design.[8]

3.4 National promotion of ex-post disclosure

Despite the lack of clear consensus in the policy debate, the Japanese national government has generally taken a position in favor of ex-post disclosure. In particular, national guidance emphasized concerns about bid rigging and lottery-based tender outcomes under ex-ante disclosure, and encouraged local governments to reconsider or discontinue ex-ante disclosure. This promotional effort took various forms, including national guidelines and repeated request letters to local governments. As a result, the share of municipalities using ex-ante disclosure declined over time, while the share using ex-post disclosure increased.[9] It is this policy-driven shift in local disclosure regimes that generates the variation exploited in our empirical analysis.

4 Data

4.1 Survey by National Ministries

While analyses of the causal effects of procurement rule changes typically rely on quasi-experimental approaches using auction-level data, such approaches require access to longitudinal data from the municipality where the policy change occurred. In this regard, Japanese local governments are required to disclose bidding results under the *Act for Promoting Proper Tendering and Contracting for Public Works* (Act No. 127 of 2000; henceforth “the Act”).

However, the Act only stipulates that such data must be made available “for at least one year.”[10] Consequently, many municipalities disclose only one or two years of recent data.

[8] These include the Fukushima Prefectural Government (Fukushima Prefectural Government, 2009), the City of Yokohama (Yokohama City Government, 2011), and Seki City in Gifu Prefecture (Seki City Government, 2022).

[9] Specifically, the proportion of municipalities adopting ex-ante disclosure declined from 53.0% in 2008 to 37.2% in 2023, while those adopting ex-post increased from 22.7% to 37.9% over the same period.

[10] Article 8, Paragraph 1 of the Act and Article 7, Paragraph 2 of the Enforcement Order mandate the disclosure of bidding results, while Article 7, Paragraph 6 of the Enforcement Order stipulates a minimum disclosure period of one year.

Although a small number of municipalities voluntarily maintain and publish longer-term records, they remain a minority. While we identified all municipalities that switched from ex-ante to ex-post disclosure between 2009 and 2023, none of them provide auction data that both cover the timing of the policy change and span a sufficient period to permit estimation using causal-inference methods.

Given the lack of auction-level data, we instead focus on municipality-level data that the Japanese national government has been collecting since the early 2000s. Specifically, since 2002, three ministries have jointly conducted annual surveys and published information on the bidding systems used by national agencies and local governments.[11,12] Although the survey items vary slightly from year to year, they have consistently collected information on the timing of price cap disclosure for all municipalities since 2008. This enables us to identify the year in which each local government changed its disclosure timing, which serves as the treatment timing in our staggered-DID analysis.

In addition to the timing of price-cap disclosure, this survey includes another important variable: the arithmetic average of the winning bid rate of all construction project auctions held in each municipality in each fiscal year. While the winning bid rate reflects the normalized winning bid in an individual auction, the average winning bid rate (AWBR) aggregates these auction-level ratios to the municipality-year level. We use the AWBR as our primary outcome because it is the only consistently available measure that summarizes bidding outcomes for the full set of municipalities over time.[13]

4.2 Sample construction

The survey covers 1,741 municipalities over 16 years, from 2008 through 2023. We classify municipalities that used ex-ante disclosure in 2008 and switched to ex-post disclosure between 2009 and 2023 as treated, and municipalities that used ex-ante disclosure throughout the sample period as "never-treated" controls.[14,15]

Some municipalities adopt mixed approaches; however, the specific criteria each municipality uses to determine which auctions are conducted under ex-ante or ex-post disclosure vary, and the survey does not capture such detailed information. Thus, it is not possible to quantify

[11] Source: https://www.mlit.go.jp/totikensangyo/const/1_6_bt_000154.html (Accessed 15 April 2025)

[12] Article 22 of the Act stipulates that the Minister of Land, Infrastructure, Transport and Tourism, the Minister for Internal Affairs and Communications, and the Minister of Finance shall make efforts to collect, organize, and provide information that contributes to better bidding practices in public construction works.

[13] The AWBR has also been used by Arai (2013), who utilized the same municipal-level survey data as our study.

[14] In a robustness exercise we also obtain estimates using "not-yet-treated" municipalities as our control group.

[15] It is worth noting briefly that examining the shift from ex-post to ex-ante is also theoretically a meaningful exercise, since it could conceivably produce an effect of similar magnitude in the opposite direction. Unfortunately, very few municipalities opted to switch from ex-post to ex-ante: only 14 units made this transition in our sample.

the extent to which each method is used, and we exclude municipalities that employed both approaches during the study period from the analysis. Municipalities that do not disclose the price cap at all are also excluded. Extremely small local governments may also hold no tenders in a given year and therefore report no AWBR.

After excluding these municipalities, we obtain 49 treatment units and 242 control units. Since each unit is observed in every period from 2008–2023, our sample is balanced, with a total of 4,656 observations. Table 1 reports group means for the AWBR and additional municipality-level characteristics, with variables averaged over years for which they are available.

Table 1: Summary statistics of key variables

	Treatment group	**Control group**
Average winning bid rate (%)	92.37 (3.70)	91.87 (3.89)
Population (thousand persons)	125.1 (213.6)	92.9 (191.1)
Labor force (thousand persons)	61.0 (105.2)	45.1 (94.4)
Unemployment rate (%)	4.6 (1.0)	5.2 (1.6)
Area (km^2)	276.6 (261.8)	196.0 (232.2)
Habitable land area (km^2)	102.2 (104.6)	76.1 (69.2)
Standard fiscal scale (billion JPY)	29.6 (51.3)	22.3 (48.6)
Fiscal strength index	0.62 (0.27)	0.57 (0.26)
Civil-engineering expenditure (billion JPY)	6.9 (15.1)	4.4 (12.0)
Number of construction firms	515.7 (852.1)	361.3 (714.5)
Number of municipalities	49	242
Total observations	784	3,872

Note: Each column reports group means and standard deviations (in parentheses) for the full sample, including Shoo Town. The data sources used to construct each variable are described in Appendix Table A2.

Before we proceed to discuss our empirical strategy, we note that one of the 49 treated municipalities (Shoo Town) appears to be a significant outlier.[16] Therefore, throughout the paper we report all treatment-effect estimates using a subsample that *excludes* this municipality.[17]

[16]Among the 49 treatment municipalities, Shoo Town experienced a sharp drop of approximately 30 percentage points in its AWBR in FY2011, followed by a similarly sharp rebound the following year.

[17]Descriptive statistics for the subsample excluding Shoo Town are reported in Appendix Table A2, and estimates obtained using the full sample appear in Appendix Figure B3.

4.3 Selection

Given that the national reform merely *encouraged*, rather than required, local governments to adopt ex-post disclosure, one should be mindful of the possibility of non-random selection into treatment as well as treatment timing. To investigate this, we report a number of additional results in Appendix A and briefly describe their insights here.

Appendix Figure A1 illustrates the geographic distribution of the treatment and control units in our sample. Although a few municipalities appear to be clustered with nearby municipalities in the same group, the sample as a whole does not exhibit a clear spatial pattern and is largely scattered across Japan. For the covariates in our sample, treated and control municipalities are broadly similar at baseline, although treated municipalities tend to be somewhat less densely populated and somewhat more fiscally strong (Appendix Table A3). Importantly, baseline AWBR does not appear to predict adoption, but we do find evidence that log population density and fiscal strength have some predictive power for eventual switching (Appendix Table A4).

Regarding treatment timing, Appendix Table A1 lists the treatment municipalities by the year they switched from ex-ante to ex-post. While the national government began promoting ex-post disclosure approximately 20 years ago, many municipalities abandoned ex-ante disclosure in the early years following the promotion, whereas fewer have done so in recent years. As a result, the number of early treatment cases is relatively large, while recent cases are fewer. This imbalance limits the precision of longer-horizon pre-treatment lead estimates in the event study, so the pre-trends exercise should be interpreted as a useful but limited diagnostic for the parallel-trends assumption. Finally, similar to selection into treatment, early and late adopters appear similar in terms of most baseline observables (AWBR, log population, and log fiscal scale), but do differ in terms of log population density and fiscal strength (Appendix Table A5).

5 Empirical strategy

5.1 Heterogeneity-robust estimator

Because municipalities switched from ex-ante to ex-post disclosure at different times, we use a DID estimator designed for staggered-adoption settings. Following recent discussions by De Chaisemartin and d'Haultfoeuille (2023) and Roth et al. (2023), our preferred approach is the estimator of Callaway and Sant'Anna (2021).[18] This estimator is robust to treatment-effect heterogeneity and avoids the problematic comparisons that can arise when conventional

[18]Our main conclusions do not change when using the approach of Borusyak et al. (2024) (Appendix Figure B7).

two-way fixed effects (TWFE) estimators are applied in staggered-adoption settings.[19]

The estimator first computes group-time average treatment effects, i.e., the average treatment effect on the treated (ATT) for each cohort of municipalities that first switch from ex-ante to ex-post disclosure in year g, evaluated in year t. In our setting, this means comparing how the AWBR evolves for municipalities that adopt ex-post disclosure in a given year with the corresponding evolution for the chosen comparison group. These cohort-time effects are then aggregated into event-time treatment effects, which trace out how the AWBR changes in the years before and after the switch in disclosure regime.

An additional advantage of this framework is that it allows the researcher to choose the comparison group. In our main analysis, we use never-treated municipalities as the control group. We show in Appendix Figure B1 that our findings are not materially affected when not-yet-treated municipalities are also used as controls.

5.2 Identification

The key identifying assumption is that, absent the switch from ex-ante to ex-post disclosure, treated municipalities would have experienced the same evolution in the AWBR as the comparison municipalities. In our setting, this means that any difference in AWBR between switching and never-treated municipalities would have remained stable over time in the absence of the reform. Since the outcome is measured at the municipality-year level, the identifying variation comes from changes in the AWBR within municipalities over time, not from cross-sectional differences in levels.

This assumption is plausible in our context for several reasons. The treatment is a change in auction rules governing disclosure timing, rather than a change targeted at specific projects, sectors, or municipalities facing unusual cost shocks. Moreover, the national government promoted the move toward ex-post disclosure as part of a broad reform agenda, which makes the timing of adoption less likely to reflect a single municipality-specific change in procurement conditions. At the same time, because local governments retained discretion over adoption, we do not treat parallel trends as self-evident. We therefore study pre-treatment event-time coefficients and also report robustness checks that condition on baseline municipality characteristics and impose overlap-based sample restrictions. As we explain below, the event-study estimates do not reveal strong evidence against this assumption, although the pre-treatment diagnostics are necessarily imprecise at longer horizons because treatment is concentrated in earlier years.

[19] A key concern is that TWFE estimators compare units treated at different times, so earlier-treated units may serve as controls for later-treated units. If treatment effects are heterogeneous over time, such comparisons can bias the estimated effect; see, e.g., Goodman-Bacon (2021).

We do not view anticipatory responses as a major threat to identification in this setting. In principle, DID estimates may be biased if potential bidders adjust their behavior in anticipation of a future change in disclosure rules (Malani and Reif, 2015). In our context, however, such effects seem unlikely: the treatment changes the disclosure rule applying to the auction at the time of bidding, and we are not aware of evidence that announcements of future rule changes affect current-period participation or bid submission in procurement auctions. Consistent with this view, the event-study estimates do not reveal a clear pattern suggestive of anticipation.

6 Results

While our primary interest lies in the overall treatment effect, we begin in Section 6.1 with DID event-study estimates. These estimates allow us to assess treatment-effect dynamics and to inspect pre-treatment coefficients as a diagnostic for the parallel-trends (PT) assumption.

6.1 Dynamic effects

In a staggered-adoption setting, the number of treatment observations decreases in periods further from the event. In our sample, longer-horizon pre-treatment leads are based on very small numbers of treated municipalities. Therefore, in Figure 2, we present event-time coefficients from -8 to $+10$. The lower bound is determined by the requirement that each displayed pre-treatment lead include at least 10 treated municipalities, while the upper bound is chosen to align the figure with our preferred overall ATT, which aggregates treatment effects over event times 1 to 10.[20]

The post-treatment estimates show little evidence of substantial dynamics. This is consistent with the prediction that the timing of disclosure would influence bid results on a per-tender basis, rather than causing a systematic increase or decline over time. The event-time 0 estimate is slightly smaller than the later estimates, which may reflect attenuation from partial-year exposure.[21]

The event-study plot does not reveal a clear or systematic pattern of differential pre-treatment movements between treated and control municipalities. Most pre-treatment coefficients are imprecisely estimated and their confidence intervals include zero, while their signs fluctuate rather than pointing to a consistent upward or downward trend. At the same time, these lead estimates should be interpreted cautiously. Because many municipalities switched

[20] Estimates over the full event-time horizon are reported in Appendix B.

[21] While Japanese government agencies typically implement new systems at the beginning of the fiscal year, some may do so during the year. If several municipal governments switched mid-year and some tenders were still conducted under ex-ante disclosure, the increase in the AWBR would be somewhat lower due to these tenders.

in the earlier years of the sample, relatively few treated units contribute to longer-horizon pre-treatment leads, limiting the event study's power as a diagnostic for the PT assumption.

Figure 2: Event-study estimates of dynamic treatment effects

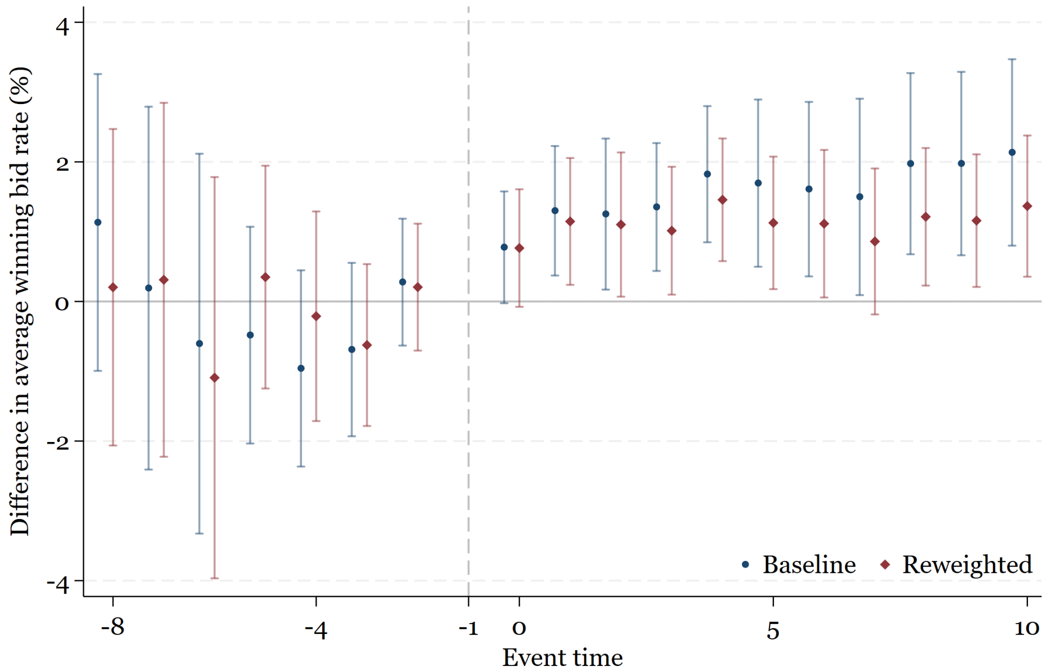

Note: The treatment units switch from ex-ante to ex-post at event time 0. The value at time -1 is set to zero since the other coefficients are measured relative to time -1 as the reference period. "Baseline" refers to unweighted estimates, while "Reweighted" refers to inverse-probability-weighted estimates, where the propensity score is estimated using pre-treatment values (i.e., reported for year 2008) of the following covariates: AWBR, log population, log population density, log standard fiscal scale and fiscal strength index. Standard errors are clustered at the municipality level, and confidence intervals are set at the 95% level.

Overall, the pre-treatment coefficients are broadly reassuring, but they do not by themselves provide strong evidence in favor of parallel trends; rather, they indicate that the event-study results do not reveal strong evidence against it. We reach the same conclusion in specifications that condition on baseline covariates (see Appendix Figures B6, B8, and B9).

6.2 Estimated treatment effect

As the estimate at event time 0 may be attenuated by partial-year treatment exposure, and estimates after event time 10 rely on relatively few treated municipalities, we aggregate the dynamic treatment effects from event times 1 to 10 to obtain our preferred overall estimate.[22]

[22] This choice excludes the contemporaneous effect at event time 0, which may be affected by within-year transitions from ex-ante to ex-post disclosure, as well as later post-treatment periods where the number of treated municipalities becomes small.

Column (1) of Table 2 reports an ATT of 1.64 percentage points. Thus, switching from ex-ante to ex-post disclosure increased the AWBR, on average, indicating that hiding the price cap raised winning bids relative to official price caps. This result suggests that the national government's campaign to promote ex-post disclosure may have inadvertently increased municipal procurement prices. Whether this effect translated into higher overall municipal spending is a separate question, to which we turn in Section 6.5.

To assess whether this estimate is sensitive to limited baseline imbalance between switching and never-treated municipalities, Table 2 also reports a set of additional estimates based on inverse-probability weighting (IPW) and sample restrictions.[23] These specifications allow us to distinguish between two concerns: whether the baseline estimate is influenced by a small number of control municipalities that receive large weights under reweighting, and whether it relies on comparisons with control municipalities that are observationally dissimilar to the treated group.[24]

Table 2: Estimated average treatment effect

	Main sample		**Trimmed sample**		**Common support**	
	(1)	(2)	(3)	(4)	(5)	(6)
	Baseline	Reweighted	Baseline	Reweighted	Baseline	Reweighted
ATT	1.638***	1.154***	1.648***	1.191***	1.509***	1.228***
	(0.518)	(0.410)	(0.517)	(0.420)	(0.520)	(0.420)

Note: All columns report the overall average treatment effect on the treated (ATT) for the average winning bid rate (AWBR), measured in percentage points, aggregated over event times 1 to 10. The "Baseline" columns report estimates without covariate adjustment, while "Reweighted" columns report estimates that reweight the control group using pre-treatment (i.e., year 2008) municipality characteristics: AWBR, log population, log population density, log standard fiscal scale, and fiscal strength index. Columns (1)–(2) use the full estimation sample, excluding Shoo Town. Columns (3)–(4) restrict the sample by excluding control municipalities with ATT weights above the 99th percentile. Columns (5)–(6) restrict the control group to municipalities whose estimated baseline propensity scores lie within the support of the treated group. Standard errors, reported in parentheses, are clustered at the municipality level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The estimates in Table 2 reveal a clear pattern. The unweighted estimates are highly stable across alternative sample restrictions, ranging from 1.51 to 1.65 percentage points. The reweighted estimates are likewise stable, ranging from 1.15 to 1.23 percentage points, and remain statistically significant at the 1% level. Thus, trimming high-weight controls and imposing common-support restrictions have relatively little effect within either weighting regime.

[23]The reweighted baseline covariates include the AWBR, log population, log population density, log standard fiscal scale, and fiscal strength index. See Appendix Figure A2 for further details.

[24]The trimmed sample excludes control municipalities with ATT weights above the 99th percentile, while the common-support restriction limits the control group to municipalities whose estimated baseline propensity scores lie within the support of the treated group; see, e.g., Crump et al. (2009) and Lechner and Strittmatter (2019).

The main sensitivity instead concerns whether treated and control municipalities are reweighted to align baseline observables, including baseline AWBR.[25] This attenuation suggests that differences in baseline observables may lead the unweighted specification to overstate the magnitude of the treatment effect. At the same time, all specifications imply a positive and economically meaningful increase in the municipality-year average winning bid rate following the switch to ex-post disclosure.

6.3 Robustness checks

In addition to the weighting and sample-restriction exercises summarized in Table 2, we examine several other dimensions of robustness, including an alternative estimator and subsample exercises that exclude specific treatment cohorts or regions. We briefly explain these here, and direct the reader to Appendix B for further details.

Appendix Figure B4 shows the results after sequentially excluding each treatment-year group. Since a relatively large number of municipalities switched to ex-post in 2009, the estimates become less stable when this group is excluded. However, even in this case, the coefficients at event times 0–4 remain positive and comparable in magnitude to the baseline estimates, where more treatment units are available. We also investigate possible regional sensitivity by dividing Japan into 12 geographic areas and re-estimating the coefficients after omitting the municipalities belonging to each region, one at a time. The results remain broadly consistent with the baseline estimates (Appendix Figure B5).

Finally, Appendix Figure B7 reports estimates from the imputation estimator of Borusyak et al. (2024), an alternative heterogeneity-robust approach to staggered-adoption designs. The post-treatment coefficients are consistently about 0.2–0.3 percentage points larger than the corresponding Callaway and Sant'Anna (2021) estimates.

6.4 Possible mechanisms

Several mechanisms are consistent with the finding that switching from ex-ante to ex-post disclosure increases the AWBR. One possibility is that withholding the price cap increases bidders' uncertainty about the valid bidding range, which may induce less aggressive (i.e., higher) bids. Another is that ex-post disclosure discourages participation by making it more costly for firms to assess whether a project is worth entering. Both mechanisms would tend to raise the winning bid rate. A third possibility is that ex-ante disclosure makes it easier

[25] This attenuation should not be read as evidence that baseline AWBR strongly predicts adoption on its own. Rather, it indicates that reweighting treated and control municipalities to improve balance on the full set of baseline observables, including the baseline outcome, yields a more conservative estimate of the treatment effect.

for bidders to infer the price floor, thereby causing bids to cluster near the lower bound and reducing the winning bid rate.

At the same time, the overall effect may reflect offsetting forces across auctions. In some settings, withholding the price cap could make collusive coordination more difficult and therefore reduce bid prices. If so, the positive average effect that we estimate would reflect the net impact of channels working in opposite directions. Unfortunately, our municipality-level data do not allow us to distinguish between these mechanisms, since we do not observe auction-level outcomes such as the number of participants, the distribution of bids, clustering near the price floor, or the frequency of unsuccessful tenders.

6.5 Impact on government expenditure

Although the estimates above indicate that switching from ex-ante to ex-post disclosure increased the municipality-level AWBR, a separate question is whether this translated into higher aggregate municipal spending on civil engineering. To examine this, we re-estimate the treatment effects using civil-engineering expenditure as the outcome and otherwise follow the same empirical design as in the main analysis. In particular, we use the same staggered-adoption DID framework and comparison-group structure as in Section 5.[26]

Figure 3 presents the corresponding dynamic estimates.[27] In contrast to the AWBR results, the expenditure coefficients do not display a clear or persistent positive shift following adoption of ex-post disclosure. The point estimates fluctuate around zero, and the confidence intervals are sufficiently wide that modest positive or negative effects cannot be ruled out. Accordingly, we interpret this exercise as providing no clear evidence that the disclosure reform changed aggregate municipal civil-engineering expenditure.[28]

One possible interpretation is that municipalities adjusted along margins other than aggregate expenditure. If local government budgets are rigid, higher normalized winning bids need not translate into higher total civil-engineering spending; instead, municipalities may respond by undertaking fewer projects or by altering the composition of projects procured. Under this interpretation, the increase in the AWBR would be absorbed through quantity or composition adjustments rather than through a rise in total expenditure.

At the same time, the expenditure results should be interpreted cautiously. The absence of a clear post-treatment expenditure response is also consistent with heterogeneous treatment effects across project types, which may offset one another in the aggregate. Our data do not

[26]For the weighted specification, we estimate adoption propensity scores using the same baseline municipality characteristics as in the main analysis.

[27]Appendix Figure B2 reports these dynamic estimates for the full event-time window.

[28]The same insights are obtained when studying the impact on log civil-engineering expenditure (not reported).

allow us to distinguish among these explanations, since we do not observe project counts or auction-level expenditure components. Thus, the expenditure analysis is best viewed as suggestive evidence that the reform did not generate a clear increase in aggregate spending, rather than as definitive evidence of budget rigidity.

Figure 3: Dynamic treatment effects on civil-engineering expenditure

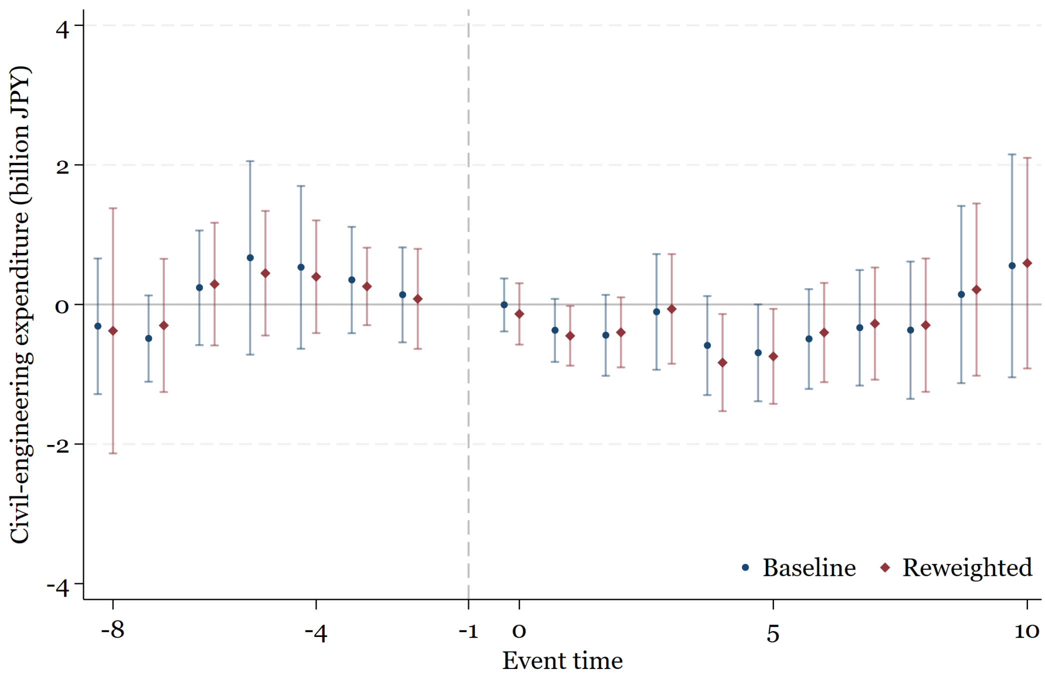

Note: This figure is analogous to Figure 2, with civil-engineering expenditure as the outcome. The treatment units switch from ex-ante to ex-post at event time 0. The value at time -1 is set to zero since the other coefficients are measured relative to time -1 as the reference period. "Baseline" refers to unweighted estimates, while "Reweighted" refers to inverse-probability-weighted estimates, where the propensity score is estimated using pre-treatment values (i.e., reported for year 2008) of the following covariates: AWBR, log population, log population density, log standard fiscal scale and fiscal strength index. Standard errors are clustered at the municipality level, and confidence intervals are set at the 95% level.

7 Conclusion

This paper estimates the effect of reserve-price disclosure in public procurement auctions using municipality-level panel data from Japan. Exploiting staggered municipal adoption of ex-post disclosure, we estimate the effect of switching from ex-ante to ex-post disclosure on the municipality-year arithmetic average of winning bid rates. Across all specifications, the estimated effect is positive and economically meaningful: the baseline estimate is about 1.6 percentage points, while reweighted specifications that improve balance on baseline observables imply effects of around 1.2 percentage points. Thus, the evidence indicates that hiding the price cap raised the average normalized winning bid across auctions conducted within a municipality-year. This conclusion is robust across a range of alternative specifications, although the aggregated nature of the outcome means that the estimates should be interpreted as reduced-form effects on a municipality-level summary measure of normalized winning bids, rather than as auction-level effects for a fixed set of comparable projects.

Placed against prior empirical evidence, this result is less anomalous than it might first appear. Although seller and procurement auctions differ, existing seller-auction evidence often finds that reserve secrecy either fails to improve or worsens outcomes for the party setting the reserve (Katkar and Reiley, 2007; Grether et al., 2016), while the limited procurement evidence is mixed (Ji and Li, 2008; Brisset et al., 2015). Moreover, theoretical settings in which secrecy can be valuable often rely on risk aversion, behavioral factors, or informational features that do not map neatly onto municipal procurement settings. Our results show that secrecy can likewise raise normalized winning bids in a reverse-auction setting.

The main limitation of the paper is data granularity. Because the available data are aggregated at the municipality level, we cannot observe auction-level outcomes such as participation, bid distributions, project counts, or unsuccessful tenders, and therefore cannot distinguish among the mechanisms that may underlie the increase in the AWBR. More detailed and consistently preserved auction-level data would make it possible to study these channels directly and to evaluate how disclosure rules affect procurement performance in greater depth. Even with this limitation, the evidence in this paper suggests that reserve-price disclosure can materially affect procurement outcomes in a high-stakes public-sector setting.

References

ANDREYANOV, P. AND E. H. CAOUI (2022): "Secret reserve prices by uninformed sellers," *Quantitative Economics*, 13, 1203–1256.

ARAI, K. (2013): "Effect of Institutions: Analysis of Japanese municipal public procurement," *International Journal of Public Administration*, 36, 638–648.

BALZER, B. AND A. ROSATO (2025): "Never say never: Optimal exclusion and reserve prices with expectations-based loss-averse buyers," *Journal of Economic Theory*, 228, 106045.

BORUSYAK, K., X. JARAVEL, AND J. SPIESS (2024): "Revisiting event-study designs: robust and efficient estimation," *Review of Economic Studies*, 91, 3253–3285.

BRISSET, K., F. COCHARD, AND J. LE GALLO (2015): "Secret versus public reserve price in an "outcry" English procurement auction: Experimental results," *International Journal of Production Economics*, 169, 285–298.

CALLAWAY, B. AND P. H. SANT'ANNA (2021): "Difference-in-differences with multiple time periods," *Journal of Econometrics*, 225, 200–230.

CHASSANG, S., K. KAWAI, J. NAKABAYASHI, AND J. ORTNER (2022): "Robust screens for noncompetitive bidding in procurement auctions," *Econometrica*, 90, 315–346.

CHASSANG, S. AND J. ORTNER (2019): "Collusion in auctions with constrained bids: Theory and evidence from public procurement," *Journal of Political Economy*, 127, 2269–2300.

CRUMP, R. K., V. J. HOTZ, G. W. IMBENS, AND O. A. MITNIK (2009): "Dealing with limited overlap in estimation of average treatment effects," *Biometrika*, 96, 187–199.

DE CHAISEMARTIN, C. AND X. D'HAULTFOEUILLE (2023): "Two-way fixed effects and differences-in-differences with heterogeneous treatment effects: A survey," *The Econometrics Journal*, 26, C1–C30.

EKLÖF, M. AND A. LUNANDER (2003): "Open outcry auctions with secret reserve prices: an empirical application to executive auctions of tenant owner's apartments in Sweden," *Journal of Econometrics*, 114, 243–260.

ELYAKIME, B., J. J. LAFFONT, P. LOISEL, AND Q. VUONG (1994): "First-price sealed-bid auctions with secret reservation prices," *Annales d'Economie et de Statistique*, 34, 115–141.

FUKUSHIMA PREFECTURAL GOVERNMENT (2009): "Yotei kakaku jigo kohyo no shiko jokyo no bunseki kekka ni tsuite," Fukushima: Procurement Oversight Division.

GOODMAN-BACON, A. (2021): "Difference-in-differences with variation in treatment timing," *Journal of Econometrics*, 225, 254–277.

GRETHER, D. M., D. PORTER, AND M. SHUM (2016): "Public vs. secret reserve prices in auctions: Evidence from combining field and lab experiments," *Working paper. Available at SSRN: https://ssrn.com/abstract=2766571 or http://dx.doi.org/10.2139/ssrn.2766571*.

HORSTMANN, I. J. AND C. LACASSE (1997): "Secret reserve prices in a bidding model with a resale option," *The American Economic Review*, 87, 663–684.

JEHIEL, P. AND L. LAMY (2015): "On absolute auctions and secret reserve prices," *The RAND Journal of Economics*, 46, 241–270.

Ji, L. and T. Li (2008): "Multi-round procurement auctions with secret reserve prices: theory and evidence," *Journal of Applied Econometrics*, 23, 897–923.

Katkar, R. and D. H. Reiley (2007): "Public versus secret reserve prices in eBay auctions: Results from a pokémon field experiment," *The BE Journal of Economic Analysis & Policy*, 6.

Kawai, K. and J. Nakabayashi (2022): "Detecting large-scale collusion in procurement auctions," *Journal of Political Economy*, 130, 1364–1411.

Kawai, K., J. Nakabayashi, J. Ortner, and S. Chassang (2023): "Using bid rotation and incumbency to detect collusion: A regression discontinuity approach," *The Review of Economic Studies*, 90, 376–403.

Kimbrough, E. O., P. Limberg, and D. Porter (2024): "Reserve price preferences and auction design," *Journal of Behavioral Finance*, 1–7.

Kusunoki, S. (2007): "Japan's government procurement regimes for public works: A comparative introduction," *Brooklyn Journal of International Law*, 32.

Lander, D. and K. Li (2026): "On the existence of equilibria with entry and trade in second-price auctions with a secret reserve price," *Economic Theory Bulletin*, 14, 6.

Lechner, M. and A. Strittmatter (2019): "Practical procedures to deal with common support problems in matching estimation," *Econometric Reviews*, 38, 193–207.

Li, H. and G. Tan (2017): "Hidden reserve prices with risk-averse bidders," *Frontiers of Economics in China*, 12, 341–370.

Li, T. and I. Perrigne (2003): "Timber sale auctions with random reserve prices," *Review of Economics and Statistics*, 85, 189–200.

Lovo, S. and C. Spaenjers (2017): "No-trade in second-price auctions with entry costs and secret reserve prices," *Economics Letters*, 156, 142–144.

Malani, A. and J. Reif (2015): "Interpreting pre-trends as anticipation: Impact on estimated treatment effects from tort reform," *Journal of Public Economics*, 124, 1–17.

Moreno, D. and J. Wooders (2017): "Reserve prices in auctions with entry when the seller is risk-averse," *Economics Letters*, 154, 6–9.

Ohashi, H. (2009): "Effects of transparency in procurement practices on government expenditure: A case study of municipal public works," *Review of Industrial Organization*, 34, 267–285.

Ohno, T. and Y. Harada (2006): "A comparison of tendering and contracting systems for public works between Japan, the United States and EU countries," *Government Auditing Review*, 13, 49–71.

Rosar, F. (2014): "Secret reserve prices in first-price auctions," *International Journal of Industrial Organization*, 37, 65–74.

Rosenkranz, S. and P. W. Schmitz (2007): "Reserve prices in auctions as reference points," *The Economic Journal*, 117, 637–653.

Roth, J., P. H. Sant'Anna, A. Bilinski, and J. Poe (2023): "What's trending in difference-in-differences? A synthesis of the recent econometrics literature," *Journal of Econometrics*, 235, 2218–2244.

SEKI CITY GOVERNMENT (2022): "Yotei kakaku jigo kohyo no shiko kekka ni tsuite," Seki: Contract Review Division.

VINCENT, D. R. (1995): "Bidding off the wall: Why reserve prices may be kept secret," *Journal of Economic Theory*, 65, 575–584.

WALLEY, M. J. AND D. R. FORTIN (2005): "Behavioral outcomes from online auctions: reserve price, reserve disclosure, and initial bidding influences in the decision process," *Journal of Business Research*, 58, 1409–1418.

YOKOHAMA CITY GOVERNMENT (2011): "Yotei kakaku no jigo kohyo no shiko no jokyo to ni tsuite," Yokohama: Contract Division 1.

APPENDIX

This appendix contains supplementary materials for *'Reserve-Price Disclosure in Public Procurement Auctions: Evidence from Japan'*. It has the following structure:

A. Data and sample information

- Figure A1: Map of Japanese municipalities by price-cap disclosure
- Table A1: List of treatment municipalities and treatment timing
- Table A2: Summary statistics of key variables (including/excluding Shoo Town)
- Table A3: Baseline balance on observables
- Table A4: Baseline correlates of adoption
- Table A5: Baseline cohort composition
- Figure A2: Baseline covariates under reweighting and trimming

B. Additional results and discussion

- Figure B1: Event-study estimates of dynamic treatment effects (all periods)
- Table B1: Event-study estimates of dynamic treatment effects (all periods)
- Figure B2: Dynamic treatment effects on civil-engineering expenditure (all periods)
- Figure B3: Event-study estimates including Shoo Town
- Figure B4: Leave-one-cohort-out analysis
- Figure B5: Leave-one-region-out analysis
- Figure B6: Event-study estimates with each baseline control
- Figure B7: Event-study estimates using BJS estimator
- Figure B8: Event-study estimates with IPW
- Figure B9: Event-study estimates under alternative sample restrictions

Appendix A Data and summary statistics

A.1 Geographic distribution of treatment and control municipalities

Figure A1: Map of Japanese municipalities by price-cap disclosure

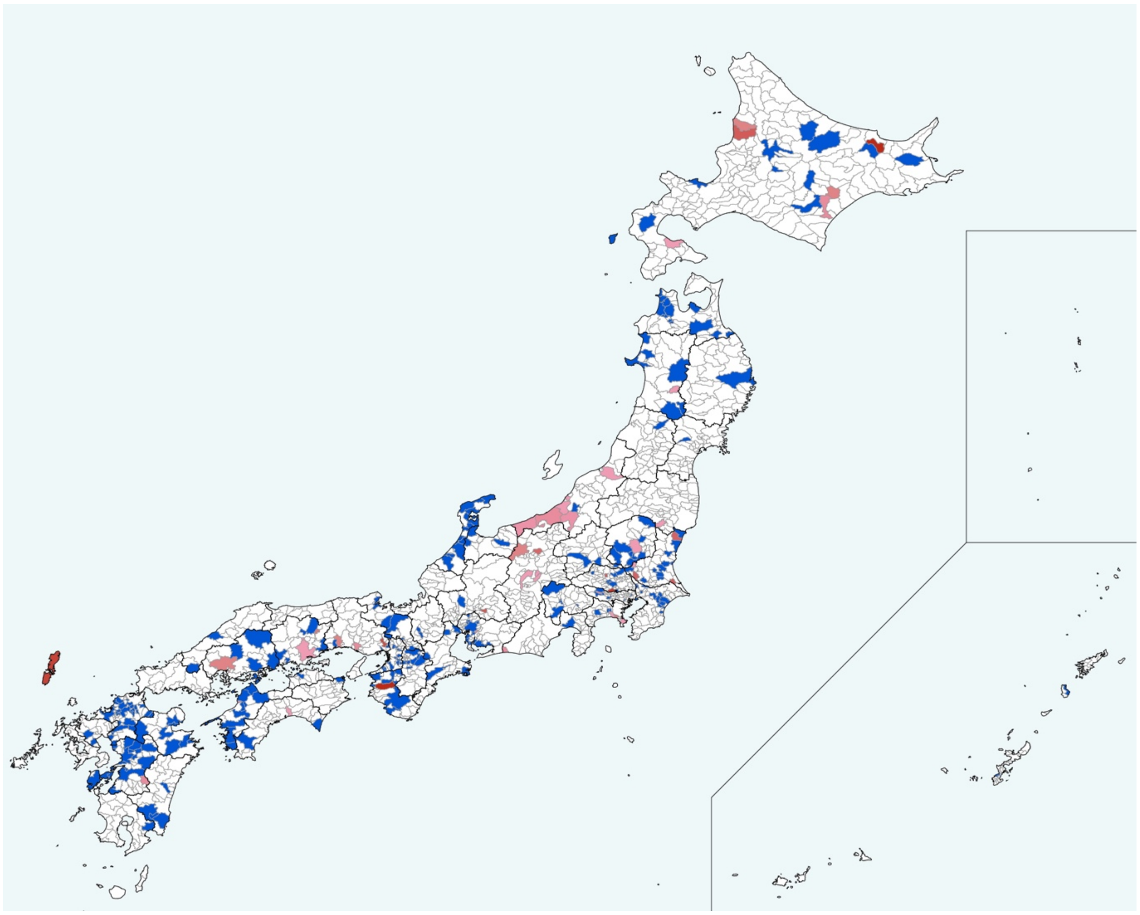

Note: The blue areas represent control municipalities that used ex-ante disclosure throughout the sample period, while the red areas represent treatment municipalities (lighter/darker shades indicate earlier/later adoption).
Source: Survey on Implementation Situation based on the Act for Promoting Proper Tendering and Contracting for Public Works, Ministry of Land, Infrastructure, Transport and Tourism; Ministry of Internal Affairs and Communications; and Ministry of Finance (2023).

A.2 Additional details on treatment municipalities

Table A1: List of treatment municipalities and treatment timing

Year	Municipalities	Total
2009	Fujisawa City, Hamura City, Kashiwazaki City, Nankoku City, Okayama City, Shibata City, Shiojiri City, Suwa City, Utsunomiya City, Yokosuka City, Yuki City, Misato Town, Mori Town, Tanagura Town, Shimosuwa Town	15
2010	Itoigawa City, Tokamachi City, Masaki Town, Hiezu Village	4
2011	Joetsu City, Kakogawa City, Kosai City, Toyoake City, Hinode Town, Makubetsu Town, Mizukami Village	7
2012	Bizen City, Fuji City, Ikeda City, Shimanto City	4
2013	Hiroshima City, Omachi City, Tatsuno City, Ikeda Town, Tomamae Town	5
2014	Kamakura City, Shoo Town	2
2015	Chikuma City	1
2016	Bando City, Takahagi City	2
2017	Itako City, Mitake Town, Namegawa Town, Obira Town	4
2018	Takarazuka City	1
2019	Tsushima City	1
2020	–	0
2021	–	0
2022	Aridagawa Town, Ozora Town	2
2023	Tokorozawa City	1
Total	–	49

Table A2: Summary statistics of key variables (including/excluding Shoo Town)

	Treatment group (Including Shoo Town)	**Treatment group** (Excluding Shoo Town)
Average winning bid rate (%)	92.37 (3.70)	92.39 (3.73)
Population (thousand persons)	125.1 (213.6)	127.5 (215.2)
Labor force (thousand persons)	61.0 (105.2)	62.2 (106.0)
Unemployment rate (%)	4.6 (1.0)	4.6 (1.0)
Area (km^2)	276.6 (261.8)	281.2 (262.6)
Habitable land area (km^2)	102.2 (104.6)	103.7 (105.2)
Standard fiscal scale (billion JPY)	29.6 (51.3)	30.1 (51.7)
Fiscal strength index	0.62 (0.27)	0.62 (0.27)
Civil-engineering expenditure (billion JPY)	6.9 (15.1)	7.1 (15.3)
Number of construction firms	515.7 (852.1)	525.3 (858.4)
Number of municipalities	49	48
Total observations	784	768

Note: Each column reports group means and standard deviations (in parentheses).
Source: Average winning bid rate: Survey on Implementation Situation based on the Act for Promoting Proper Tendering and Contracting for Public Works, Ministry of Land, Infrastructure, Transport and Tourism; Ministry of Internal Affairs and Communications; and Ministry of Finance (years: FY2008–FY2023); population: Basic Resident Registration, Ministry of Internal Affairs and Communications (years: 2008–2023); labor force population and unemployment rate: Population Census of Japan, Ministry of Internal Affairs and Communications (averaged from 2010, 2015 and 2020); area and habitable land area: The Report of Statistical Reports on the Land Area by Prefectures and Municipalities in Japan, Ministry of Land, Infrastructure, Transport and Tourism (years: 2008–2023); standard fiscal scale, fiscal strength index, and civil-engineering expenditure: Municipal Settlement Status Survey, Ministry of Internal Affairs and Communications (years: FY2008–FY2023); number of construction firms: Economic Census, Ministry of Internal Affairs and Communications (averaged from 2011, 2014, 2016 and 2021).

Table A3: Baseline balance on observables

	(1) Control	(2) Treated	(3) Difference	(4) p-value
Average winning bid rate (%)	89.78	89.01	0.77	0.413
Log population	10.79	10.79	−0.01	0.976
Log population density	1.55	1.16	0.39	0.153
Log standard fiscal scale	16.31	16.40	−0.09	0.617
Fiscal strength index	0.63	0.68	−0.05	0.263

Note: Columns (1) and (2) report baseline means for control municipalities and municipalities treated during the sample period (i.e., by 2023), respectively. Column (3) reports the difference in means (control minus treated). Column (4) reports the p-value from a two-sample t-test with unequal variances. All covariates are measured in the baseline year (2008).

Table A4: Baseline correlates of adoption

	(1) Baseline covariates	(2) + Baseline AWBR	(3) + Expenditure
Log population	−0.128	−0.135	−0.141
	(0.116)	(0.116)	(0.116)
Log population density	−0.061***	−0.065***	−0.062***
	(0.0219)	(0.0220)	(0.0218)
Log standard fiscal scale	0.156	0.157	0.091
	(0.119)	(0.119)	(0.140)
Fiscal strength index	0.375***	0.388***	0.348***
	(0.111)	(0.111)	(0.118)
Average winning bid rate (%)		−0.006	−0.006
		(0.004)	(0.004)
Log civil-engineering expenditure			0.065
			(0.056)
Observations	291	291	291
R-squared	0.057	0.063	0.067
p-value: joint test	0.0006	0.0006	0.0007

Note: All columns report linear probability models, where the dependent variable is an indicator equal to one if a municipality adopts ex-post disclosure by 2023. All covariates are measured in the baseline year (2008). Column (1) includes baseline covariates only. Column (2) additionally includes baseline average winning bid rate (AWBR). Column (3) further adds baseline log civil-engineering expenditure. The bottom row reports the p-value for a joint test that all reported slope coefficients in a given column are equal to zero. Robust standard errors are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A5: Baseline cohort composition

	(1) Early adopters	(2) Late adopters	(3) Never treated
Average winning bid rate (%)	88.60	89.47	89.78
	(6.18)	(5.86)	(5.61)
Log population	10.93	10.64	10.79
	(1.39)	(1.43)	(1.09)
Log population density	1.42	0.87	1.55
	(1.55)	(1.93)	(1.57)
Log standard fiscal scale	16.46	16.33	16.31
	(1.22)	(1.15)	(0.92)
Fiscal strength index	0.73	0.62	0.63
	(0.29)	(0.32)	(0.29)

Note: The table reports baseline means, with standard deviations in parentheses. All covariates are measured in the baseline year (2008). "Early adopters" are treated municipalities that switched to ex-post disclosure in a year less than or equal to the median adoption year among treated municipalities; "late adopters" are treated municipalities that switched in a year above the median; "never treated" municipalities are those that do not adopt ex-post disclosure during the sample period.

Figure A2: Baseline covariates under reweighting and trimming

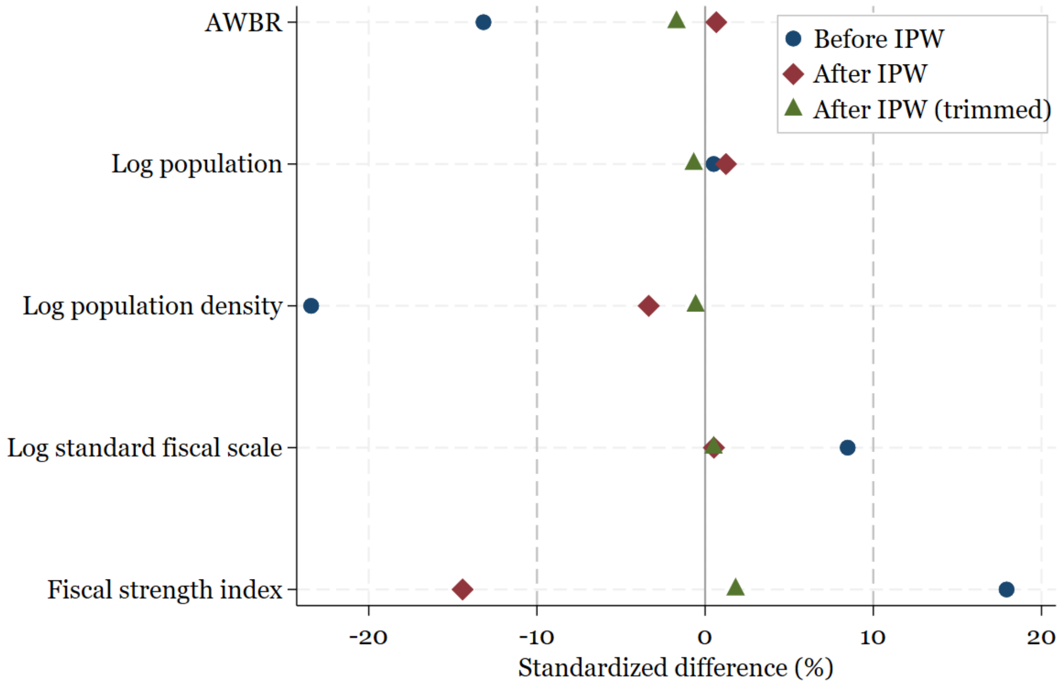

Note: Each point reports the standardized difference in means between treated and control municipalities, computed as the difference in group means divided by the pooled standard deviation of the corresponding covariate. All covariates are measured in the baseline year (2008). "After IPW" uses inverse-probability weights based on a logit model of (in-sample) treatment adoption as a function of baseline covariates: AWBR, log population, log population density, log standard fiscal scale, and fiscal strength index. "After IPW (trimmed)" additionally excludes control observations with ATT weights above the 99th percentile. Smaller absolute standardized differences indicate better covariate balance.

Appendix B Additional results and discussion

B.1 Main estimates

Figure B1: Event-study estimates of dynamic treatment effects (all periods)

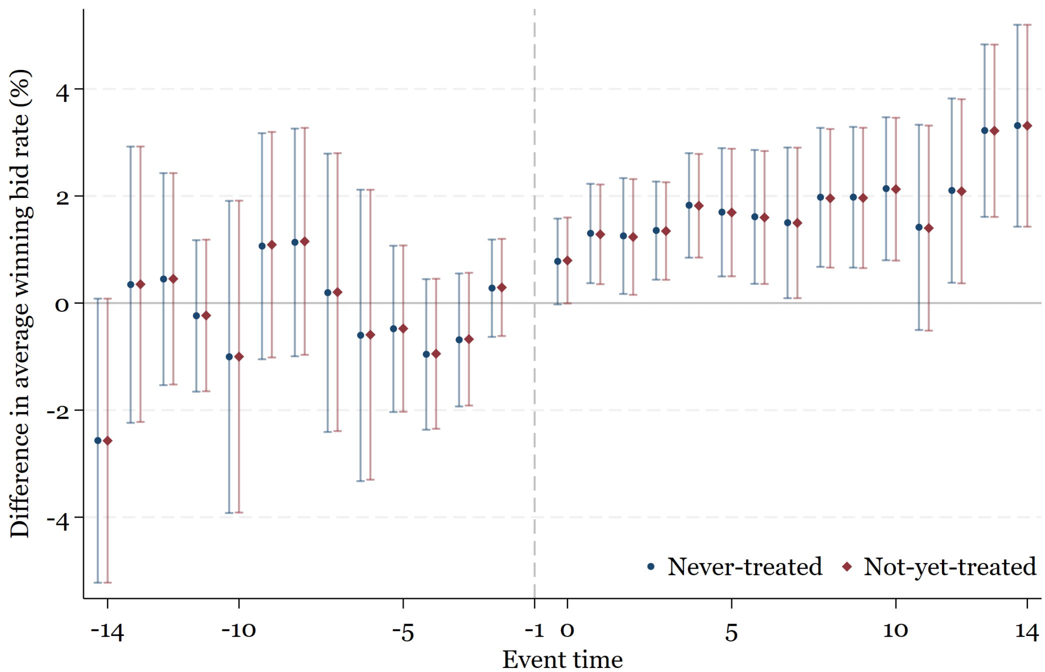

Note: This figure reports estimates using (i) only never-treated units as the control ("Never-treated"), and (ii) both never-treated and not-yet-treated units ("Not-yet-treated"). Shoo Town is excluded from the sample. The treatment units switch from ex-ante to ex-post at event time 0. The value at time -1 is set to zero since the other coefficients are measured relative to time -1 as the baseline period. Standard errors are clustered at the municipality level, and confidence intervals are set at the 95% level.

Table B1: Event-study estimates of dynamic treatment effects (all periods)

Pre-treatment			**Post-treatment**		
Time	Treated units	Coefficient estimate	Time	Treated units	Coefficient estimate
−15	1	−0.624 [−1.199, −0.050]	0	48	0.776 [−0.025, 1.577]
−14	3	−2.571 [−5.224, 0.082]	1	47	1.300 [0.372, 2.228]
−13	3	0.343 [−2.237, 2.922]	2	45	1.252 [0.170, 2.334]
−12	3	0.447 [−1.533, 2.427]	3	45	1.353 [0.437, 2.270]
−11	4	−0.240 [−1.655, 1.175]	4	45	1.824 [0.848, 2.800]
−10	5	−1.006 [−3.921, 1.909]	5	44	1.695 [0.497, 2.894]
−9	9	1.062 [−1.049, 3.173]	6	43	1.609 [0.359, 2.859]
−8	11	1.132 [−0.994, 3.259]	7	39	1.498 [0.091, 2.906]
−7	12	0.191 [−2.408, 2.791]	8	37	1.975 [0.676, 3.273]
−6	13	−0.604 [−3.325, 2.117]	9	36	1.976 [0.662, 3.290]
−5	18	−0.483 [−2.035, 1.070]	10	35	2.135 [0.800, 3.471]
−4	22	−0.959 [−2.365, 0.447]	11	30	1.414 [−0.503, 3.331]
−3	29	−0.689 [−1.931, 0.553]	12	26	2.100 [0.379, 3.821]
−2	33	0.277 [−0.632, 1.186]	13	19	3.220 [1.610, 4.831]
−1	48	– –	14	15	3.312 [1.427, 5.198]
Average:		−0.266 [−1.414, 0.882]			1.829 [0.780, 2.879]

Note: These estimates correspond to those presented in Figure B1 (“Never-treated”, i.e., only using never-treated units as the control group). They also correspond to Figure 2 (“Baseline”), expanded to the full event-time window. Shoo Town is excluded from the sample. 95% confidence intervals are shown in brackets.

Figure B2: Dynamic treatment effects on civil-engineering expenditure (all periods)

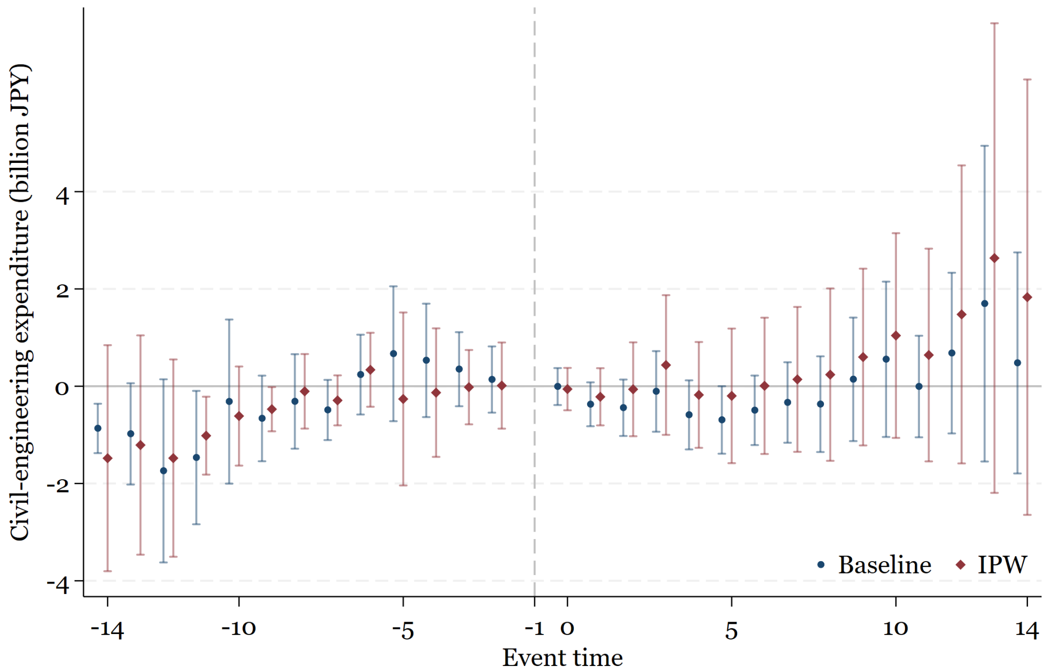

Note: This figure reports the estimates corresponding to Figure 3 over the full event-time window. Shoo Town is excluded from the sample. The treatment units switch from ex-ante to ex-post at event time 0. The value at time −1 is set to zero since the other coefficients are measured relative to time −1 as the baseline period. Standard errors are clustered at the municipality level, and confidence intervals are set at the 95% level.

B.2 Sensitivity analysis

Figure B3: Event-study estimates including Shoo Town

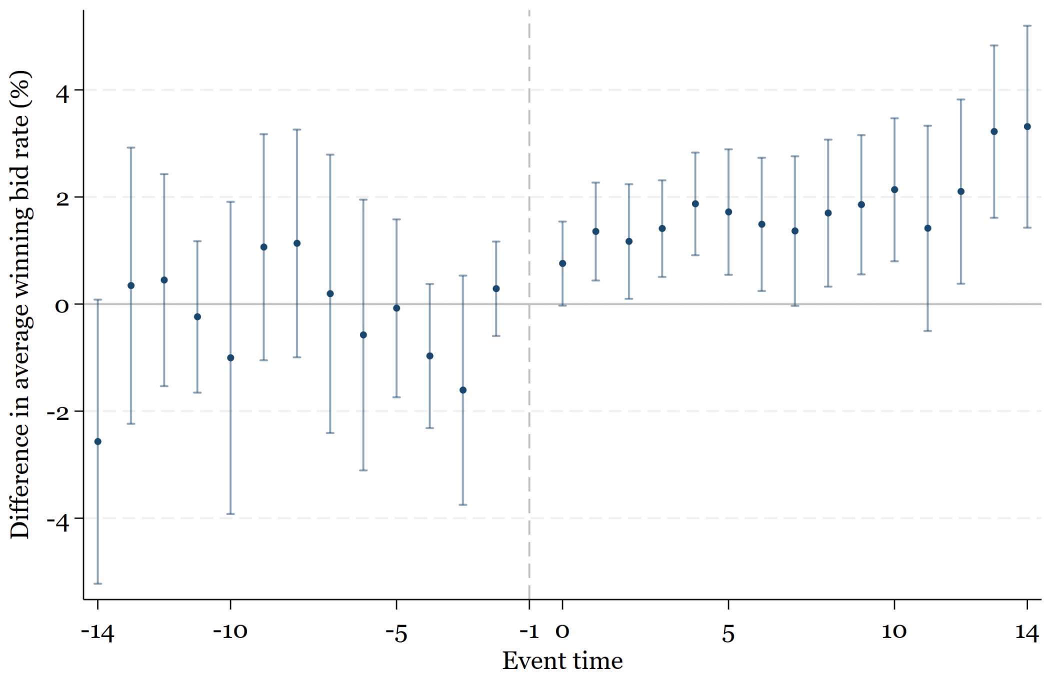

Note: This exercise examines the sensitivity of the main results to including Shoo Town, which is excluded from Figure 2. The treatment units switch from ex-ante to ex-post at event time 0. The value at time -1 is set to zero since the other coefficients are measured relative to time -1 as the baseline period. Standard errors are clustered at the municipality level, and confidence intervals are set at the 95% level.

Figure B4: Leave-one-cohort-out analysis

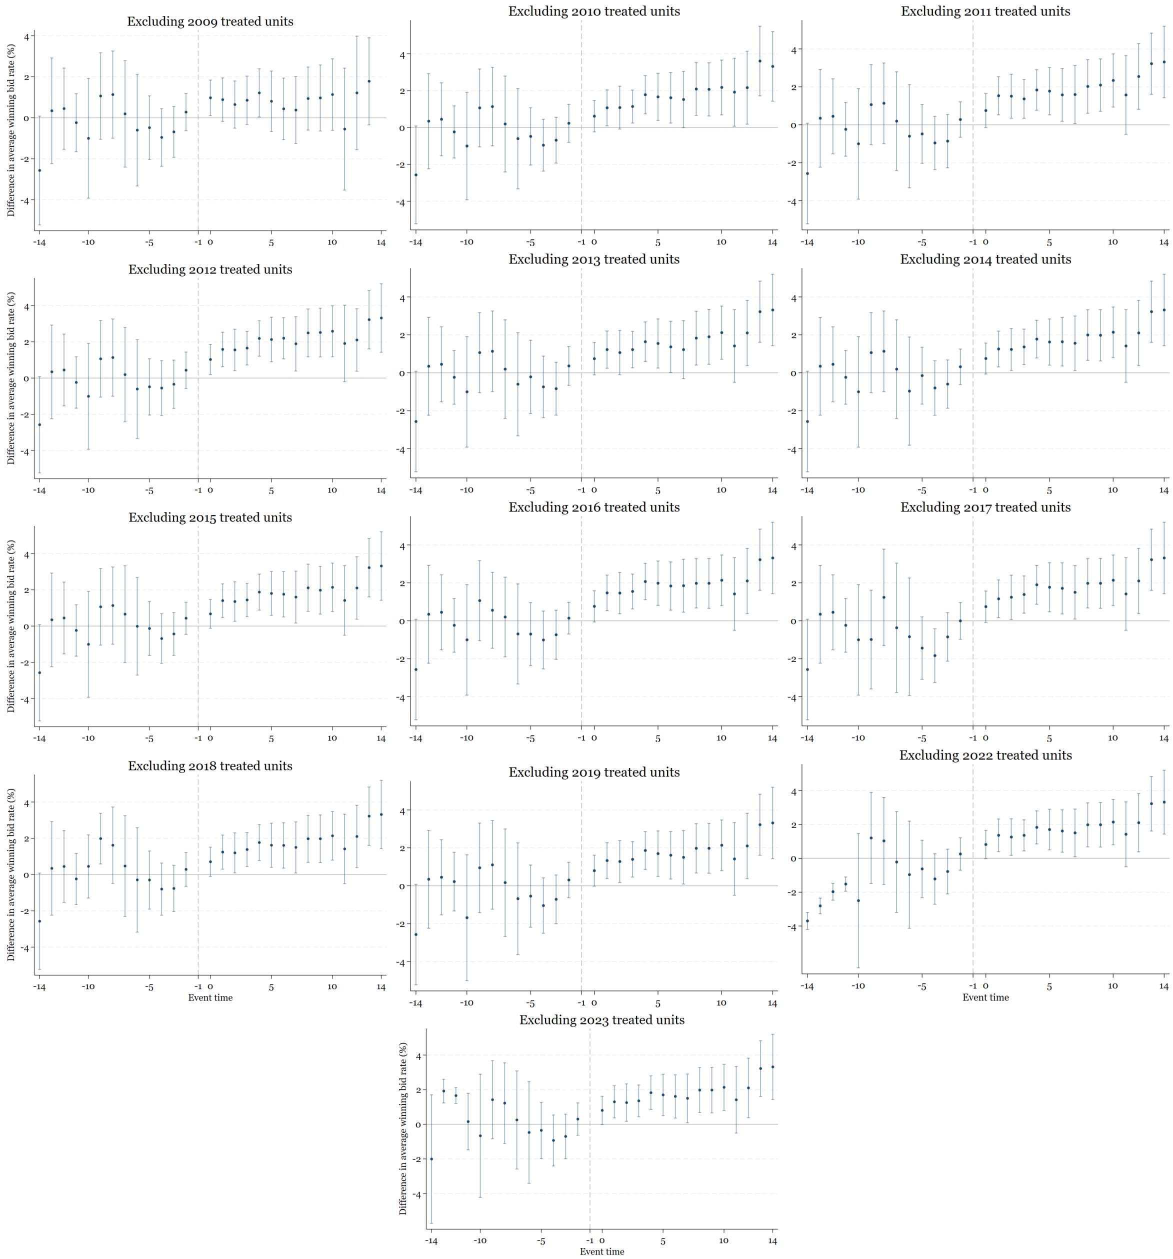

Note: This exercise tests the sensitivity of our main results with respect to the year of adoption, by reporting estimates excluding all municipalities treated in the year listed (in each panel). The treatment units switch from ex-ante to ex-post at event time 0. The value at time -1 is set to zero since the other coefficients are measured relative to time -1 as the baseline period. Standard errors are clustered at the municipality level, and confidence intervals are set at the 95% level.

Figure B5: Leave-one-region-out analysis

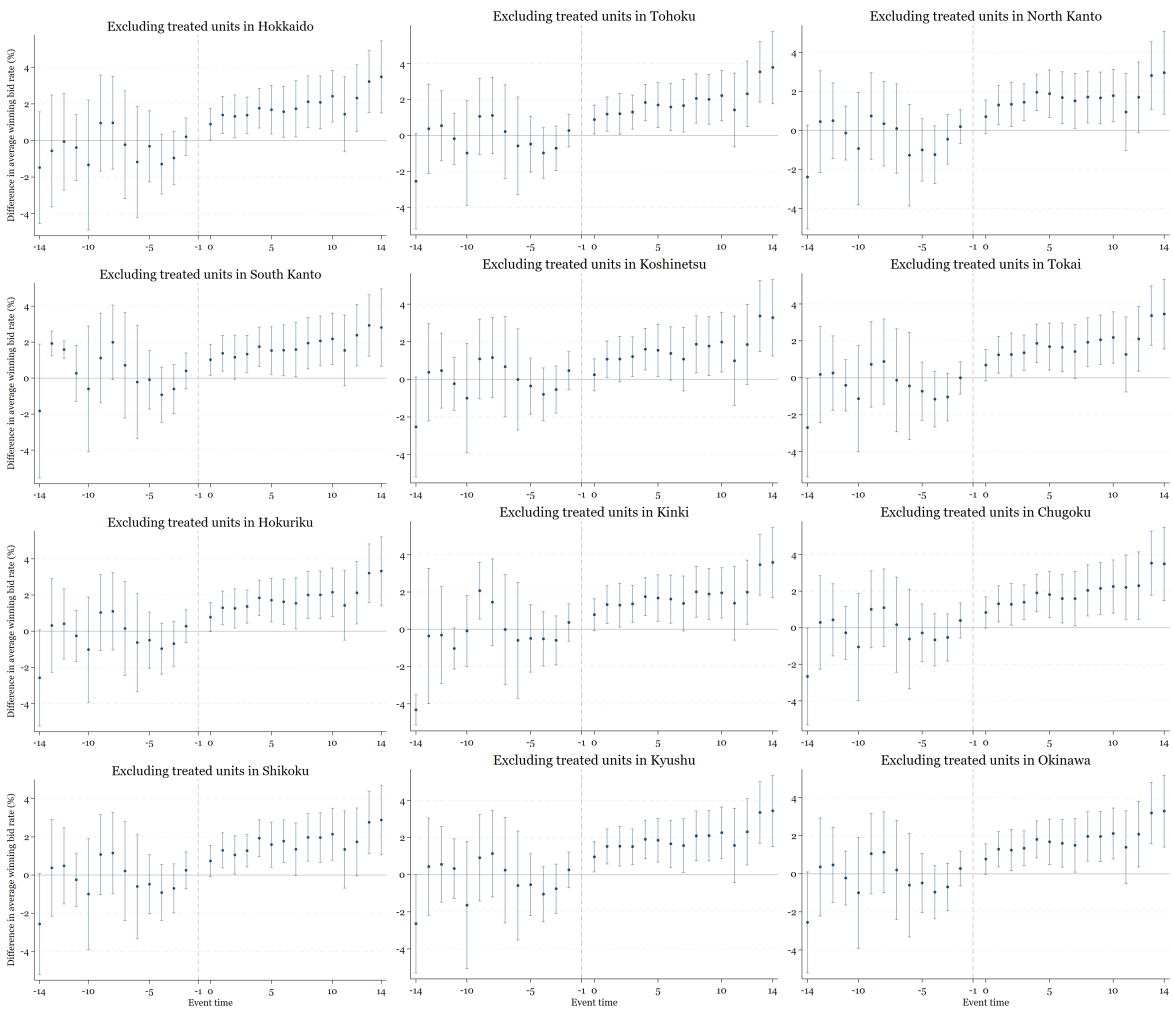

Note: This exercise tests the sensitivity of our main results with respect to geographic region, by reporting estimates excluding all treated municipalities from a specific region (in each panel). The treatment units switch from ex-ante to ex-post at event time 0. The value at time −1 is set to zero since the other coefficients are measured relative to time −1 as the baseline period. Standard errors are clustered at the municipality level, and confidence intervals are set at the 95% level.

B.3 Alternative specifications

B.3.1 Sensitivity to baseline controls

Figure B6: Event-study estimates with each baseline control

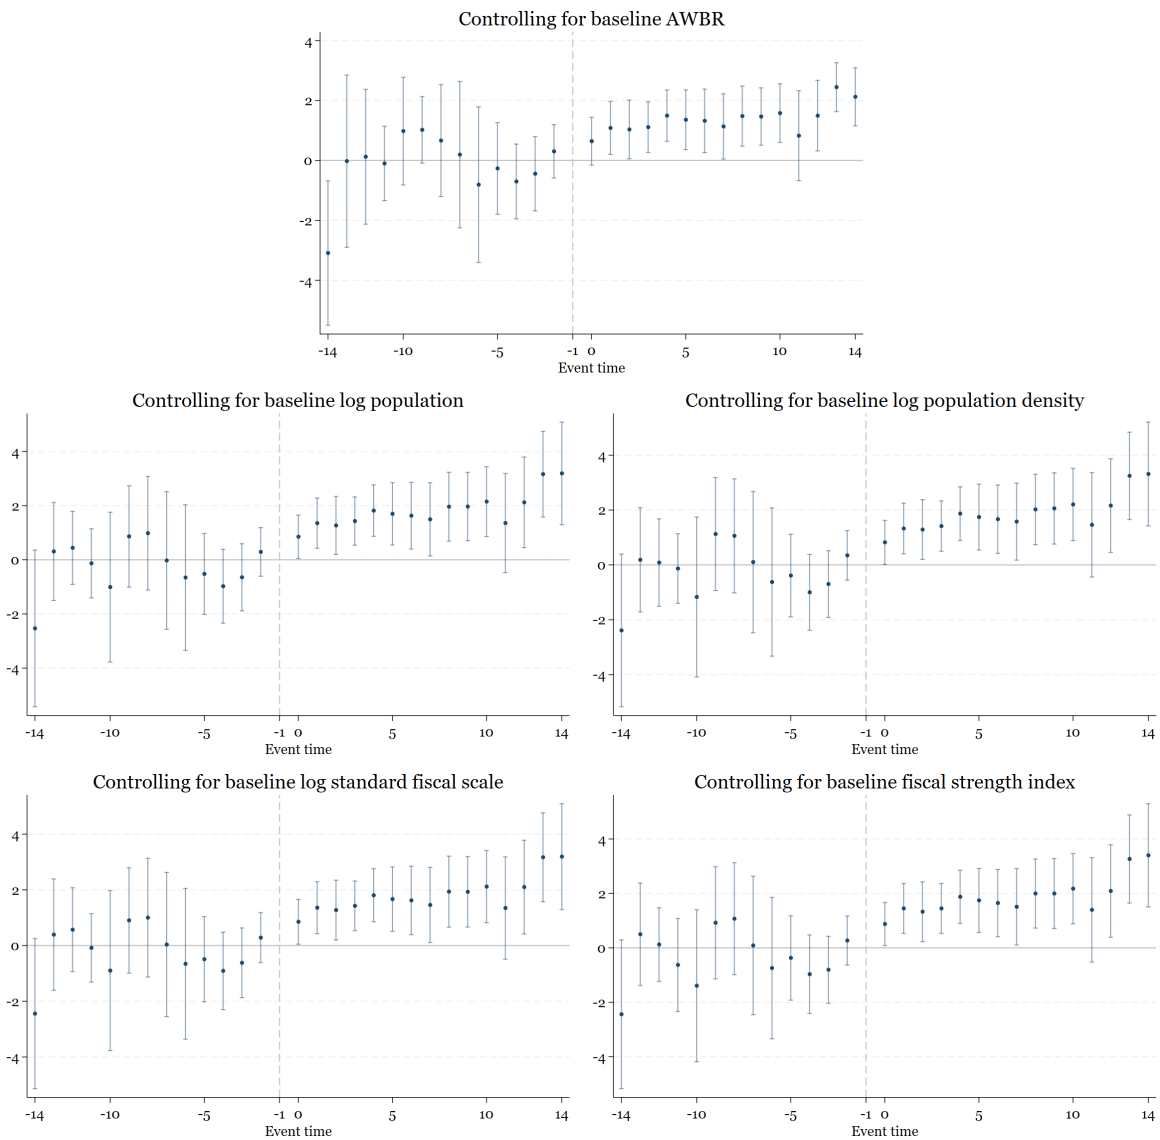

Note: Each panel of this figure reports estimates including one baseline control (i.e., AWBR, log population, log population density, log standard fiscal scale, and fiscal strength index). Shoo Town is excluded from the sample. The treatment units switch from ex-ante to ex-post at event time 0. The value at time -1 is set to zero since the other coefficients are measured relative to time -1 as the baseline period. Standard errors are clustered at the municipality level, and confidence intervals are set at the 95% level.

B.3.2 Estimates using Borusyak-Jaravel-Spiess (BJS) estimator

Figure B7: Event-study estimates using BJS estimator

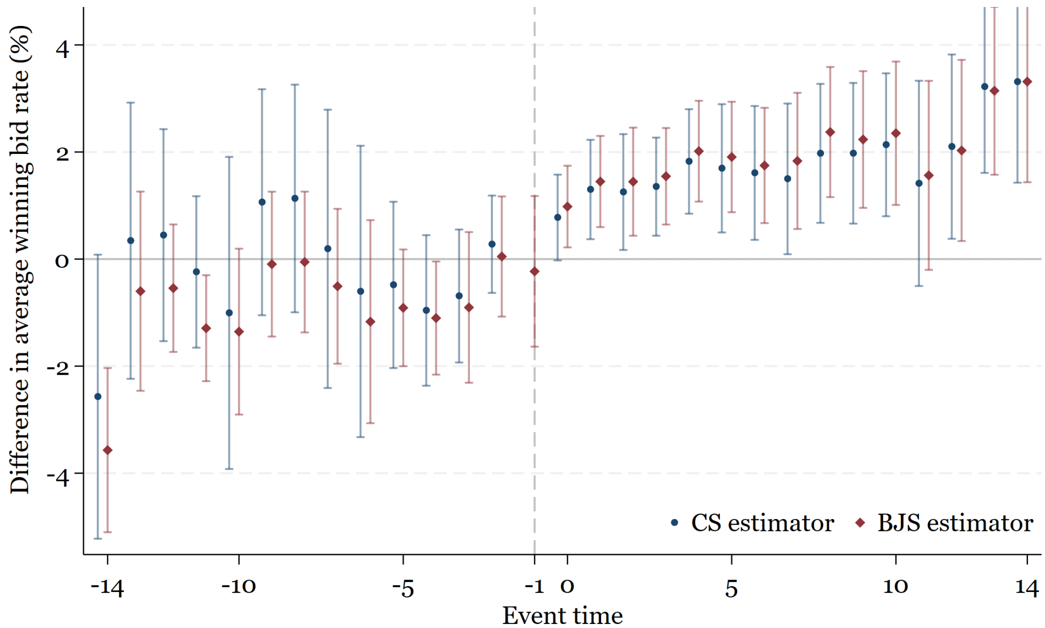

Note: This figure reports estimates using the Borusyak et al. (2024) estimator, which we also refer to as the BJS estimator, together with our baseline estimates using the Callaway and Sant'Anna (2021) estimator. Shoo Town is excluded from the sample. The treatment units switch from ex-ante to ex-post at event time 0. The value at time -1 is set to zero since the other coefficients are measured relative to time -1 as the baseline period. Standard errors are clustered at the municipality level, and confidence intervals are set at the 95% level.

B.3.3 Estimates with inverse-probability weighting (IPW)

Figure B8: Event-study estimates with IPW

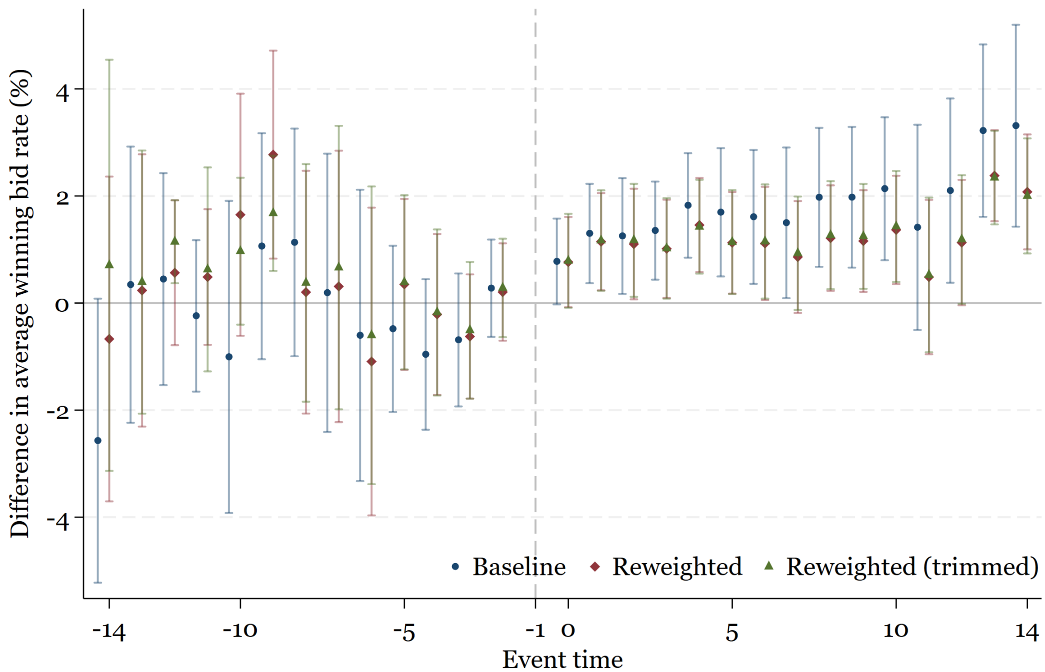

Note: The "Baseline" specification uses the full estimation sample without additional covariates. The "Reweighted" specification re-estimates the event study using inverse-probability weighting based on baseline municipality characteristics (AWBR, log population, log population density, log standard fiscal scale, and fiscal strength index). The "Reweighted (trimmed)" specification additionally excludes control municipalities with ATT weights above the 99th percentile. Shoo Town is excluded from the sample. The treatment units switch from ex-ante to ex-post at event time 0. The value at time -1 is set to zero since the other coefficients are measured relative to time -1 as the baseline period. Standard errors are clustered at the municipality level, and confidence intervals are set at the 95% level.

B.3.4 Estimates under alternative sample restrictions

Figure B9: Event-study estimates under alternative sample restrictions

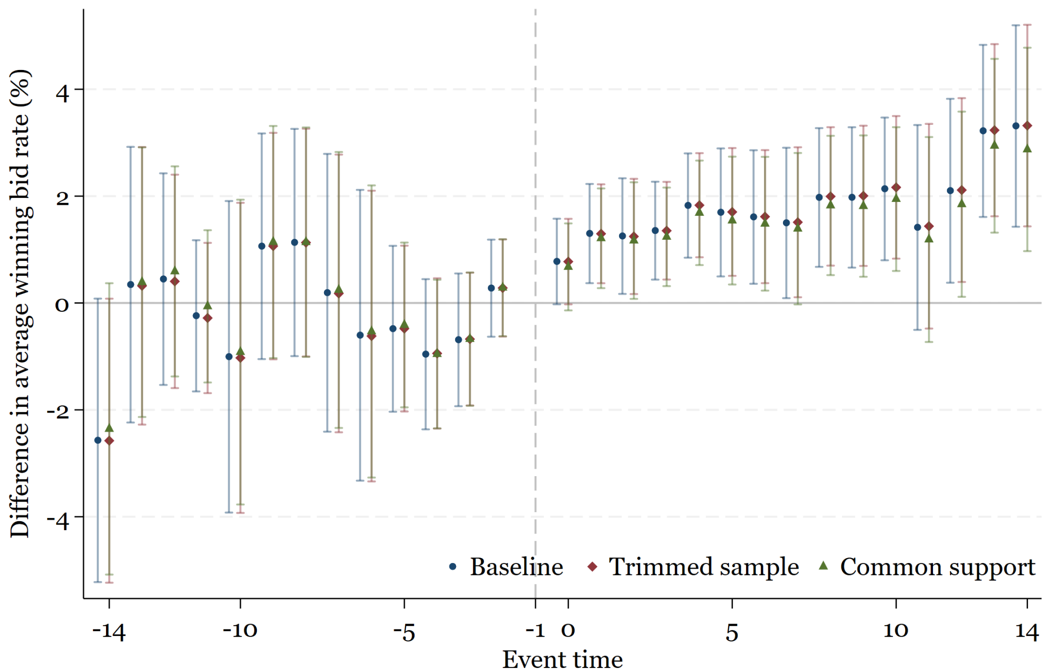

Note: The "Baseline" specification uses the full estimation sample, excluding Shoo Town. The "Trimmed sample" specification excludes control municipalities with ATT weights above the 99th percentile. The "Common support" specification restricts the control group to municipalities whose estimated baseline propensity scores lie within the support of the treated group. The treatment units switch from ex-ante to ex-post at event time 0. The value at time −1 is set to zero since the other coefficients are measured relative to time −1 as the baseline period. Standard errors are clustered at the municipality level, and confidence intervals are set at the 95% level.