

The AI Procurement Boom: State Demand After ChatGPT

Kolbe A. Dumas

Department of Political Science, University of Houston

Working paper, October 2026. Not peer reviewed.

Abstract

Purpose

This article examines whether ChatGPT 3.5’s release was associated with increased state demand for artificial intelligence (AI) capabilities. Procurement solicitations make this demand observable before agencies contract, deploy, or evaluate those tools.

Design/methodology/approach

I analyze 267 AI-related procurement opportunities across all 50 U.S. states from January 2017 through May 2026, using a dataset constructed from public procurement records. An interrupted time-series design estimates changes in pooled monthly solicitation counts, with a linear Poisson model, a quadratic comparison, and state fixed-effects checks.

Findings

Solicitation activity increased following the release, with stronger evidence for faster growth than an immediate jump. In the primary model, estimated monthly growth increased from approximately 1.8 to 4.0 percent. The size and shape of this increase depend on the underlying trend specification, particularly when states are allowed their own trends. The pattern is consistent with bounded learning as public buyers formally explore capabilities while evidence about their administrative performance remains incomplete.

Originality

This article introduces an original, hand-collected dataset of state AI procurement records that have not previously been examined together. The dataset allows procurement demand to be tracked before and after ChatGPT’s release.

Keywords: artificial intelligence; public procurement; bounded learning; state government; technological salience.

1 Introduction

Public agencies begin moving toward artificial intelligence (AI) through procurement requests that define needs and invite vendors to propose capabilities. These requests make administrative interest observable before agencies execute contracts, deploy systems, or evaluate their performance. Since the release of ChatGPT 3.5, promises of speed, cost efficiency, process improvement, and improved well-being have made AI more attractive, while questions of data privacy, vendor opacity, misuse, and systemic bias complicate its evaluation. Existing research identifies several incentives behind bureaucratic interest in AI, including fascination with new technology, efficiency gains, process improvement, and career advancement (Andersson et al., 2025). The question for public procurement is how buyers begin formally exploring these tools while reliable evidence about their administrative consequences remains limited.

In principle, procurement decisions about technically opaque tools would benefit from extensive learning through peer implementation and feedback. During the pandemic, for example, Chen et al. (2024) find that the National Association of State Workforce Agencies (NASWA) convened weekly meetings where agencies shared operational challenges and coping strategies; chatbot use to relieve caseloads was one recurring topic. Such exchanges offer agencies information about how other public organizations respond to similar pressures. Agencies also face incentives to act while that information remains incomplete. When they fail to recalibrate under rapidly changing conditions, administrative reputation, funding, and political support can erode (Brodin, 2021).

This article examines state AI procurement following ChatGPT through a bounded learning framework, in which public buyers rely on partial, noisy, vendor-filtered, and time-constrained signals. Such information can make a tool appear useful or feasible before systematic evidence from peer implementation is available. Solicitations provide an observable record of how agencies act within that information environment; they allow public buyers to search for capabilities, scope potential uses, and seek information from suppliers before committing to implementation.

The release of ChatGPT 3.5 provides a setting for examining how heightened technological visibility enters procurement. Earlier research on e-procurement documented variation in how U.S. states adopted tools to organize purchasing, with technically and legally complex innovations spreading more slowly (Moon, 2005). Connecting to that, public buyers encounter new technologies within established organizational and procedural constraints. Greater awareness can expand the capabilities they consider, while procurement provides a formal route for exploring their possible use.

The release of ChatGPT 3.5 stands out as a technological salience shock because it lowered barriers to experimentation and brought AI into everyday professional use. Surveys indicate that generative AI spread at work as rapidly as the personal computer, while overall adoption outpaced both personal computers and the early internet (Bick et al., 2024). Its perceived potential for automation, information access, and efficiency gains made AI easier to consider for administrative tasks. For public buyers, greater familiarity could make these capabilities appear feasible and worth exploring through procurement, even while uncertainty about their performance remained.

Did this increased visibility coincide with greater formal demand for AI capabilities across U.S. state governments? Research on bureaucratic careers suggests that administrative behavior can reflect incentives for professional advancement alongside organizational performance (Brösamle, 2012). In interviews with U.S. city employees, Johnson et al. (2024) report one participant describing departments that want to “buy the cool thing,” improve operations, and secure promotions. Efficiency goals, administrative reputation, career considerations, and modernization expectations can thus create incentives to explore AI through procurement. Under these pressures, public buyers may issue solicitations while evidence about the tools’ administrative consequences is still developing.

To examine this question, I analyze over 200 AI-related procurement opportunities across all 50 U.S. states from January 2017 through May 2026, using a dataset constructed from public procurement records. An interrupted time-series (ITS) design estimates changes in monthly solicitation activity following ChatGPT’s release, accounting for the underlying time trend and recurring seasonality. Because the release was national and individual state events remain sparse, the primary analysis uses pooled monthly counts, with state-month models providing alternative specifications. The headline results show higher solicitation activity and faster growth after the release. The contribution to procurement research is to observe demand for an emerging technology at the solicitation stage, when public buyers begin searching for, scoping, and formally requesting capabilities before contracts and implementation outcomes become observable.

2 Theory

2.1 Technological Salience Shocks

Public buyers allocate limited attention among competing administrative needs and possible technological solutions. Salience describes how visible and relevant an innovation becomes within that decision environment. The diffusion literature connects the spread of innovations

to how governments prioritize information according to its perceived importance, salience, and urgency (Boushey, 2010). ChatGPT’s public release made experimentation possible without agency-specific development or a prior procurement process. Greater familiarity could make AI appear more feasible for administrative tasks and encourage buyers to consider capabilities they had previously overlooked. Consistent with punctuated-equilibrium accounts of government attention and institutional adjustment (Jones and Baumgartner, 2012), I expect this increase in visibility to produce a change in the trajectory of AI procurement activity following the release.

Following the broader treatment of AI in public administration research (Sun and Medaglia, 2019; Haesevoets et al., 2025), I use AI as an encompassing category that includes machine learning, natural language processing, and generative AI. Technical distinctions among these tools remain important for evaluating any particular procurement. The expectation here is that ChatGPT’s visibility can draw attention to AI capabilities beyond text generation. Public buyers who become interested in generative AI may also reconsider existing tools for prediction, classification, or language processing. I therefore treat ChatGPT as a salience shock to AI broadly and examine whether solicitation activity across that category increased following its release.

2.2 Procurement as Exploratory Administrative Demand

Procurement solicitations provide a concrete way to observe how public buyers begin pursuing AI capabilities. These documents translate administrative interest into requests that suppliers can respond to, whether agencies have identified a specific need or are still exploring possible applications. The solicitations in this dataset range from New Hampshire’s Judicial Branch chatbot project to California’s request for generative AI solutions for employment projections and recession forecasting. Georgia’s request for an AI solicitation evaluation tool extends this interest to the procurement process itself. These examples show how public buyers pursue AI across different administrative functions; solicitations record the capabilities agencies seek and the problems they want vendors to help address.

As such, these instruments need to be distinguished from adoption and implementation. A solicitation establishes that an agency has formally sought information or proposals; it does not establish that a contract was signed, a system was deployed, or administrative performance improved. Budget limits, internal approvals, and changing priorities can shape whether a request is issued and whether it proceeds further. This paper therefore treats solicitation activity as an indicator of formal administrative demand. Public buyers can use procurement to clarify needs, assess supplier offerings, and explore feasible solutions while

the decision to acquire or deploy a tool remains unsettled. Connecting solicitations to executed contracts and subsequent use would allow future research to examine which requests lead to implementation.

What drives public buyers to issue solicitations for AI capabilities? Some seek efficiency gains or relief from routine administrative work; others aim to improve service quality or respond to modernization pressures. Red tape can consume organizational resources and make alternatives to existing procedures more attractive (Bozeman, 1993). Agencies may also seek tools that make services easier to navigate, especially where administrative burdens weigh heavily on applicants with limited cognitive resources (Christensen et al., 2020). These expectations connect to what Bozeman and Youtie (2020) call robotic bureaucracy, where automated systems manage administrative and compliance tasks. Their analysis also cautions that automation can shift work onto those interacting with the system. Public buyers may therefore pursue AI because they expect it to reduce costs or burdens, while the distribution and extent of those benefits remain uncertain.

Organizational relationships and the placement of technical expertise can also shape which possibilities reach procurement. Network research shows that relationships supporting routine organizational change can differ from those enabling departures from the institutional status quo (Battilana and Casciaro, 2012). In public sector AI research, Selten and Klievink (2024) find that separate data science teams can strengthen technical expertise, while integrating those teams into operational departments can improve alignment with existing administrative work. These arrangements may affect which needs agencies identify and which capabilities public buyers consider worth pursuing. Most procurement decisions likely reflect a blend of these organizational conditions and individual motivations. Under bounded rationality, bureaucrats seek workable choices while facing incomplete information, competing priorities, and limits on their ability to evaluate alternatives (Simon, 1991). Greater AI salience can thus encourage formal exploration even when buyers are still clarifying what they need and how a proposed tool would fit their work.

2.3 Bounded Learning and Bureaucratic Choice Under Uncertainty

In the public policy diffusion literature, “learning” refers to a process in which policymakers assess outcomes from prior adoption, including whether an innovation produced measurable gains or maintained political support (Shipan and Volden, 2008). Learning about these outcomes takes time, especially when a tool’s operation is proprietary and difficult to evaluate outside the organization. Correctional Offender Management Profiling for Alternative

Sanctions (COMPAS), a tool used to estimate recidivism risk, illustrates the problem. A 2012 New York study reported favorable predictive performance for its probation population (Lansing, 2012). In 2016, ProPublica reported racial differences in prediction errors among defendants in Broward County, Florida (Larson et al., 2016). These studies concern different populations and settings; favorable validation in one context leaves questions about performance and consequences elsewhere. Public buyers considering comparable tools therefore face an evidence problem that extends beyond whether an early evaluation reports acceptable predictive accuracy.

This gap between apparent promise and fuller evidence is where bounded learning becomes relevant. Drawing from work linking bounded rationality and diffusion, Weyland (2005) argues that bold innovations can attract disproportionate attention and spread because they appear promising before their success has been demonstrated. Public administration research similarly emphasizes the cognitive and organizational constraints shaping bureaucratic decisions (Battaglio et al., 2019; Schwarz et al., 2022). I use bounded learning to describe how public buyers evaluate AI through incomplete information, vendor claims, and signals available within limited decision timeframes. Greater visibility can make these signals easier to encounter and more compelling without making them more reliable. Buyers may consequently begin procurement exploration while evidence about administrative performance is still developing. Evidence from peer implementation can enter this process as it develops, while limits on its availability, relevance, and interpretation continue to shape decisions.

Public buyers weigh expected efficiency and service improvements alongside reputational concerns, career incentives, and demands from political principals. These incentives operate within agencies that must also manage competing deadlines, program requirements, budget limits, and legal responsibilities. Research on public sector AI identifies organizational resources, technical expertise, and regulatory conditions as factors shaping whether agencies can pursue these possibilities (Madan and Ashok, 2023; Selten and Klievink, 2024). Bureaucrats seek to optimize under uncertainty, incomplete information, and institutional constraint. A tool that appears administratively useful may still carry costs or risks that discourage procurement; conversely, pressure to improve services may encourage exploration before its performance is fully understood. Bounded learning concerns how buyers assess these tradeoffs with the information and resources available to them.

Vendors can further limit the information available to public buyers during this assessment. In interviews with city employees, Johnson et al. (2024) document instances where vendors declined to disclose model information, invoked proprietary protections, or deflected questions about system performance and training data. Participants also described limited

opportunities to customize tools that suppliers intended to sell as standard products. These experiences illustrate how buyers can seek information through procurement while remaining dependent on what suppliers are willing to provide. Mulligan and Bamberger (2019) similarly argue that procurement can leave consequential system design choices largely in vendors’ hands when agencies lack independent expertise. For buyers facing limited staff capacity and competing responsibilities, vendor explanations may become practical shortcuts for evaluating possible solutions. Procurement can therefore support learning while leaving buyers with an incomplete basis for judging whether a proposed tool will meet their administrative needs.

Planning for Medicaid work requirements under the 2025 reconciliation law illustrates how new policy demands can make automation administratively attractive. KFF notes that states must process income-based Medicaid applications within 45 days; states processing larger shares within 30 days likely have greater automation, while slower processing may reflect manual review or backlogs (KFF, 2026). Additional verification requirements may place further pressure on these processes. Agencies facing high caseloads and limited staffing may consequently seek automated tools to handle new responsibilities within existing administrative timelines. These pressures can make technological exploration a practical response to service demands, even when agencies have limited capacity to evaluate available tools. Automation here may involve conventional data matching; whether buyers pursue AI capabilities depends on the needs they identify and the solutions they consider.

Procurement gives public buyers a way to act on these pressures while learning about possible solutions. Solicitations allow agencies to clarify needs, request supplier information, and assess proposed capabilities before committing to implementation. However, reliance on suppliers for both technical knowledge and potential solutions can reinforce the information and power asymmetries associated with government dependence on private vendors (Zhang et al., 2026). As AI becomes more visible and appears more feasible, buyers may become more willing to begin this process using the limited evidence available to them. Increased solicitation activity can thus emerge while agencies are still learning what tools could address their needs and how those tools might perform. Evidence from peer implementation may strengthen subsequent decisions, but formal exploration can begin before that evidence becomes sufficiently developed or relevant to the buyer’s circumstances.

2.4 Hypothesis

The preceding argument suggests that greater AI salience should make public buyers more willing to translate administrative needs into procurement requests. Buyers can begin this

exploration while evidence about performance remains incomplete, using available information to assess whether further investigation is worthwhile. Because solicitation preparation requires time and internal coordination, the response may develop gradually as agencies clarify needs and navigate review.

H1 (Procurement Demand). The release of ChatGPT 3.5 is associated with an increase in U.S. state demand for AI capabilities, measured through monthly counts of AI-related procurement solicitations.

I evaluate this expectation using an interrupted time-series design that estimates changes in solicitation levels and subsequent growth relative to the projected pre-release trajectory. These are different ways the expected increase may appear; the hypothesis requires neither an immediate jump nor a uniform administrative lag. Increased solicitation activity would be consistent with buyers pursuing AI under bounded learning. Establishing how buyers learned or why they issued particular requests would require additional evidence from their decision processes.

3 Methods

3.1 Data and Measurement

I constructed a dataset of 267 AI-related procurement opportunities issued by state agencies across all 50 U.S. states between January 2017 and May 2026. The analysis excludes procurement opportunities issued by university campuses and local governments. The dependent variable is the monthly count of AI-related procurement opportunities issued by each state. Data collection involved scraping state procurement portals, agency archives, and publicly available procurement records.

Where information was lacking, targeted agency inquiries and public records requests sought missing documents and clarification of release dates. Available notices, solicitation documents, attachments, and amendments were used to establish each opportunity’s scope and timing. Historical retrieval varies across states because portals differ in retention, indexing, and access to closed solicitations.

Searches used terms including artificial intelligence, machine learning, deep learning, natural language processing, large language models, generative AI, and chatbot variants. Eligibility depended on the capabilities described in procurement notices and associated buyer documents. AI could be a central requirement, an optional feature, or a supporting capability within a broader request. This approach captures agencies exploring AI alongside

other solutions, including requests where AI appears in the scope rather than the solicitation title. Generic uses of intelligence and conventional analytics or chatbot references were assessed for substantive AI relevance.

The counting unit is a distinct procurement opportunity, assigned to its earliest verified release month. Linked stages of the same opportunity, such as an RFI followed by an RFP, are counted once. Amendments, addenda, and reposts remain associated with that opportunity unless they establish a separate procurement request. Repeated purchases with distinct scope can contribute additional observations, while uncertain later increments are excluded from the primary series and examined in a sensitivity check. Canceled solicitations remain included because the agency had already issued a formal request. I assign release timing using the issuance or publication month recorded in procurement documents. Response deadlines and contract award dates alone are insufficient to establish when a request became public.

3.2 Time Window and Aggregation

The analysis begins in 2017, the first year with recorded opportunities in the dataset, and ends in May 2026. Starting in January retains the full calendar year in which those first requests appeared. Alternative specifications begin in 2016 and 2018 to assess sensitivity to the starting point. December 2022 is the first full month following ChatGPT’s late-November release, leaving 71 months before the breakpoint and 42 months afterward.

The resulting panel contains 5,650 state-month observations, approximately 95.6 percent of which record no AI-related procurement opportunity. Solicitation activity is sparse because agencies issue these requests intermittently, with substantial variation across states. The primary analysis sums state counts into a monthly national series to examine whether procurement demand across state governments changed after ChatGPT. State-month fixed-effects models retain the underlying geographic variation and assess sensitivity to persistent differences in procurement activity and state-specific trends.

3.3 Model Specification

I estimate a Poisson interrupted time-series model of the pooled monthly solicitation count. The primary specification includes a linear time trend, calendar-month fixed effects, and two terms capturing changes starting with the first month after ChatGPT 3.5’s release, December 2022. The level term estimates whether expected solicitation activity changed at the breakpoint; the growth term estimates whether activity subsequently grew faster or slower than the underlying trend. Because the model uses a log link, the linear trend

represents a constant proportional growth rate rather than the same number of additional requests each month. I estimate all coefficients jointly using observations from the full analysis period.

An alternative specification adds squared time since December 2022 to allow the subsequent growth rate to change over time. This term is zero before the breakpoint, preserving the linear baseline while allowing growth after release to accelerate or moderate. This accommodates procurement demand building as AI attracts greater attention and agencies prepare requests, while also permitting growth to slow as the period progresses. Both specifications include recurring seasonal controls and use confidence intervals adjusted for unequal variance and serial dependence in monthly counts.

3.4 Identification and Interpretation

The ITS design assumes that, without ChatGPT’s release, solicitation activity would have continued along the underlying trend after accounting for recurring seasonality. This continuation provides the baseline against which changes beginning in December 2022 are assessed. Because the release was national, the analysis has no untreated comparison group. Changes in vendor offerings, policy demands, or the availability of procurement records could also influence the observed pattern. I therefore interpret the estimates as changes associated with the release, conditional on the trend specification and reasonably comparable data coverage over time.

I assess sensitivity to these assumptions using alternative starting dates, counting rules, and state-month specifications. Placebo tests apply the models to earlier dates using only observations before December 2022, while the event-time figure shows how solicitation activity differs across periods surrounding the release. These checks assess whether the findings depend on the underlying trend, a small number of high-activity states, or decisions about which opportunities to count. The observed changes can be consistent with public buyers responding under bounded learning, although identifying their motivations and information sources would require evidence from agency decision processes.

3.5 AI Assistance

I used OpenAI’s ChatGPT and Codex to assist with search planning, procurement portal navigation, document retrieval, text extraction, and screening of eligibility, dates, and linked procurement records. These tools also supported preparation of records requests, dataset construction code, statistical programming, checks of model estimates, and data visualization. Manuscript assistance included organization, language revision, and formatting of the

text, tables, and figures. I set the inclusion and counting rules, made the recorded manual decisions, and retained responsibility for the analysis and interpretation.

Earlier screening used GPT-5 to flag borderline procurement records for review. Later reconstruction sessions used models recorded under the user-facing labels 5.6 sol, 6 sol, 6.1 sol, and 6 astra. Exact backend versions and model assignments for individual records were not consistently retained. Assistance occurred through iterative research sessions rather than a standardized classifier applied uniformly across records. The preserved source evidence, decision records, and construction and analysis code document how the final counts and estimates were produced.

4 Results

4.1 Main Estimates

AI-related solicitation activity increased substantially following ChatGPT’s release, particularly during the later years of the observation period. Figure 1 plots a trailing three-month average to show the pattern without individual monthly fluctuations dominating the series. Recorded opportunities increased from 25 in 2023 to 57 in 2024 and 75 in 2025, with another 39 during the first five months of 2026. Across the full observation period, the monthly average was approximately 0.99 before December 2022 and 4.69 afterward. The models assess whether this increase exceeded the growth already occurring before release.

Table I reports the linear and quadratic specifications using incidence rate ratios (IRRs). Values above one indicate higher expected counts, or faster growth when applied to a trend term. The primary linear model estimates monthly growth of approximately 4 percent after release, compared with about 1.8 percent along the underlying trend. The estimated change in monthly growth is positive, with an IRR of 1.022 and a 95 percent confidence interval of [1.004, 1.040]. The level estimate is considerably less precise, with an IRR of 1.076 and an interval of [0.617, 1.874]. The primary model supports faster growth in procurement demand following release, with limited evidence of a discrete jump in December 2022.


Allowing growth to change over time produces a pattern of faster initial growth followed by moderation. In the quadratic specification, the initial change in monthly growth has an IRR of 1.070, with a 95 percent confidence interval of [1.023, 1.120]. Combined with the underlying trend, this implies initial growth of approximately 9 percent per month. The estimated growth rate subsequently declines, although the confidence interval for this slow-down only narrowly excludes zero. This specification suggests that procurement demand

Table I: Main Interrupted Time-Series Estimates

	Monthly AI Solicitation Counts	
	Linear	Quadratic
<i>Underlying Trend</i>		
Baseline Monthly Growth	1.018** [1.005, 1.031]	1.018** [1.006, 1.031]
<i>Changes After Release</i>		
Level Change	1.076 [0.617, 1.874]	0.710 [0.368, 1.373]
Change in Monthly Growth	1.022* [1.004, 1.040]	1.070** [1.023, 1.120]
Change in Growth Over Time	—	-0.146* [-0.289, -0.003]
<i>Model Fit and Sample</i>		
AIC	378.4	377.7
BIC	419.3	421.3
Observations	113	113
Solicitations	267	267

Note: IRRs with 95% CIs; squared-time estimates use the log scale per year squared, and quadratic growth refers to the breakpoint. Poisson with month effects and serial-correlation-robust uncertainty. * $p < .05$, ** $p < .01$, *** $p < .001$.

Figure 1: AI Procurement Solicitation Activity



Note: Trailing three-month mean of observed counts; the guide marks December 2022. The series ends in May 2026.

grew more quickly early in the period following release, with the pace easing in later years.


Figure 2 examines the timing more directly by comparing solicitation activity across periods surrounding the release. Estimates are measured against the six preceding months, June through November 2022, while accounting for the underlying trend and recurring seasonality. The early post-release estimates are positive but imprecise, with clearer differences emerging from the second year onward. The estimated difference reaches approximately 2.65 additional monthly solicitations during months 12–23 and 5.08 during months 24–35. These larger differences in later periods indicate that increased solicitation activity persisted well beyond the initial release.

4.2 Sensitivity Checks

Beginning the analysis in 2018, 2019, or 2020 preserves the positive growth estimate in the linear model, with confidence intervals remaining above one. The quadratic specification also retains faster initial growth across these alternative windows. Evidence of subsequent moderation becomes less precise when the earliest years are excluded.

With a common underlying trend, state fixed-effects models reproduce the pooled estimates while accounting for persistent differences in procurement activity. Allowing each

Figure 2: Solicitation Activity Around ChatGPT's Release



Note: Count-scale bin differences relative to the six pre-release months, adjusted for trend and month effects. Bars show 95% confidence intervals with serial-correlation-robust uncertainty.

state its own underlying trend, however, brings the linear growth-change estimate close to one, with an IRR of 0.999 and a 95 percent confidence interval of [0.982, 1.016]. The evidence for a constant increase in growth is therefore sensitive to how existing state trajectories are modeled. Under the same check, the quadratic specification retains faster initial growth, with an IRR of 1.078 and an interval of [1.018, 1.142], followed by moderation over time.

The growth estimate changes little under alternative counting rules and coverage restrictions. These checks include counting uncertain additional opportunities, treating separately dated releases as additional observations, and excluding records with timing cautions or outstanding document recovery. Across the checks, the linear growth IRR ranges from 1.020 to 1.024, close to the primary estimate of 1.022. Confidence intervals remain above one for both the linear growth change and the quadratic model’s initial growth change.

Removing the five states with the highest solicitation totals leaves 182 opportunities and a positive linear growth change, with an IRR of 1.031 and a 95 percent confidence interval of [1.007, 1.056]. The quadratic model’s initial growth estimate remains positive but becomes less precise, with an IRR of 1.045 and an interval of [0.988, 1.107]. Evidence of subsequent moderation also weakens when these states are excluded. Faster growth remains detectable beyond the highest-activity states, although the smaller sample provides less certainty about how that growth changed over time.

Placebo evidence varies across the two trend specifications, as reported in the appendix. Using only observations before December 2022, the linear model detects declines in growth at artificial breaks in December 2020 and December 2021. The March 2022 placebo shows no clear change in growth under that specification. Joint tests of the quadratic model’s growth terms show no statistically detectable change at any of these three dates. The linear placebo results warrant caution when projecting a constant growth rate from the period before release.

An additional diagnostic, reported in the appendix, tests how well monthly counts distinguish periods before and after the release. Separation at December 2022 is stronger than at any tested pre-release date and remains statistically detectable under tests that reshuffle individual months or blocks of months. After accounting for the search across earlier dates, the strongest pre-release separation is not statistically detectable. This test concerns recorded activity, which can also change because of underlying growth or differences in record coverage.

5 Discussion

Public buyers can use procurement to respond to emerging technologies while decisions about acquisition and implementation remain open. The rise in AI-related solicitations documents increased formal demand; its association with ChatGPT is consistent with greater technological visibility encouraging agencies to explore available capabilities. ChatGPT could make AI appear useful or feasible for administrative tasks without resolving questions about how particular systems would perform. The solicitation stage is where that interest can become a formal request for supplier information or proposed solutions.

Bounded learning helps explain how buyers can pursue these possibilities while working with incomplete evidence. Supplier responses may clarify available capabilities, but buyers remain dependent on what vendors disclose about their systems and proposed uses (Mulligan and Bamberger, 2019). This dependence can be especially consequential in sole-source procurement, where vendor information may help justify pursuing a particular supplier without competing proposals. Georgia’s 2024 procurement of Acusensus’s Heads-Up enforcement system provides an example. In a letter attached to the Department of Public Safety’s sole-source notice, Acusensus claimed to be the only supplier offering the combined imaging, AI detection, and real-time enforcement capabilities for identifying mobile phone and seat belt violations.¹ The vendor supplied information about the technology alongside a claim about why competing suppliers could not meet the same requirements. Buyers must evaluate both claims with the technical knowledge and market information available to them. Evidence from peer implementation can add information as experience accumulates, although another agency’s results may not establish whether a tool will meet the buyer’s own needs. The findings are consistent with procurement exploration under these constraints. Establishing which information buyers relied upon, or what motivated their requests, would require agency decision records and interviews alongside the solicitation counts.

Several motivations could produce the increased procurement demand observed here. Public buyers may seek efficiency gains, relief from administrative work, or improvements in service delivery (Bozeman and Youtie, 2020; Chen et al., 2024). Career incentives and pressure to demonstrate modernization may also encourage officials to pursue emerging technologies (Brösamle, 2012). These motivations can operate together, with the same request reflecting both an organizational need and an opportunity for professional advancement. The solicitation counts provide evidence of increased formal activity, but distinguishing these motivations requires closer examination of how agencies identified needs and decided to pursue

¹Acusensus Inc., sole-source letter to the Georgia Department of Public Safety, March 18, 2024, attached to Georgia Procurement Registry event PE-46600-SS-2024-000000007.

particular capabilities. Bounded learning describes the conditions under which buyers assess these possibilities, allowing for different incentives across agencies.

For procurement research, the findings bring attention to demand formation before acquisition and implementation. Agencies issuing requests begin defining the problems they want technology to address and the capabilities they expect suppliers to provide. These choices can shape subsequent procurement even while buyers are still learning what available systems can do. Examining this stage makes it possible to study how technological visibility enters public purchasing decisions and how supplier information influences the options buyers consider. It also places procurement within AI governance from the outset, when agencies are defining requirements and deciding what evidence to request from vendors.

The ITS estimates depend on how procurement demand would have evolved without ChatGPT’s release. Without an untreated comparison group, changes in vendor offerings, policy demands, or record availability could also contribute to the observed increase. The linear growth increase disappears when states are allowed their own underlying trends, although the quadratic specification retains faster initial growth. Earlier placebo results also vary across these specifications. The increase in recorded demand is clear, but its magnitude relative to the projected baseline depends on how existing procurement trajectories are modeled. These qualifications limit how confidently the estimated change can be attributed specifically to ChatGPT.

Historical coverage also varies because state portals differ in how they retain and provide access to procurement records. Targeted inquiries and records requests helped recover missing information, but incomplete archives remain a limitation. Solicitations capture agencies pursuing capabilities, with some requests later canceled or never converted into contracts. Purchases through existing agreements or below public posting thresholds may also occur without a separate solicitation appearing in the dataset. These channels leave part of government AI acquisition outside the activity examined here. Connecting requests to contracts and agency decision records would provide a fuller account of which capabilities buyers acquired, how they evaluated vendor claims, and whether the resulting systems entered administrative use.

6 Conclusion

This article examined state demand for AI capabilities following the release of ChatGPT 3.5. Using 267 procurement opportunities across all 50 states from 2017 through May 2026, the analysis documents higher solicitation activity and faster growth after release. The estimated change depends on how existing procurement trends are modeled, and the evidence

is stronger for subsequent growth than an immediate jump. Bounded learning provides an interpretation of this response, as public buyers can pursue promising capabilities while information about their usefulness remains incomplete. Procurement makes this exploration observable through the requests agencies issue, before those requests become contracts or systems enter administrative use.

The purpose of this analysis was limited to examining how overall procurement activity responded to the emergence of a newly salient technology. In this case, ChatGPT’s release was followed by increased state demand for AI capabilities. The focus is on the observable change in procurement activity and how that change developed over time. Vendor dependence and the governance challenges discussed above make this response consequential, while closer examination of individual procurements is needed to understand how buyers navigated them. Further research can examine which types of AI agencies seek, the channels they use to acquire those capabilities, and the safeguards they require from suppliers.

Future research should follow these opportunities through contract execution and implementation, examining how buyers assessed supplier claims and whether acquired tools met their needs. Agency decision records and interviews would help explain which information shaped procurement choices and how organizational pressures influenced them. Municipal governments provide another setting for examining these decisions under different service demands and purchasing constraints. For procurement practice, the findings draw attention to the stage when agencies define needs, assess available capabilities, and decide what evidence suppliers must provide. AI governance needs to examine those decisions while requirements remain open and buyers can still question the evidence supporting a proposed purchase.

Acknowledgments

The author thanks Ling Zhu and Michael Kistner for substantial comments on earlier versions of this manuscript. The author also thanks participants in the Southern Political Science Association conference-within-a-conference on large language models in political science, as well as faculty and graduate students at the University of Houston Graduate Student Research Series, for comments on earlier presentations of this work.

References

Andersson, P. E., Arbin, K. and Rosenqvist, C. (2025), ‘Assessing the value of artificial intelligence (AI) in governmental public procurement’, *Journal of Public Procurement*

25(1), 120–139.

- Battaglio, R. P., Belardinelli, P., Bellé, N. and Cantarelli, P. (2019), ‘Behavioral Public Administration *ad fontes*: A Synthesis of Research on Bounded Rationality, Cognitive Biases, and Nudging in Public Organizations’, *Public Administration Review* **79**(3), 304–320.
- Battilana, J. and Casciaro, T. (2012), ‘Change Agents, Networks, and Institutions: A Contingency Theory of Organizational Change’, *Academy of Management Journal* **55**(2), 381–398.
- Bick, A., Blandin, A. and Deming, D. J. (2024), *The Rapid Adoption of Generative AI*, number no. w32966 in ‘NBER working paper series’, National Bureau of Economic Research, Cambridge, Mass.
- Boushey, G. (2010), *Policy Diffusion Dynamics in America*, Cambridge University Press.
- Bozeman, B. (1993), ‘A Theory of Government “Red Tape”’, *Journal of Public Administration Research and Theory: J-PART* **3**(3), 273–303.
- Bozeman, B. and Youtie, J. (2020), ‘Robotic Bureaucracy: Administrative Burden and Red Tape in University Research’, *Public Administration Review* **80**(1), 157–162.
- Brodkin, E. Z. (2021), ‘Street-Level Organizations at the Front Lines of Crises’, *Journal of Comparative Policy Analysis: Research and Practice* **23**(1), 16–29.
- Brösamle, K. J. (2012), ‘Careering Bureaucrats and Bureaucrats’ Careers’, *German Journal of Human Resource Management* **26**(2), 195–198.
- Chen, T., Gascó-Hernandez, M. and Esteve, M. (2024), ‘The Adoption and Implementation of Artificial Intelligence Chatbots in Public Organizations: Evidence from U.S. State Governments’, *The American Review of Public Administration* **54**(3), 255–270.
- Christensen, J., Aarøe, L., Baekgaard, M., Herd, P. and Moynihan, D. P. (2020), ‘Human Capital and Administrative Burden: The Role of Cognitive Resources in Citizen-State Interactions’, *Public Administration Review* **80**(1), 127–136.
- Haesevoets, T., Verschuere, B. and Roets, A. (2025), ‘AI adoption in public administration: Perspectives of public sector managers and public sector non-managerial employees’, *Government Information Quarterly* **42**(2), 102029.

- Johnson, N., Silva, E., Leon, H., Eslami, M., Schwanke, B., Dotan, R. and Heidari, H. (2024), ‘Public Procurement for Responsible AI? Understanding U.S. Cities’ Practices, Challenges, and Needs’. arXiv:2411.04994 [cs] version: 1.
- Jones, B. D. and Baumgartner, F. R. (2012), ‘From There to Here: Punctuated Equilibrium to the General Punctuation Thesis to a Theory of Government Information Processing’, *Policy Studies Journal* **40**(1), 1–20.
- KFF (2026), ‘Tracking implementation of the 2025 reconciliation law: Medicaid work requirements, state and national data and policies’. Accessed October 1, 2026.
- Lansing, S. (2012), New york state COMPAS-probation risk and need assessment study: Examining the recidivism scale’s effectiveness and predictive accuracy, Technical report, New York State Division of Criminal Justice Services.
- Larson, J., Mattu, S., Kirchner, L. and Angwin, J. (2016), ‘How we analyzed the COMPAS recidivism algorithm’.
- Madan, R. and Ashok, M. (2023), ‘AI adoption and diffusion in public administration: A systematic literature review and future research agenda’, *Government Information Quarterly* **40**(1), 101774.
- Moon, M. J. (2005), ‘E-procurement management in state governments: Diffusion of e-procurement practices and its determinants’, *Journal of Public Procurement* **5**(1), 54–72.
- Mulligan, D. K. and Bamberger, K. A. (2019), ‘PROCUREMENT AS POLICY: ADMINISTRATIVE PROCESS FOR MACHINE LEARNING.’, *Berkeley Technology Law Journal* **34**(3), 773–851.
- Schwarz, G., Christensen, T. and Zhu, X. (2022), ‘Bounded Rationality, Satisficing, Artificial Intelligence, and Decision-Making in Public Organizations: The Contributions of Herbert Simon’, *Public Administration Review* **82**(5), 902–904.
- Selten, F. and Klievink, B. (2024), ‘Organizing public sector AI adoption: Navigating between separation and integration’, *Government Information Quarterly* **41**(1), 101885.
- Shipan, C. R. and Volden, C. (2008), ‘The Mechanisms of Policy Diffusion’, *American Journal of Political Science* **52**(4), 840–857.
- Simon, H. A. (1991), ‘Bounded Rationality and Organizational Learning’, *Organization Science* **2**(1), 125–134.

- Sun, T. Q. and Medaglia, R. (2019), ‘Mapping the challenges of Artificial Intelligence in the public sector: Evidence from public healthcare’, *Government Information Quarterly* **36**(2), 368–383.
- Weyland, K. (2005), ‘Theories of Policy Diffusion Lessons from Latin American Pension Reform’, *World Politics* **57**(2), 262–295.
- Zhang, B., Coyle, D., Kim, J. Y., Himmelreich, J. and Gascó-Hernandez, M. (2026), ‘AI, the Public Sector, and Policymaking’.

Appendix

A Model Definitions and Inference

The primary analysis uses 113 monthly observations from January 2017 through May 2026. Let Y_t denote the pooled monthly count of AI-related procurement opportunities, and let t measure months relative to December 2022. Define $D_t = 1$ beginning in December 2022 and zero beforehand. The linear and quadratic specifications share the following conditional mean:

$$\log(\mu_t) = \alpha + \beta(t/12) + \delta D_t + \gamma t D_t + \kappa(t/12)^2 D_t + s_{m(t)}, \quad (\text{A.1})$$

where $\mu_t = E(Y_t)$ and $s_{m(t)}$ represents calendar-month effects. The primary linear model sets $\kappa = 0$; the comparison estimates this term only for the period after release. Both models estimate all coefficients jointly using the full analysis window.

The level-change IRR is $\exp(\delta)$, and the growth-change IRR is $\exp(\gamma)$. The underlying trend is measured in years, so the baseline monthly growth IRR is $\exp(\beta/12)$. The squared-time estimate κ remains on the log-count scale per year squared. Its inclusion allows growth after release to change over time, so the quadratic model’s reported growth IRR is not a constant rate throughout that period. Estimates of the projected baseline set the release-related terms to zero while retaining the coefficients estimated from the full series.

Pooled models use a Newey–West covariance estimator with bandwidth three, without prewhitening and with the finite-sample adjustment. This bandwidth concerns serial dependence in the uncertainty calculation, not a delay in the procurement response. Confidence intervals use the normal approximation, with $\hat{b} \pm 1.96 \text{SE}(\hat{b})$ calculated on the coefficient scale before exponentiating terms reported as IRRs. Significance stars use two-sided normal tests based on the same covariance matrix. Table VII varies the bandwidth without changing the fitted coefficients or December 2022 breakpoint.

The pooled Pearson dispersion ratios are 1.13 for the linear model and 1.10 for the quadratic model. A twelve-month Ljung–Box test of Pearson residuals yields $p = .064$ and $p = .021$, respectively. Residual serial dependence therefore remains a concern, particularly for the quadratic specification. The robust covariance calculation addresses uncertainty under weak dependence; it does not establish that the mean specification captures every source of temporal variation.

B State Fixed Effects

Table II reports Poisson models using the 5,650 state-month observations. All specifications include state fixed effects and calendar-month effects, with uncertainty clustered by state. The common-trend specifications retain the pooled model's time terms. In this balanced panel, their release-related coefficients match the pooled coefficients, although the uncertainty calculation differs. They are an alternative way to account for persistent state differences, rather than independent confirmation from another dataset.

The remaining specifications allow each state its own underlying linear trend. Under this check, the linear growth-change estimate is 0.999 with a 95 percent confidence interval of [0.982, 1.016]. The quadratic model retains a positive initial growth change of 1.078 [1.018, 1.142], followed by moderation. Evidence of a constant increase in growth is consequently sensitive to whether states share a common underlying trajectory. State clustering permits dependence within a state, but it does not account for all possible common shocks across states.

Table II: State Fixed-Effects Estimates

	Monthly AI Solicitation Counts			
	Linear	Linear + State Trends	Quadratic	Quadratic + State Trends
<i>Changes After Release</i>				
Level Change	1.076 [0.530, 2.183]	1.198 [0.591, 2.428]	0.710 [0.279, 1.811]	0.616 [0.240, 1.582]
Change in Monthly Growth	1.022** [1.006, 1.038]	0.999 [0.982, 1.016]	1.070* [1.012, 1.132]	1.078* [1.018, 1.142]
Change in Growth Over Time	—	—	-0.146 [-0.318, 0.025]	-0.249** [-0.423, -0.075]
<i>Sample</i>				
Observations	5,650	5,650	5,650	5,650
Solicitations	267	267	267	267

Note: IRRs with 95% CIs; squared-time estimates use the log scale per year squared, and quadratic growth refers to the breakpoint. Poisson with state and month effects; state-clustered uncertainty. * $p < .05$, ** $p < .01$, *** $p < .001$.

C Analysis Windows

Table III changes the starting year while retaining May 2026 as the endpoint and December 2022 as the breakpoint. The 2018, 2019 and 2020 windows retain positive linear growth changes with intervals above one. Extending the series backward to include 2016, which has no recorded opportunities, produces an IRR of 1.015 [0.997, 1.033]. The quadratic initial growth estimate remains positive across these windows, while subsequent moderation becomes less precise under later starts. These comparisons make the influence of the early observation period visible without treating the absence of recovered records as proof that no procurement activity occurred.

Table III: Analysis Window Sensitivity

	Monthly AI Solicitation Counts			
	2016 Start	2018 Start	2019 Start	2020 Start
<i>Linear</i>				
Level Change	0.938 [0.541, 1.626]	1.196 [0.667, 2.143]	1.204 [0.623, 2.327]	1.382 [0.688, 2.776]
Change in Monthly Growth	1.015 [0.997, 1.033]	1.028** [1.008, 1.048]	1.028* [1.001, 1.055]	1.039* [1.009, 1.070]
Observations	125	101	89	77
Solicitations	267	263	253	243
<i>Quadratic</i>				
Level Change	0.617 [0.323, 1.181]	0.796 [0.397, 1.595]	0.795 [0.366, 1.725]	0.911 [0.406, 2.046]
Change in Monthly Growth	1.063** [1.017, 1.113]	1.075** [1.027, 1.126]	1.076** [1.023, 1.132]	1.088** [1.030, 1.149]
Change in Growth Over Time	-0.148* [-0.290, -0.005]	-0.143 [-0.288, 0.002]	-0.145 [-0.293, 0.002]	-0.145 [-0.297, 0.007]
Observations	125	101	89	77
Solicitations	267	263	253	243

Note: IRRs with 95% CIs; squared-time estimates use the log scale per year squared, and quadratic growth refers to the breakpoint. Poisson with month effects and serial-correlation-robust uncertainty. All windows end in May 2026. * $p < .05$, ** $p < .01$, *** $p < .001$.

D Counting, Coverage and Concentration

Table IV examines alternative counting decisions and restrictions on the retained records. Including seven uncertain additional opportunities increases the sample from 267 to 274. The separate-dated-release variant contains 280 records and counts later public releases separately, changing the unit from distinct opportunities to releases. The remaining checks exclude eight named-product cases whose identification depended on specific capability judgments, nine opportunities issued by central university systems, eight opportunities with cautions about the first version containing AI-related scope, or thirteen opportunities with outstanding fuller-document recovery. These subsets can overlap and are not separate estimates of missing procurement activity.

Campus-level university opportunities remain excluded throughout the analysis; the university-system check removes statewide system issuers retained in the primary series. Timing cautions concern whether AI-related scope was present in the earliest available version or became visible in later material. Fuller-source gaps concern incomplete recovery of supporting documents for otherwise retained opportunities. Across these variants, the linear growth-change IRR ranges from 1.020 to 1.024, and the quadratic initial growth-change IRR ranges from 1.066 to 1.076. All corresponding intervals remain above one, although these restrictions cannot remove every possible coverage difference across states or years.

Table IV: Counting and Coverage Sensitivity

		Monthly AI Solicitation Counts	
		Solicitations	Linear Quadratic
Additional Uncertain Opportunities	274	1.021*	1.070**
		[1.002, 1.040]	[1.021, 1.121]
Separate Dated Releases	280	1.020*	1.070**
		[1.001, 1.040]	[1.023, 1.119]
Exclude Named-Product Exceptions	259	1.021*	1.066**
		[1.003, 1.039]	[1.021, 1.113]
Exclude Central University Systems	258	1.024*	1.072**
		[1.005, 1.043]	[1.027, 1.120]
Exclude Timing Cautions	259	1.023*	1.075**
		[1.004, 1.043]	[1.025, 1.128]
Exclude Fuller-Source Gaps	254	1.020*	1.076**
		[1.002, 1.038]	[1.030, 1.125]

Note: Growth IRRs with 95% CIs; quadratic growth refers to the breakpoint. Serial-correlation-robust uncertainty; separate dated releases changes the counting unit. * $p < .05$, ** $p < .01$, *** $p < .001$.

Table V examines whether the findings depend on the highest-activity states. Georgia has the largest retained total, with 20 opportunities. The five highest totals are Georgia, Texas, California, New York and Massachusetts, with 85 opportunities combined. Removing those states leaves 182 opportunities and a positive linear growth change of 1.031 [1.007, 1.056]. The quadratic initial growth estimate becomes less precise, at 1.045 [0.988, 1.107], and its moderation term is also indistinguishable from zero. The remaining states retain evidence of faster growth under the linear specification, with less certainty about its shape over time.

Table V: Sensitivity to High-Activity States

	Monthly AI Solicitation Counts		
	All States	Omit Highest	Omit Highest Five
<i>Linear</i>			
Level Change	1.076 [0.617, 1.874]	1.225 [0.624, 2.403]	1.320 [0.520, 3.350]
Change in Monthly Growth	1.022* [1.004, 1.040]	1.023* [1.003, 1.043]	1.031* [1.007, 1.056]
Observations	113	113	113
Solicitations	267	247	182
<i>Quadratic</i>			
Level Change	0.710 [0.368, 1.373]	0.855 [0.404, 1.811]	1.168 [0.404, 3.379]
Change in Monthly Growth	1.070** [1.023, 1.120]	1.065** [1.016, 1.116]	1.045 [0.988, 1.107]
Change in Growth Over Time	-0.146* [-0.289, -0.003]	-0.128 [-0.282, 0.026]	-0.042 [-0.213, 0.130]
Observations	113	113	113
Solicitations	267	247	182

Note: IRRs with 95% CIs; squared-time estimates use the log scale per year squared, and quadratic growth refers to the breakpoint. Poisson with month effects and serial-correlation-robust uncertainty. * $p < .05$, ** $p < .01$, *** $p < .001$.

E Pre-Release Placebos

Table VI applies artificial breaks in December 2020, December 2021 and March 2022 using only January 2017 through November 2022 observations. The linear model detects negative growth changes at the first two dates, with $p = .047$ and $p = .048$, while the March 2022 growth change is indistinguishable from zero. Joint tests of the quadratic model’s growth and squared-time terms yield $p = .077$, $.123$ and $.848$. These are unadjusted tests at three specified dates; the null results do not establish a perfectly stable pre-release trajectory. The differences across specifications indicate that inference about earlier breaks depends on the assumed trend form.

Table VI: Pre-Release Placebo Breaks

	Monthly AI Solicitation Counts		
	December 2020	December 2021	March 2022
<i>Linear</i>			
Level Change	1.355 [0.620, 2.960]	1.502 [0.515, 4.383]	0.443 [0.153, 1.283]
Change in Monthly Growth	0.944* [0.893, 0.999]	0.854* [0.731, 0.999]	0.966 [0.769, 1.214]
Joint Growth Test (p)	0.047	0.048	0.769
<i>Quadratic</i>			
Level Change	0.998 [0.347, 2.875]	1.834 [0.544, 6.182]	0.316 [0.052, 1.919]
Change in Monthly Growth	1.025 [0.830, 1.266]	0.758 [0.408, 1.407]	1.276 [0.402, 4.048]
Change in Growth Over Time	-0.522 [-1.780, 0.737]	1.587 [-6.312, 9.486]	-4.922 [-24.709, 14.866]
Joint Growth Test (p)	0.077	0.123	0.848

Note: IRRs with 95% CIs; squared-time estimates use the log scale per year squared, and quadratic growth refers to the breakpoint. Poisson with month effects and serial-correlation-robust uncertainty. All placebo fits end in November 2022. * $p < .05$, ** $p < .01$, *** $p < .001$.

F Uncertainty Calculation

Table VII recalculates confidence intervals using Newey–West bandwidths of zero, three, six and twelve months. The coefficients remain unchanged because these checks alter the covariance estimate rather than the fitted mean. The linear growth-change and quadratic initial growth-change intervals remain above one under every bandwidth. The quadratic moderation estimate is less precise with bandwidth zero, where its interval includes zero. Evidence for faster growth is stable under these uncertainty calculations, while the precision of the changing growth rate varies.

Table VII: Sensitivity to the Uncertainty Calculation

	Monthly AI Solicitation Counts			
	Bandwidth 0	Bandwidth 3	Bandwidth 6	Bandwidth 12
<i>Linear</i>				
Level Change	1.076 [0.591, 1.957]	1.076 [0.617, 1.874]	1.076 [0.661, 1.751]	1.076 [0.696, 1.662]
Change in Monthly Growth	1.022* [1.004, 1.040]	1.022* [1.004, 1.040]	1.022* [1.004, 1.040]	1.022* [1.003, 1.041]
<i>Quadratic</i>				
Level Change	0.710 [0.341, 1.479]	0.710 [0.368, 1.373]	0.710 [0.410, 1.230]	0.710 [0.422, 1.195]
Change in Monthly Growth	1.070** [1.019, 1.124]	1.070** [1.023, 1.120]	1.070*** [1.033, 1.109]	1.070*** [1.037, 1.105]
Change in Growth Over Time	-0.146 [-0.298, 0.005]	-0.146* [-0.289, -0.003]	-0.146* [-0.258, -0.034]	-0.146*** [-0.233, -0.060]

Note: IRRs with 95% CIs; squared-time estimates use the log scale per year squared, and quadratic growth refers to the breakpoint. Bandwidth changes Newey–West covariance, with December 2022 fixed. * $p < .05$, ** $p < .01$, *** $p < .001$.

G Count-Separation Diagnostic

Table VIII reports an ODA-style threshold-classification diagnostic using recorded monthly counts. For a specified date, the procedure selects the count threshold and classification direction that best distinguish months before and after that date. Balanced accuracy averages the share correctly classified in each period. Effect strength for sensitivity (ESS) is defined as $100(2 \times \text{balanced accuracy} - 1)$, with zero representing chance-level balanced accuracy and 100 representing perfect classification. The December 2022 ESS of 60.36 corresponds to balanced accuracy of approximately 80.2 percent. It measures separation between count distributions, rather than a percentage increase in procurement or a causal effect.

The actual-date calculation uses all 113 months, while the earlier-date search uses only the 71 pre-release months. The search tests 60 candidate dates with at least six months on each side. At every date, the threshold and direction are selected again. Permutation tests use 1,999 reshufflings of individual months or blocks of three, six and twelve months. The earlier-date test compares the largest observed ESS with the largest ESS found across the same candidate dates in each reshuffling, accounting for the date search. Reported permutation probabilities use the adjustment $(1 + \text{exceedances})/(1 + 1,999)$.

Separation at December 2022 is statistically detectable under each reshuffling scheme, while the largest earlier separation is not detectable after accounting for the date search. The calculations use different sample lengths and horizons, and the procedure does not remove underlying trends or seasonality. Block reshuffling preserves ordering within blocks while relying on exchangeability across blocks. This diagnostic supports a difference in recorded activity across periods, but it does not resolve the regression trend sensitivity or identify why demand changed.

Table VIII: Count-Separation Diagnostic

	Monthly AI Solicitation Counts			
	Shuffled Months	3-Month Blocks	6-Month Blocks	12-Month Blocks
December 2022 ESS	60.36	60.36	60.36	60.36
December 2022 Test (p)	0.0005	0.0005	0.0005	0.0030
Largest Pre-Release ESS	46.65	46.65	46.65	46.65
Pre-Release Search Test (p)	0.133	0.166	0.333	0.417

Note: ESS is chance-corrected classification accuracy; 1,999 permutations. Pre-release search p-values account for testing 60 candidate dates.

H Fitted Paths and Confidence Intervals

Figure 3 compares expected monthly activity with the projected baseline under each specification. To prevent recurring seasonality from obscuring the comparison, the plotted prediction for each date averages over all twelve calendar-month categories. The baseline uses the same fitted model while setting the level, growth-change and squared-time terms to zero. Both paths therefore depend on shared coefficients estimated using the full analysis window. The baseline is a model-based projection, not a separate untreated series or a model fitted only to pre-release observations.

Let $\mathbf{x}_{tm}$ contain the model predictors for date t and calendar-month category m , and let $\hat{\boldsymbol{\theta}}$ and $\hat{V}$ denote the estimated coefficients and their Newey–West covariance matrix. The plotted mean and gradient of its logarithm are


$$\begin{aligned}\bar{\mu}_t &= \frac{1}{12} \sum_{m=1}^{12} \exp(\mathbf{x}_{tm}^\top \hat{\boldsymbol{\theta}}), \\ \mathbf{g}_t &= \frac{\sum_{m=1}^{12} \exp(\mathbf{x}_{tm}^\top \hat{\boldsymbol{\theta}}) \mathbf{x}_{tm}}{\sum_{m=1}^{12} \exp(\mathbf{x}_{tm}^\top \hat{\boldsymbol{\theta}})}.\end{aligned}\tag{H.1}$$

The pointwise 95 percent interval is calculated as

$$\exp\left\{\log(\bar{\mu}_t) \pm 1.96\sqrt{\mathbf{g}_t^\top \hat{V} \mathbf{g}_t}\right\}.\tag{H.2}$$

This delta-method calculation uses the covariance among all fitted coefficients, including the calendar-month terms. The intervals are symmetric on the log scale and positive without clipping, with legitimate asymmetry after transformation to expected counts. They describe uncertainty in each expected mean, rather than prediction intervals for individual monthly counts or simultaneous bands across the entire series. Predictions across dates share parameters, so these pointwise intervals should not be summed to construct uncertainty for cumulative differences.

Figure 3: Estimated Activity and Projected Baseline



Note: Expected monthly counts averaged over calendar-month effects, with pointwise 95% confidence intervals. The vertical guide marks December 2022.