Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Risk-Aware AI Procurement in the U.S. Public Sector: Stochastic Optimization and Technology Governance

Audrey Rah

arahimi@uh.edu

University of Houston https://orcid.org/0000-0001-7501-7809

David W Johnson

University of Houston - Downtown

Research Article

Keywords: Artificial intelligence, public procurement, appropriations, encumbrance, consumptionbased pricing, financial risk, Texas government, cloud services, stochastic programming, conditional value-atrisk, distributionally robust optimization

Posted Date: November 11th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11679788/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Risk-Aware AI Procurement in the U.S. Public Sector: Stochastic Optimization and Technology Governance

Audrey Rah[1,*] · David W. Johnson[2]

[1] Department of Electrical and Computer Engineering, University of Houston, Houston, TX, USA

[2] Assistant Vice President, University of Houston, UIT Technology Services and Support, Houston, TX, USA

[*] Corresponding author: arahimi@uh.edu

David W. Johnson: dwjohnson@uh.edu

ORCID: Audrey Rah — 0000-0001-7501-7809

***Abstract*—Public institutions acquire artificial intelligence (AI) services under established procurement, appropriation, and governance requirements, while commercial offers often state a rate and leave the quantity to later use. Executive leaders, procurement officers, financial authorities, and governance professionals need coordinated decision support so that a required review, a funding figure, and a management control are considered together. The objective is to give these decision-makers transparent, comparable financial estimates for AI service arrangements under uncertain demand. Operational controls and optimization are complementary parts of the proposed solution. The framework maps the reviewed rules onto subscription, consumption-based, prepaid-credit, and hybrid arrangements and labels each step as a legal requirement, a point the reviewed sources do not settle, or a recommended management practice. Ceilings, quotas, usage alerts, incremental funding, and named human approval are recommended practices, not Texas mandates. The proposed governance-constrained, risk-aware decision optimization framework supports executive and institutional decision-making under uncertain AI consumption, financial exposure, and public-sector procurement requirements, while preserving human oversight and legally authorized decision authority. The model reports expected expenditure, conditional value-at-risk, demand ambiguity, budget exposure, and process timing. The experiments use synthetic data and characterize model behavior rather than institutional savings. Solver outputs are planning estimates; legal determinations and purchase authority remain with authorized officials.**

***Index Terms*—Artificial intelligence, public procurement, appropriations, encumbrance, consumption-based pricing, financial risk, Texas government, cloud services, stochastic programming, conditional value-at-risk, distributionally robust optimization.**

***Managerial Relevance Statement*—**Engineering and technology managers who fund variable-cost AI services need numbers that separate planning estimates from legally binding obligations. This paper supplies a reproducible risk-aware budgeting architecture—expected expenditure, conditional value-at-risk, Wasserstein demand ambiguity, funding caps, and authorization limits—for comparing subscription, consumption, prepaid-credit, and hybrid arrangements under identical assumptions. The outputs are financial-clarity tools for officers already bound by Texas appropriation, procurement, cloud-security, and AI-governance stops. Purchase and adoption decisions remain with authorized officers, and the reported figures are synthetic planning results.

I. Introduction

Public institutions are buying AI as commercial software, as a cloud service, and as a technology service whose invoice depends on use. A subscription states a price for a term or a seat. A consumption, token, credit, or other metered price states a rate and leaves the quantity open. Hybrid offers combine a committed base with metered overage. Texas purchasing law was written for automated information systems and for commodity technology, not for a named class of generative-AI invoices [1], [2]. The 89th Legislature nevertheless directed the Department of Information Resources (DIR) to adopt an AI code of ethics and minimum standards for heightened-scrutiny systems, and it prohibited specified governmental uses of AI [2], [3], [4].

Building on our previous study of human responsibility and organizational outcomes following AI adoption [5], this paper examines the financial, procurement, and governance decisions that precede deployment, connecting organizational consequences with risk-aware institutional decision-making.

The operational problem is narrower than "AI is hard to buy." Several Texas decisions require a number before or at award: which competitive procedure applies, whether a commodity contract is lawful, whether a purchase is a major information system, whether funds may be encumbered, and whether unencumbered appropriation will lapse [1], [6], [7], [8], [9]. Metered AI makes that number depend on future demand, model choice, and vendor rate changes. This paper asks which of those decision rules are satisfied by existing Texas text, and which remain unspecified when the price is a rate rather than a sum.

A wrong classification changes the procedure. Under Tex. Gov't Code § 2157.068, for a state-agency commodity purchase that section covers, the number of vendors solicited, the availability of an alternative purchasing method, and a $10 million contract limit all turn on contract value [1], [10]. Under Tex. Gov't Code §§ 403.021 and 403.071, money is paid from an appropriation only within statutory windows, and unencumbered nonconstruction balances lapse [7], [8]. An agency that cannot say, at award, what it is legally obligated to pay cannot apply those rules mechanically. Federal oversight has already recorded a related estimating problem for AI acquisitions [11]. Texas institutions need an account of their own rules before importing federal practice.

Business process optimization is incorporated here as a process-oriented planning approach for variable-cost AI procurement. It considers the administrative workflow, including the sequence of activities, the dependencies among them, the approval gates the file already contains, and the overlaps that remain feasible when a cited rule does not order two activities. The objective is to find feasible process changes that reduce cycle time and avoidable delay while keeping statutory requirements, institutional responsibilities, and human authority outside the solver. The solver does not replace the officer who classifies the buy, signs the order, or certifies funds.

Disconnected budgeting and procurement decisions can select a funded amount that the chosen administrative path cannot execute on time, or a process configuration whose prepaid outlay is inconsistent with the officer's risk tolerance and authorization limit. Joint optimization is valuable for financial clarity when those couplings can change the selected package, because it then enforces them in one feasible planning estimate: the package, the quantities, the funded amount, and—in the extended model—the activity schedule must be admissible together under the same objective and the same caps. When the financially preferred package is already the administratively cheapest feasible package, finance-first, admin-first, and integrated selection coincide, and the integrated model certifies that coincidence rather than changing the award. Process timing is therefore analyzed together with procurement expenditure, uncertain demand, and financial tail risk. Two-stage stochastic programming supplies the split between a decision made before demand is known and a recourse decision made after demand is observed [12]. Conditional value-at-risk supplies a tail-risk measure and a linear representation of that measure on a finite scenario set [13]. The first numerical model evaluates three predefined packages with assumed days and hours. The extended model in Section IX-E schedules the activities in a simulated procurement file and places a finite-support Wasserstein ball on the training demands, so that process timing, risk sensitivity, and demand ambiguity interact inside a single mixed-integer program. The computation optimizes that integrated decision for one file; queues of several files are evaluated by list scheduling, and the contract-value and encumbrance questions left open in Section VIII remain with the responsible officers.

The gap emerges against the federal record. Office of Management and Budget (OMB) Memorandum M-25-22 and U.S. Government Accountability Office (GAO) report GAO-26-107859 address federal acquisition of AI [11], [14]. GSA Acquisition Letter MV-21-06 addresses federal purchase of cloud computing on a consumption basis [15]. This review did not identify a Texas statute, Comptroller policy, or DIR rule that performs the same function for metered AI: defining contract value for a usage-priced service, stating what portion of a ceiling is a binding encumbrance, or requiring usage alerts. The absence of such a document in the sources searched is a gap in published statewide guidance; individual agencies may maintain internal procedures outside those sources.

Given this gap, the study asks seven related questions, which move from the legal setting to the planning model. The first three concern the Texas rules themselves. They ask which Texas appropriation, procurement, cloud, and AI-governance rules determine how a public institution may buy an AI service; which of those rules presuppose a contract value or a binding obligation that can be stated at award; and where subscription, consumption, credit, token, and hybrid prices fit those rules, and where the sources leave the fit unspecified. The fourth question turns to the federal record. It asks which federal instruments address the same fiscal problem, and which of their controls are available to Texas institutions only as management practice. The fifth question asks what interim operating procedure an officer can follow where statewide Texas sources leave contract value or encumbrance unspecified, without describing that procedure as a legal mandate. The last two questions concern the planning model. The sixth asks what planning calculation can quantify expected expenditure, expenditure-tail risk, budget exposure, and a funded amount for alternative price mixes and process configurations under uncertain demand, without treating the calculation as a determination of contract value, encumbrance, or a purchase authorization. The seventh asks, on a simulated file, when an activity schedule, or a finite-support Wasserstein ball, changes the package and the funded amount selected by the nominal two-stage planning model under identical assumptions.

In answering these questions, the paper keeps five distinctions explicit throughout. (i) Contract value is the figure a cited purchasing rule uses for procedure and thresholds; expected expenditure is a planning estimate of cash outlay under uncertain use. (ii) A legally binding financial obligation is what an awarded instrument already requires the entity to pay; a planning estimate is not that obligation. (iii) A budget ceiling or not-to-exceed figure is an exposure limit; it is not, by itself, realized expenditure. (iv) Financial risk (including conditional value-at-risk of expenditure) is not the same object as governance compliance (TX-RAMP, heightened scrutiny, nondelegable human approval). (v) Reported numerical tables are synthetic or otherwise labeled, and they characterize model behavior rather than measured institutional

outcomes. The model identifies mathematically efficient or feasible procurement configurations under stated assumptions; purchase decisions remain with authorized officers.

These questions and distinctions define the paper's three contributions. The first contribution is a decision framework for buying and budgeting a variable-cost AI service when the rate is known and the quantity is not. The framework is the sequence, the decision table, and the decision rule in Section VIII. The review of Texas text serves as the input to that framework. The framework keeps three labels separate. A Texas requirement is stated only when a cited statute, rule, or Comptroller or DIR instruction already imposes it. An unresolved issue is a point the reviewed statewide sources do not settle, including the contract value of a meter and whether a usage ceiling is fully encumberable. A recommended management practice is a control this paper proposes where those sources are silent. Ceilings, user quotas, usage alerts, incremental funding, and a named human approval before spend moves past the funded amount are in that third class. They are not Texas legal requirements. The General Services Administration (GSA) consumption procedure is the analogy for that class: a transferable management pattern drawn from one federal schedule rather than from Texas law. The supporting map shows that Chapter 2157 covers technology services regardless of fee method while its thresholds are written as contract value [1], and that 1 Tex. Admin. Code § 212.1 and the Version 4.0 procurement guide value the binding purchase order rather than a nonbinding fund reservation [16], [17]. GAO-26-107859 provides federal evidence of AI cost uncertainty and is used here only for that federal context [11].

The second, and central modeling, contribution is the integrated risk-aware decision model for financial planning under demand uncertainty and governance constraints in Sections IX, IX-E, and X. The contribution is a formulation of the procurement decision that is solved, analyzed, and interpreted as one problem. The model jointly selects a permissible package among subscription, consumption, prepaid-credit, and hybrid arrangements, the associated quantities, a funded amount, and—in the extended model—an activity schedule, and it reports expected procurement expenditure, the expenditure tail (CVaR), unmet demand, service coverage, and cycle time under identical assumptions. The formulation draws on established foundations, namely two-stage recourse [12], the Rockafellar–Uryasev linearization of conditional value-at-risk [13], a finite-support 1-Wasserstein dual on the training atoms [18], and resource-constrained list scheduling, with a time-indexed mixed-integer program available when the list is not already optimal. The analytical results characterize how that dual behaves inside the procurement model. Corollaries 1 and 2 show that, for a fixed loss, the worst-case expectation is invariant to the radius exactly when the recourse expenditure is constant on the separated training atoms, and two fixed profiles then cross at a critical radius. On the published zero-risk-weight profiles that radius is 639.6 demand units and falls between the radius factors already reported. Proposition 1 derives it in closed form, as the exposed profile's radius-zero advantage divided by the token price, and Proposition 2 states when summing the two separate worst-case suprema is exact and when it is conservative relative to the single-distribution worst case over the same ambiguity set. On the frozen hybrid profile that conservatism is at most $7.2. A prepaid plan with no token use is one sufficient case of a constant loss. Remark 1 formalizes that a hard governance gate removes plans from the feasible set, so a conditional-value-at-risk comparison cannot waive a failed security gate. The officer-supplied funding cap and human-authorization limit enter as binding constraints, so a mathematically admissible planning estimate is consistent for budgeting, authorization exposure, and scheduling at once. That coupling is what disconnected budgeting and procurement decisions cannot guarantee. The ambiguity layer is specified on the training atoms, and Section IX-E states the scope of that specification precisely [18]. The computational comparisons identify where risk sensitivity, demand ambiguity, or scheduling changes the optimum, and where those layers leave a prepaid optimum unchanged because the relevant constraints are inactive or the expenditure is nonrandom. Identical optima under those conditions indicate solution stability. A finance-first sequence, an admin-first sequence, and the integrated rule are compared on identical draws in Section X-F2; that comparison is a sensitivity analysis of additive process charges built from established methods. Section XI positions the formulation against established methods.

The third contribution, developed computationally in Section XII and detailed in the supplementary material, extends that foundation with a governance-constrained multi-objective layer: a complete COBIT 2019 mapping of all forty objectives, hard security and residual-risk gates that cannot be traded away by penalties, vendor selection over a simulated catalog, epsilon-constraint Pareto enumeration, ablation, and decision-switching experiments on the same seeds as the validated baseline. Governance constraints remain distinct from financial-risk measures: a plan can be CVaR-efficient and still infeasible under a failed security gate.

These three contributions are bounded by the scope that Section XIII states for the central manuscript. Exploratory process-cost sensitivity experiments on a generic simulated workflow, including AI adoption comparisons, are reported in the supplementary material and do not alter the validated financial baselines.

The remainder of this paper is organized as follows. Section II reviews the Texas appropriation, procurement, cloud-security, and AI-governance rules that bound a purchase, and it retains the federal instruments as a comparison baseline. Section III characterizes subscription, consumption, credit, and hybrid pricing and defines the planning cost functions. Section IV documents the source-selection protocol, the planning-model definitions, and the uncertainty and evaluation protocol. Sections V, VI, and VII report the findings on procurement, budget-

ing, contracting, governance, and cost-control constraints, analyze where AI pricing strains the Texas rules, and compare the federal evidence. Section VIII presents the operational interim control procedure, its decision rule, and the points at which the optimization model may be called. Section IX formulates the integrated risk-aware optimization model, including the activity network, the finite-support Wasserstein layer, and the analytical results on its critical radius and risk charge. Section X reports the computational evaluation, including the base comparison, the sensitivity analyses, the extended experiment, the thirty-seed validation, ablation, and decision-grid studies, and a labeled comparison of finance-first, admin-first, and integrated selection. Section XI positions the formulation against established methods, and Section XII adds the governance-constrained multi-objective extension. Section XIII states the scope and applicability of the paper. Section XIV discusses the implications of the findings, Section XV sets out the limitations, Section XVI describes directions for future research, and Section XVII concludes. The supplementary material provides the full COBIT 2019 mapping and the complete epsilon-constraint tables.

We begin with the Texas institutional and regulatory setting, because those rules determine what a purchase may state at award and therefore what any later calculation is permitted to inform.

II. Institutional and Regulatory Background

A. Texas Public-Sector Procurement

Chapter 2157 governs purchase of automated information systems, including a service related to the automation of an information system [1]. A state agency must use the commodity method in § 2157.068 or another purchasing method the Comptroller designates for best value [1]. "Best value" under § 2157.003 means the lowest overall cost of an automated information system, determined by considering factors that include purchase price, compatibility to exchange existing data, and capacity for expansion and upgrade [1]. Those factors matter at renewal. They do not price a meter.

Section 2157.068 defines commodity items to include technology services, "regardless of the method of charging fees" [1]. DIR must, to the greatest extent practicable, negotiate a specific price for commonly purchased commodity items, and it must document why a percentage discount from list price is best value [1]. A token rate is neither of those named tools. DIR may charge an administrative fee that becomes part of the customer's cost [1].

For a state agency commodity purchase, § 2157.068(e-1) sets procedure by contract value: at or below $50,000, a direct award is permitted; above $50,000 through $1 million, the agency must request pricing from at least three listed vendors in the category; above $1 million through $10 million, it must request pricing from at least six, or from all vendors if fewer than six are listed [1]. Subsection (e-4) permits, above $5 million through $10 million, a Comptroller-designated method under § 2157.006(a)(2). Subsection (e-2) states that a state agency may not enter a contract to purchase a commodity item if the contract value exceeds $10 million, with a bulk-purchase exception in (e-3) [1]. Subsection (f) authorizes DIR to require by rule that a covered purchase be made under a DIR contract unless the agency obtains a DIR exemption or express prior approval from the Legislative Budget Board (LBB), or DIR certifies in writing that the item is not available under an existing DIR contract; under subsection (g), the agency must request the DIR exemption before seeking LBB approval [1]. DIR states that state agencies are required to purchase hardware, software, and technical services through DIR under § 2157.068, and that a cooperative-contract order is typically a purchase order to the vendor that cites the DIR contract number [19].

When a DIR commodity contract requires a statement of work (SOW) to initiate services, § 2157.0685 defines that SOW, for contracts of more than $50,000, as the document stating requirements, including deliverables, performance specifications, and other requirements specific to the vendor. The agency must consult DIR before sending the SOW to a vendor. The SOW is not valid, and money may not be paid under it, unless DIR first signs it [1], [20].

Delegated open-market purchasing is a different track. Under 34 Tex. Admin. Code § 20.82, competitive bidding is not required at or below $10,000; informal competition applies above $10,000 through $25,000; and formal solicitation from the Centralized Master Bidders List applies above $25,000. Goods purchases estimated at or below $50,000 and service purchases estimated at or below $100,000 are within the general delegation. The agency's cost estimate "must be developed in good faith using a method that is reasonable under the circumstances" [21]. Items available on a term or cooperative contract are outside that delegated authority, except below the contract's minimum order quantity [21]. *Requirement.* The good-faith estimate rule governs delegated purchases. It does not, by its own text, govern a purchase made on a DIR cooperative contract.

Agencies must plan and document covered procurements consistently with the state procurement manual and contract management guide, including the assessment of need, choice of method, specifications, evaluation criteria, the conduct of the purchase, and management of the resulting contract [22]. The Comptroller's current guide is Version 4.0, which the Statewide Procurement Division describes as incorporating procurement legislation from the 89th Legislature [16]. Section 20.181 refers to "the state procurement manual and contract management guide" without naming a version, and this paper treats Version 4.0 as that document. Quotations from the guide are taken from Version 4.0, publication 96-1809, and are limited to sentences that bear on contract value, statements of work, data-center services, change control, indefinite contracts, and cross-year funding.

Section 2157.002 limits subchapters A, B, and D of Chapter 2157 to a state agency to which Chapter 2054

applies [1]. Section 2054.003 defines "state agency," except as otherwise provided, to include a university system or an institution of higher education as defined by Education Code § 61.003 [6]. Section 2054.0075 provides that Chapter 2054 does not apply to a public junior college or junior-college district except as necessary for information-security standards and participation in specified shared technology services [6]. Local governments may use a Chapter 2157 method, and doing so satisfies a state competitive-bidding requirement for that purchase [1]. They are not the "state agency" addressed by the mandatory value bands in § 2157.068(e-1). Eligible entities may purchase commodity items through DIR [1].

Inclusion in the § 2054.003 definition does not, by itself, put every institution of higher education under those bands or under the statement-of-work signature. Education Code § 51.9335(d) provides that, subject to § 51.9337, Subtitle D, Title 10, Government Code, which contains Chapter 2157, does not apply to an acquisition of goods or services under § 51.9335, except provisions and rules on historically underutilized businesses and on procurement from persons with disabilities. The same subsection says the institution may, but is not required to, buy under that subtitle [10]. Section 51.9335(e) uses the Education Code § 61.003 meaning of "institution of higher education," including a school of veterinary medicine and a health-care facility operated by a medical or dental unit, and it excludes The University of Texas M. D. Anderson Cancer Center and a public junior college. Sections 51.9337(a) and (j) condition the exemption. An institution may not use the § 51.9335 authority unless it meets the required standards, and a suspension returns it to the state-agency acquisition laws, including Subtitle D [10]. DIR states the same customer limit in its procurement guidance: the § 2157.068 solicitation counts apply to state agencies and do not apply to higher-education customers, and DIR's statement-of-work review is addressed to state agencies [20], [23]. A § 2157.068 or § 2157.0685 duty is therefore stated below as a requirement only for a buyer that section covers. It is not a requirement for an institution acquiring under § 51.9335 while that authority is in force. Whether a named university, medical unit, or junior college is inside the exemption is an institutional fact this paper does not decide.

B. State Budgeting and Appropriations

The Comptroller's eXpendit guidance states that the Texas Constitution prohibits a state agency from spending money in the state treasury unless the Constitution or the Legislature has appropriated it, and that the Constitution limits appropriations of money to be drawn from the treasury to two-year terms [9]. The same guidance states that General Appropriations Act (GAA) appropriations are usually effective for one fiscal year, from September 1 through August 31, and that this one-year practice is legislative policy rather than the constitutional maximum [9]. The Legislature enacts a GAA once each biennium. Riders limit appropriated amounts. They do not themselves create additional appropriations [9], [24].

Section 403.071(b) provides that a claim may not be paid from an appropriation unless it is presented to the Comptroller not later than two years after the end of the fiscal year for which the appropriation was made, with longer windows for specified construction, repair, and grant categories that are not the ordinary case of an AI service invoice [7], [9]. Section 403.021 requires reporting of payables and binding encumbrances into the Uniform Statewide Accounting System, and it requires the Comptroller, on November 1, to lapse unencumbered nonconstruction appropriation balances for prior appropriation years based on the reported encumbrances and payables [7]. A claim is legally payable from an earlier appropriation only if that appropriation was encumbered to pay the claim before the appropriation expired [7].

eXpendit defines an encumbrance as an amount the agency is obligated to pay through a binding agreement. It states that an encumbrance can be made only for an awarded contract, not for an anticipated contract or a contract under negotiation, and that earmarking or budgeting an appropriation is not an encumbrance [9]. An appropriation may be encumbered only during its life. The unencumbered balance lapses [9]. Comptroller Accounting Policy Statement (APS) 018 tells agencies to avoid over-encumbering. For the binding-encumbrance certification, an outstanding encumbrance is a contract, agreement, or other action that legally obligates state funds [8]. The Manual of Accounts describes general-ledger account 9001 as commitments related to unperformed contracts for goods or services: anticipated expenditures evidenced by a contract or purchase order or determined by administrative action [25]. Read together, these sources tie encumbrance to an obligation that already exists. They do not describe a reservation of funds for consumption that no contract yet requires the state to pay.

Capital-budget riders identify and cap the categories and amounts an agency may spend for capital purposes and prohibit using those amounts for other purposes [24]. APS 026 includes, among capital categories, acquisition of information-resource technologies, including purchase of or contract for computer time, facility resources, maintenance, and training, and cybersecurity projects. For data-center services and shared technology services provided by DIR under Chapter 2054, including public and private cloud services, the guidance states that appropriations are available for assets with a biennial project cost or unit cost of more than $100,000 [24]. Whether a particular AI subscription falls inside an agency's capital rider depends on the GAA text for that agency and on the cost used to classify the purchase. This paper does not make that classification for any agency.

eXpendit also restates a last-quarter limit: during the last quarter of a fiscal year, a state agency using GAA funds may not expend more than one-third of the funds appropriated for that fiscal year, with stated exemptions, including expenditures contracted for in previous quarters

[9]. The Comptroller attributes the limit to GAA Article IX, Section 6.06 [9]. The rider is appropriations-act text and can change. The sources reviewed do not decide whether a later invoice for metered use, under an order placed in an earlier quarter, is an "expenditure contracted for" in that earlier quarter.

C. Technology, Cloud, and AI Governance

Section 2157.007 incorporates the National Institute of Standards and Technology (NIST) cloud-computing definition in Special Publication 800-145 as it existed on January 1, 2015 [1], [26]. A state agency must consider cloud options, including security benefits and cost savings, when buying for a major information-resources project, and must consider cloud options and compatibility when developing new software applications. The statute also provides, in text amended by a separate 2019 act, that the system or project must be capable of deployment and operation on cloud services unless integration, security, or cost makes that infeasible, in which case the agency must give advance notice to the LBB or the quality-assurance team [1]. Projects that use cloud options must meet or exceed state cybersecurity standards [1].

Senate Bill 475 enacted the cloud risk-and-authorization program as Tex. Gov't Code § 2054.0593 and assigned it to the Department of Information Resources [27]. House Bill 150, Section 16, transferred that section to Chapter 2063, redesignated it as § 2063.408, and amended it, effective September 1, 2025. Amended subsection (b) directs Texas Cyber Command, rather than the department, to establish the program [28], [29]. A state agency still may not enter or renew a contract for cloud computing services subject to the program unless the vendor demonstrates compliance, and it must require the vendor to maintain program compliance and certification throughout the term [29]. Section 50(c) of the same act directs the department to continue duties under Chapters 2054 and 2059 as they existed immediately before September 1, 2025, until the date stated in the memorandum required by Section 50(e). Section 50(d) requires the cybersecurity functions under Chapter 2063 to transfer to the command not later than December 31, 2026. This paper does not have that memorandum and does not state whether the continuation date has passed [28]. The information-security rules, including TX-RAMP, were transferred effective August 1, 2026, from 1 Tex. Admin. Code Chapter 202 (Title 1, Part 10) to Chapter 477 (Title 1, Part 19). The Texas Secretary of State's Texas Register transfer notice lists that transfer of the existing rules [30]. The rule language summarized below is that of the former Chapter 202 sections, which the notice moves without a stated change in text. Title 1, § 477.27 (formerly § 202.27), requires a state agency contracting for cloud services that store, process, or transmit agency data to confirm TX-RAMP certification before entering or renewing a covered contract on or after January 1, 2022, and to require the vendor to maintain certification [31], [32]. Section 477.5 (formerly § 202.5) sets Level 1 and Level 2 baselines according to impact category [33]. A parallel IHE rule appears at § 477.77 (formerly § 202.77) [34]. Whether a particular institution of higher education is covered by the agency rule, the IHE rule, both, or neither for a named cloud buy is a coverage fact for that buyer; this paper does not decide it for any named campus.

A "major information resources project" under § 2054.003 includes an information-resources project in a biennial operating plan whose development costs exceed $5 million and that takes at least a year to reach operations, involves more than one agency, or substantially alters work methods or client services, plus any project the General Appropriations Act designates as a major information resources project, and any project designated for additional monitoring under § 2261.258(a)(1) whose development costs exceed $5 million [6]. The same section defines a "project" as an initiative with defined parameters, specific objectives, planned activities, a scheduled completion date, and an established budget [6]. Section 2054.008 requires written LBB notice, within 30 days after award, of a contract for a major information system, defined to include computers aggregating more than $100,000, a computer-related service including software that costs more than $100,000, or telecommunications equipment that costs more than $100,000. A university system or institution of higher education gives that notice only if the cost exceeds $1 million [6].

Senate Bill 1964, effective September 1, 2025, defines an AI system as a machine-based system that infers from information how to generate outputs such as predictions, content, recommendations, or decisions [2]. A heightened-scrutiny AI system is one specifically intended to make, or to be a controlling factor in, a consequential decision about access to a government service, with stated exclusions for narrow procedural tasks, improvement of a completed human activity, preparatory tasks, and detection of decision patterns [2]. DIR must adopt, and state agencies and local governments must adopt, an AI code of ethics aligned with the NIST AI Risk Management Framework (AI RMF 1.0), covering human oversight, fairness and accuracy, transparency, data privacy and security, redress, and the frequency of evaluations [2], [35]. Minimum standards for heightened-scrutiny systems must include assessment of known security risks, performance metrics, and transparency measures before deployment and at material change, and must address contractual risk-management duties for vendors [2].

DIR adopted 1 Tex. Admin. Code Chapter 219 effective March 18, 2026 [4], [36]. Section 219.11 requires covered entities to follow the code when they procure, develop, deploy, or use AI systems, including documented evaluation over the life cycle and vendor contract duties [4]. Section 219.22 requires a written risk assessment before development, procurement, deployment, or use of a heightened-scrutiny system, and again at a material change, and the AI risk officer must approve or deny deployment [4].

Section 219.24 adds an acceptable-use policy and a vendor duty to implement the NIST AI RMF or a comparable standard for those systems [4]. Section 2054.708 separately requires an impact assessment when a state agency, or its vendor, deploys or uses a heightened-scrutiny system [2]. In the Texas Register adoption discussion, DIR described the rule's vendor impact-assessment provision as permissive rather than mandatory [4]. This paper does not resolve any difference in force between that discussion and the statute. None of these duties states a price, a quota, or an encumbrance amount.

House Bill 149, effective January 1, 2026, adds Business and Commerce Code Chapter 552. "Governmental entity" excludes a hospital district and an institution of higher education [3]. The chapter restricts specified uses, including governmental social scoring and specified biometric identification, and it is enforceable by the attorney general [3]. It does not set a price, an encumbrance method, or a procurement threshold. Section 2054.705 assigns an advisory board to identify procurement-streamlining opportunities. Section 2054.706 directs DIR to administer a pre-deployment sandbox [2]. The DIR AI resources page reviewed for this study did not describe an operating sandbox [36]. Absence from that page does not establish that the program is unimplemented. Neither provision states how to value or encumber metered use.

D. Federal Comparison Baseline

Federal executive agencies obligate and expend appropriations under the Antideficiency Act, 31 U.S.C. § 1341, which prohibits obligations in excess of or in advance of an appropriation, and under 31 U.S.C. § 1502, which confines a time-limited appropriation to expenses properly incurred during its period of availability or to completion of contracts properly made and obligated in that period [37], [38]. Texas agencies are not subject to those federal sections. The structural similarity is that both systems separate an appropriation's availability from a later invoice, and both require an obligation if old-year funds are to pay a later claim [7], [38].

OMB Memorandum M-25-22 (April 3, 2025) is the current executive-branch memorandum on AI acquisition. It rescinds Memorandum M-24-18 [14]. It tells covered federal agencies to acquire effective and trustworthy AI in a timely and cost-effective manner; to address vendor lock-in through data portability, interoperability, and pricing transparency; to include contract terms on intellectual property, government data, privacy, testing, and high-impact risk practices; and, to the extent practicable, to evaluate value during administration, including operation and maintenance cost [14]. It encourages performance-based statements of objectives rather than statements of work that limit solutions, and it directs GSA to publish acquisition guides and a repository of AI acquisition lessons [14]. It does not bind Texas institutions.

III. AI Service Pricing and Procurement Characteristics

A. Pricing Patterns in the Reviewed Sources

The sources reviewed here support four pricing patterns. They do not show which pattern dominates Texas spending. No Texas invoice sample was available.

Subscription. GAO reported that selected federal agencies often acquired AI as a service and that most reviewed agreements included licenses, frequently for software procured as a service [11]. A subscription can be a determinate quantity: a named term, a named number of seats, or a named bundle. Section 2157.003's "purchase price" and § 2157.068's value bands can be applied to that sum if the sum is the contract value [1]. Renewals still face appropriation availability [7], [9].

Consumption and tokens. GSA defines consumption-based cloud, for its Federal Supply Schedule procedure, as an offering that is metered and accrues charges on a predetermined periodic basis, such as per second, hour, week, or month, and that is billed on actual usage during an elapsed period at predetermined prices or discounts [15]. GSA states that cloud quantities are often difficult to estimate because demand spikes, and that agencies have often obligated funds at the start so that capacity is available [15]. On September 10, 2026, GSA announced a OneGov agreement, expected to take effect October 1, 2026, offering federal, state, local, and tribal governments discounted token-based usage of specified models, with no platform fee, minimum order, or spend commitment, and with ordering through GSA's Multiple Award Schedule [39]. That announcement is evidence of a federal vehicle using token prices rather than of a Texas legal rule. This paper does not determine whether a Texas agency may use that vehicle without a DIR exemption under § 2157.068(f) [1], [39].

Credits and hybrid prices. GSA states that industry stakeholders described combinations of committed and consumption prices as ordinary commercial practice, and that its special ordering procedures allow combinations in one task order [15]. GAO reported that GSA's USAi effort pursued usage pricing rather than a license for every employee, with structured prices so that simple prompts and more complex model features were not priced alike [11]. A credit pool is, on these descriptions, a prepaid or committed quantity of metered units. The Texas sources reviewed do not define credits or tokens.

Uncertain demand and cost escalation. GAO found that officials at all five selected federal agencies described AI pricing and overall cost as a challenge, including uncertainty about how vendors price licenses and services, in a market the officials described as changing quickly [11]. Army officials told GAO that buyers can underestimate cost by counting model training and omitting cloud infrastructure and computing power over time, because they are more familiar with hardware acquisitions [11]. GSA officials told GAO that testing mattered in deciding whether a new model version justified a higher price [11]. NIST

AI RMF 1.0 is a risk framework, not a price schedule. Texas incorporates it for ethics and heightened-scrutiny governance [2], [4], [35]. It does not supply a forecast of tokens.

Interpretation. A subscription reduces quantity risk only for the subscribed quantity. It does not remove rate changes at renewal, seat growth, or overage. A pure meter removes the duty to pay for unused units and removes the ability to state total expenditure at award unless the contract also states a ceiling. Texas text reviewed here does not choose between those designs.

B. Planning Cost Functions

The four price forms need planning identities before a numerical model can use them. The identities in this subsection are accounting definitions for that model. They are not formulas for statutory contract value, and they are not a reading of § 2157.068 or of 1 Tex. Admin. Code § 212.1. Let $d \geq 0$ be realized demand in service units over the planning horizon. Let $q^{\mathrm{sub}} \geq 0$ be subscribed capacity, $q^{\mathrm{cred}} \geq 0$ a prepaid credit pool, and $u^{\mathrm{tok}} \geq 0$ token units actually served. The unit prices p^{sub}, p^{cred}, and p^{tok} are model inputs for the numerical study rather than Texas rates.

$$C^{\mathrm{sub}}(q^{\mathrm{sub}}) = p^{\mathrm{sub}} q^{\mathrm{sub}}, \tag{1}$$

$$C^{\mathrm{cred}}(q^{\mathrm{cred}}) = p^{\mathrm{cred}} q^{\mathrm{cred}}, \tag{2}$$

$$C^{\mathrm{tok}}(u^{\mathrm{tok}}) = p^{\mathrm{tok}} u^{\mathrm{tok}}, \tag{3}$$

$$C^{\mathrm{hyb}} = C^{\mathrm{sub}}(q^{\mathrm{sub}}) + C^{\mathrm{cred}}(q^{\mathrm{cred}}) + C^{\mathrm{tok}}(u^{\mathrm{tok}}). \tag{4}$$

Equations (1), (2), (3), and (4) are those identities. Subscription and prepaid credits are paid whether or not the units are used. Token expenditure is paid only for units served. Service from the subscription cannot exceed q^{sub}. Service from credits cannot exceed q^{cred}. A hybrid plan may use all three, with unmet demand $m \geq 0$. Equation (5) is the material balance:

$$s^{\mathrm{sub}} + s^{\mathrm{cred}} + u^{\mathrm{tok}} + m = d. \tag{5}$$

The expenditure that later enters the risk measure is C^{hyb}, not C^{hyb} plus a penalty for m. Unmet demand is a separate term. No identity in this subsection converts C^{hyb} into the amount the agency is legally obligated to pay. Table I places each identity against the decision an officer has to make at award. The open cells are the points the reviewed Texas texts do not settle. A subscription can name the quantity. A meter cannot. Credits can name the pool and still leave later use unknown. The hybrid row keeps the committed base and leaves the overage unresolved unless the order makes that overage a present obligation.

IV. Research Methodology

A. Source-Selection Protocol

The study is a documentary analysis of primary legal and administrative sources. The principal search was conducted on September 22, 2026. A follow-up search on September 23, 2026, tested seven questions against current official text: the contract-value method for pay-per-use, token, credit, or other metered technology; whether a not-to-exceed ceiling is fully encumberable when payment is due only for use; how the Uniform Statewide Accounting System (USAS) should record a usage-based technology obligation; whether statewide guidance requires usage alerts; whether DIR or the procurement guide contains consumption, cloud, token, or credit examples; when an amendment or further DIR approval is required after a change in consumption, rate, model, or scope; and whether 2025–2026 guidance issued after the AI statutes resolves those questions. The search targets were the Texas Constitution and Statutes site, the Texas Legislature's enrolled-bill record for the 87th and 89th Legislatures, the Texas Register, the Comptroller's fiscal-policy pages and Procurement and Contract Management Guide, Version 4.0, 1 Tex. Admin. Code Chapters 202 and 212 and 34 Tex. Admin. Code Chapter 20, DIR's buying, statement-of-work, cloud, and AI pages, the U.S. Code, OMB memoranda, GAO reports, GSA acquisition letters and news releases, and NIST publications incorporated by Texas or cited by those federal documents.

B. Inclusion and Exclusion

Priority 1 sources are Texas statutes, the Texas Administrative Code, Comptroller fiscal policies, and DIR rules, manuals, and procurement pages. Priority 2 sources are federal statutes and OMB, GAO, GSA, and NIST documents. A DIR-posted bulk-purchase agreement is cited only as contract text that DIR has published, and it is labeled as such [40]. The GSA OneGov release is cited only for the pricing mechanism GSA itself describes [39]. Blogs, vendor marketing pages, consulting promotions, news articles, and secondary summaries were excluded. The planning layer cites two methodological sources. Rockafellar and Uryasev supply the linear representation of conditional value-at-risk on a scenario set [13]. Birge and Louveaux supply the two-stage recourse model [12]. Neither work is a Texas legal authority or a source of Texas expenditure data.

C. Coding

Each retained provision was coded to the decision it controls: estimating value, choosing a vehicle, obligating funds, limiting a fiscal year, imposing a ceiling, reviewing security or AI risk, monitoring after award, amending scope, or closing and renewing. Codes were then grouped into the taxonomy in Fig. 1. A statement was marked as a requirement only when the source uses mandatory language or states a condition of payment, validity, or legal authority. A statement was marked as an interpretation when it describes how two provisions interact and the sources do not themselves state that interaction. Silence was recorded when the searched sources do not address the point.

TABLE I
PROCUREMENT ASSUMPTIONS AND FOUR AI PRICE FORMS. THE VALUE-BAND AND STATEMENT-OF-WORK ROWS APPLY TO A BUYER COVERED BY §§ 2157.068 AND 2157.0685. THEY ARE NOT DUTIES OF AN ACQUISITION UNDER EDUCATION CODE § 51.9335 WHILE § 51.9337 LEAVES THAT AUTHORITY IN FORCE [10].

Decision at award	Subscription	Consumption or tokens	Credits	Hybrid
Quantity known	A term or a seat count can be stated	Quantity remains a rate until the units are used	The pool size can be stated even when later use cannot	The base can be stated. The overage cannot
Contract value for a threshold	Often the stated price	Gov't Code § 2157.068 states no token formula. 1 TAC § 212.1 counts the binding purchase order, not a nonbinding encumbrance document [16], [17]	The face amount legally owed, or expected redemption. The meter formula is unspecified	The base is stated. Overage counts only if it is part of the obligated amount
Binding encumbrance	Fits an awarded, priced order	Fits usage only when that usage is an obligation	Fits the sum legally owed for the credits	Fits the committed base. Overage depends on the duty to pay
Appropriation lapse	Unused obligated funds follow the encumbrance rules	The balance lapses if it is not a present legal obligation [7], [9]	Prepaid credits may obligate the pool	Unobligated overage lapses
DIR signature and statement of work	For a buyer covered by § 2157.0685, only when the contract requires a statement of work over $50,000 [1], [10]	Same covered-buyer limit	Same covered-buyer limit	Same covered-buyer limit

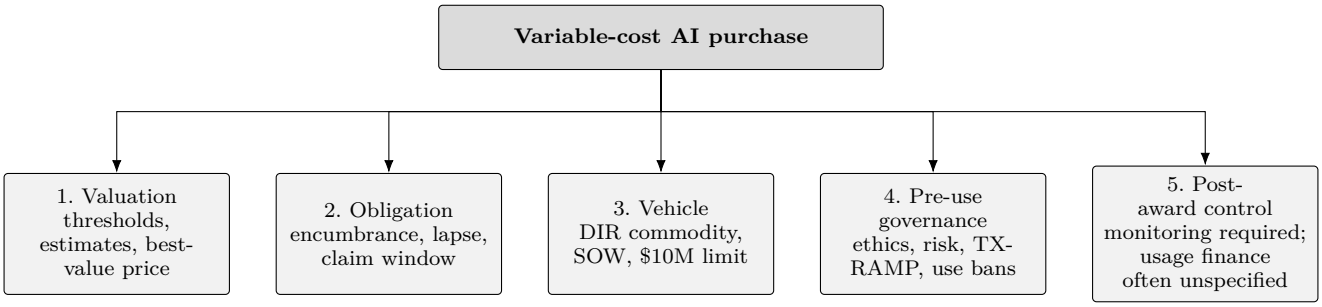


Fig. 1. Taxonomy of constraints on a variable-cost AI purchase. Categories 1–4 restate duties found in Texas sources. Category 5 combines duties the sources state with controls they do not state.

D. Validation

Operative sentences were checked against the statute, rule, or agency document retrieved on the access date. Where an agency page did not display the operative text, the study used the enrolled bill, the Texas Register adoption, or the Comptroller's published restatement, and it identifies which of those is quoted. The study does not use Texas invoice data, interviews, or institutional procurement records, and it does not estimate how often these valuation, encumbrance, or overrun problems occur.

E. Search Dates and Currency

The legal cutoff for the follow-up search is September 23, 2026. Chapter 219 was effective March 18, 2026 [4]. House Bill 149 was effective January 1, 2026 [3]. Senate Bill 1964 was effective September 1, 2025 [2]. The commodity-value definitions in 1 Tex. Admin. Code § 212.1 were amended effective February 10, 2022, and are restated in the Version 4.0 procurement guide [16], [17]. GAO-26-107859 is dated April 13, 2026 [11]. The GSA OneGov announcement is dated September 10, 2026 [39]. GSA's consumption-ordering page was rechecked on September 23, 2026; the page still publishes Acquisition Letter MV-21-06 and Supplement 1 and states that the letter remains in effect until rescinded or incorporated into the Federal Acquisition Regulation or the GSA acquisition manual [15]. Later amendments are outside the record.

F. Planning-Model Definitions

The planning calculation takes the documentary labels as given. It does not reclassify a statement as a

requirement, an interpretation, an unresolved question, or a recommended control.

Demand D is a nonnegative random variable in service units. The calculation uses a finite equiprobable scenario set S, with realization d_s and probability $\pi_s = 1/|S|$. No scenario probability is estimated from a Texas invoice file. The base distribution is lognormal with a stated mean and coefficient of variation (CV). That distribution is an assumption. Demand is measured in service units.

Decisions made before D is known are the subscribed capacity q^{sub}, the credit pool q^{cred}, the pre-authorized token budget R in dollars, and a binary configuration choice y_k for $k \in K$. The set K contains three configurations derived from Section VIII: hybrid, which may use committed units and tokens; committed, which may use subscription and credits only; and meter, which may use tokens only. A configuration that would award before the use screen, or that would increase token funding without a human authorization, is not in K.

Decisions made after d_s is known are units served by each instrument, unmet demand m_s, incremental token funding z_s, and a binary authorization a_s. Recourse cannot spend above the officer-supplied cap, and it cannot spend incremental token funds unless $a_s = 1$ and the chosen configuration allows a meter.

Let p^{sub}, p^{cred}, and p^{tok} be the unit prices in dollars per service unit, and let u_s^{tok} be the token units served in scenario s. Equation (6) is the procurement expenditure in scenario s, in dollars:

$$L_s = p^{\text{sub}} q^{\text{sub}} + p^{\text{cred}} q^{\text{cred}} + p^{\text{tok}} u_s^{\text{tok}}. \tag{6}$$

There is no encumbrance variable. L_s is a planning outlay under those prices, not a certification that L_s, R, or the cap is legally obligated.

G. Uncertainty, Objective, and Evaluation Protocol

Let $\beta \in (0,1)$ be the tail probability level and $\lambda \in [0,1]$ the weight on tail expenditure. Conditional value-at-risk (CVaR) is the expected expenditure in the upper tail beyond β. Equation (7) is the Rockafellar–Uryasev linearization on the finite scenario set, using the auxiliary η and excesses $v_s \geq 0$ [13]:

$$\text{CVaR}_\beta(L) = \min_{\eta} \ \eta + \frac{1}{(1-\beta)|S|} \sum_{s \in S} v_s, \tag{7}$$

subject to $v_s \geq L_s - \eta$. Equation (8) is the planning objective. Its constraint system is in Section IX:

$$\begin{aligned} \min \quad & (1-\lambda) \sum_{s \in S} \pi_s L_s + \lambda \, \text{CVaR}_\beta(L) \\ & + w_m \sum_{s \in S} \pi_s m_s + \sum_{k \in K} y_k (r_E E_k + r_T T_k) \\ & + \sum_{s \in S} \pi_s a_s (r_E e_a + r_T t_a). \end{aligned} \tag{8}$$

The first line is expected procurement expenditure and expenditure tail risk. The second line is the penalty w_m per unmet unit and the business-process term for the chosen configuration, using assumed staff hours E_k and cycle days T_k. The third line is an incremental authorization, which adds e_a hours and t_a days. The weights r_E and r_T are assumed values rather than salaries or delay costs observed at any institution. Cycle time and staff effort were not available from the legal sources, so both are labeled assumptions in Section X.

The objective does not weight a legal-compliance score, because the documentary sections do not supply one. Cap feasibility is reported as a constraint check. It is not a finding that a buy complies with Texas law.

The evaluation protocol has six steps. Draw a training scenario set and a disjoint test set from the same distribution, with fixed seeds. Solve the mixed-integer linear program on the training set. Freeze the first-stage quantities and configuration. Re-solve only the recourse on the test set under the same objective. Execute four comparison rules on that same test set: a fixed budget that spends the officer-supplied cap on credits; a deterministic expected-demand plan that solves the model at the mean demand with $\lambda = 0$ and then serves realized demand up to the frozen plan; a subscription-only plan sized to the mean; and a pay-as-you-go plan that uses tokens, pre-authorizes the assumed fraction of the cap, and may request authorized incremental funding. Record, for every rule, expected expenditure, CVaR_β, expected unmet demand, the share of test demand served, expected days, expected hours, the share of scenarios that use authorization, and the share of scenarios that satisfy the cap, the authorization rule, and the selected configuration. Repeat one factor at a time for demand volatility, token price, the officer-supplied cap, λ, and β. The protocol is specified so that those steps can be rerun from the script cited in Section X. The results of that run are reported there. Mak, Morton, and Wood give a Monte Carlo confidence interval for the optimality gap of a candidate first-stage solution [41]. The scores below are paired out-of-sample values of frozen decisions rather than those gap intervals. The extended experiment keeps the same freeze-and-test step and adds a separate duration sample, a staff capacity, an ambiguity radius, and a delay-weight sweep. Its script and its sample sizes are stated in Section X-E, and it complements the first protocol.

V. Findings: Constraints Established by Texas Sources

A. Procurement Constraints

Three procurement constraints are established by Texas text.

First, commodity procedure depends on contract value. The bands in § 2157.068(e-1), the alternative in (e-4), and the prohibition in (e-2) are mandatory for the state-agency purchases they cover [1]. They are not a requirement for an institution of higher education acquiring under Education Code § 51.9335 while § 51.9337 leaves that authority in force [10], [23]. *Requirement.* A buyer covered by § 2157.068 must locate the contract in a value band. *Interpretation.* The statute does not say whether the value of

a metered service is the not-to-exceed ceiling, the expected spend, the historical spend, or the maximum technically possible consumption. That silence is the estimating problem. It is not a finding that agencies misclassify contracts.

Second, the fee method is not itself a bar to the commodity category. Technology services are defined without regard to how fees are charged [1]. *Requirement.* A usage-priced technology service can be a commodity item if it meets the rest of the definition and DIR has established demand and a contract. *Interpretation.* Subsection (b) then tells DIR to negotiate a specific price where practicable, and it regulates percentage discounts from list price [1]. Specific prices and list discounts are the negotiation tools the subsection names. A token rate is neither term. The chapter therefore admits variable fees in the definition and speaks, in the negotiation standard, the language of determinate prices. The sources do not say that a token rate is noncompliant.

Third, pre-award competition and documentation assume a specification. Section 2157.003 evaluates lowest overall cost, including purchase price [1]. The lifecycle factors in § 2157.125, including operation and maintenance of hardware and software, apply to acquisitions under the competitive-sealed-proposal subchapter [1]. They are not, by the section's placement, the evaluation rule for every DIR purchase order. Section 20.181 requires a documented need, method, specification, and evaluation criteria for procurements subject to that chapter [22]. Section 20.82(b)(6) requires a good-faith, reasonable cost estimate for delegated purchases only [21]. *Interpretation.* A DIR purchase can satisfy the statute's demand for a value without satisfying the delegated-purchase estimate rule, because that rule is on a different track. Statewide text still does not prescribe the estimating method for a DIR consumption order.

DIR's ordering description adds an administrative fact: the customer typically issues a purchase order citing the DIR contract [19]. *Requirement.* Where the master contract requires an SOW above $50,000, DIR must sign before payment [1], [20], [42]. *Requirement.* Under 1 Tex. Admin. Code § 212.1, a DIR customer's fully executed purchase order, and a purchase-order change notice, constitute the contract. A document executed only to encumber funds, if it is not a binding transaction, does not constitute a contract for purposes of determining dollar value. A statement of work, license, maintenance agreement, or service agreement without an accompanying purchase order likewise does not [17]. The Version 4.0 guide tells an agency to use its best judgment to determine the monetary value of the anticipated information-technology commodity contract, and it states that the dollar value is the total value over the term together with modifications, renewals, and extensions [16]. The guide cites, for that total-value instruction, Tex. Gov't Code §§ 322.020 and 2054.008 and 34 Tex. Admin. Code § 20.25(b)(13). Section 20.25(b)(13) defines contract value, for that administrative chapter, as the estimated dollar amount a state agency may be obligated to pay under the contract and all executed and proposed amendments, extensions, and renewals [16], [43]. The guide also warns that other oversight bodies may not use the same standard [16].

Unresolved. No statewide Texas authority located in the reviewed sources specifies how to calculate that estimated obligation when the price is a token rate, a credit, or another meter. The 2025–2026 AI statutes and Chapter 219 do not supply the missing formula [2], [3], [4]. *Interpretation.* Section 212.1 does answer a narrower question. A fund reservation that is not itself a binding transaction is excluded from the dollar value used for the commodity threshold. The purchase order that creates the payment duty is included. A ceiling number and an encumbrance entry can therefore be different documents, and only the binding one counts toward that dollar value.

B. Budgeting Constraints

The budgeting constraints are older than AI and apply to any uncertain service.

Requirement. Treasury money is spent only under an appropriation, and the Comptroller states that such appropriations are constitutionally limited to two-year terms, with GAA appropriations commonly limited to one fiscal year [9]. A multi-year AI subscription cannot be funded, beyond that availability, by an appropriation that has not been made. A renewal is a new claim on whatever appropriation is then available.

Requirement. Unencumbered nonconstruction balances lapse, and a later claim is payable from an earlier appropriation only if that appropriation was encumbered for the claim before it expired [7]. eXpendit states that anticipated contracts and mere earmarks are not encumbrances [9]. APS 018 warns against over-encumbrance and defines a binding encumbrance as action that legally obligates state funds [8].

Interpretation. These rules create a double bind for metered AI, without any Comptroller ruling that says so. If the contract obligates the state only for measured use, usage that has not occurred is not yet a legal obligation, and the appropriation for that possible usage is not encumbered by an earmark. It lapses at the statutory point if still unobligated [7], [9]. If the agency encumbers a high ceiling "so the service does not stop," it may record as obligated an amount the contract does not require it to pay. APS 018 tells agencies not to over-encumber and to report obligations accurately [8]. The Manual of Accounts allows commitments evidenced by a contract, purchase order, or administrative action [25]. It does not say that a nonbinding forecast is account 9001. The correct USAS amount for a usage ceiling is therefore not settled by the AI statutes. It is a fiscal-interpretation question the searched policies do not answer for cloud or AI.

Capital-budget classification adds a second numeric test. APS 026 places DIR data-center and shared-technology services, including public and private cloud, into the capital discussion when biennial project cost or unit cost exceeds $100,000, and it lists computer time

among information-resource technology acquisitions [24]. *Interpretation.* An AI service bought as ordinary DIR commodity software is not, by that sentence alone, a DIR shared-technology capital item. An AI service that is DIR cloud, or that an agency classifies as computer time or other information-resource technology, can hit a rider cap if the estimated cost crosses the rider's threshold. Because the estimate is the uncertain number, the capital classification inherits the valuation problem. This study did not read each agency's GAA rider.

The last-quarter one-third limit is a timing constraint on GAA expenditures, with an exemption for expenditures contracted in earlier quarters [9]. *Interpretation.* A consumption spike in June through August can press that limit if the spike is a new expenditure. The exemption may cover it if the governing order was placed earlier. The Comptroller's restatement does not apply the exemption to metered technology.

C. Contracting Constraints

Section 2157.068(e-2)'s $10 million commodity-contract limit is a hard statutory ceiling for the contracts it covers, subject to the bulk-purchase exception [1]. The Version 4.0 guide restates that a state agency, not including a state institution of higher education, may not buy a commodity item through the DIR cooperative-contracts program if the contract value exceeds $10 million, unless the purchase is a DIR coordinated bulk purchase [1], [16]. That exclusion agrees with Education Code § 51.9335(d) for an institution acquiring under that section while the authority is in force [10]. *Interpretation.* Placing a not-to-exceed figure at or below $10 million does not, by itself, demonstrate compliance. The figure that matters under § 212.1 and the guide is the amount the agency may be obligated to pay over the term, including modifications, renewals, and extensions, and it excludes a nonbinding encumbrance document [16], [17]. A contract with no ceiling and no other statement of that amount has no number to test. The guide states, as general policy, that indefinite contracts are generally prohibited [16]. Actual invoices that later exceed $10 million raise a question the statute does not answer: whether value is fixed at award or tracked in performance. This paper does not treat a later overrun as a violation. It treats the missing choice among ceiling, expected use, and cumulative invoices as an ambiguity.

SOW validity is binary for a buyer covered by § 2157.0685. Without DIR's signature, money may not be paid [1], [10]. A change in rate, model, or usage limit that requires a new SOW cannot be paid under an unsigned revision. The statute does not define which consumption changes require a new SOW.

Section 2054.008 imposes a post-award notice when a major-information-system contract costs more than $100,000, or more than $1 million for a university system or institution of higher education [6]. *Interpretation.* Notice follows award and uses a cost figure. It does not tell the agency how to compute cost for a meter. Undercounting cost could omit a required notice. The statute does not discuss estimates or amendments.

A DIR-posted 2026 bulk-purchase agreement illustrates emerging contract text, not a general rule. The posted terms state that the contractor shall not use customer data to train AI technologies without the customer's prior written consent, and that AI use in performance must comply with the code of ethics under § 2054.702 and the standards under § 2054.703 [40]. *Requirement.* Those sentences bind the parties to that agreement. They do not establish a statewide clause library for token contracts.

D. Governance and Oversight Constraints

Governance duties are now specific. They are largely about risk, disclosure, and security, not about the invoice.

Requirement. Before procurement or deployment of a heightened-scrutiny system, § 219.22 requires a written risk assessment and a deployment decision by the AI risk officer [4]. *Requirement.* Section 219.11 applies ethical duties, including vendor contract terms and documented evaluation, to AI systems generally, not only to heightened-scrutiny systems [4]. *Requirement.* Covered cloud contracts require demonstrated TX-RAMP compliance before award or renewal [29], [31]. *Requirement.* Chapter 552 prohibits specified governmental uses by covered governmental entities [3]. Institutions of higher education are outside that chapter's definition of governmental entity [3]. They are not thereby outside Chapter 2054 or Chapter 219, which use different coverage.

Interpretation. These duties add lead time and contract terms before an order is lawful to deploy. They do not stabilize a token price. GAO reported that officials at five federal agencies described traditional acquisition timelines as poorly aligned with the speed of AI change, including long market-research and approval periods [11]. That is federal interview evidence from a nongeneralizable sample. It is not a measurement of Texas DIR, TX-RAMP, or Chapter 219 processing times. The Texas sources establish that the steps exist. They do not establish how many days the steps take or that products change faster than those days.

Major-information-resources oversight is keyed to development cost above $5 million and to project attributes, including a scheduled completion date and an established budget [6]. *Interpretation.* Ongoing token consumption may never be a "development cost" and may lack a completion date. If so, large operating AI spend can sit outside major-project oversight while still sitting inside commodity, appropriation, and AI-ethics rules. The statute does not say this. It is a consequence of reading the project definition against a service that has no planned end.

Section 2054.705 assigns an advisory board to identify procurement-streamlining opportunities [2]. *Requirement.* The assignment exists. It is not a legislative finding that current thresholds are unworkable.

E. Operational Cost-Control Constraints

Texas sources require some post-award control and are silent on others.

Requirement. Section 219.11 requires formal accuracy monitoring over the life cycle and documented regular evaluations of whether the system still benefits the population served [4]. Section 20.181 requires contract management to be documented and ties planning documents to later vendor-performance evaluation under § 20.509 [22]. Section 403.071(h) authorizes the Comptroller to require agencies to retain documentation for post-payment audit and to respond to invalid claims, including by seeking a refund or reducing remaining appropriations [7].

The sources reviewed do not require a consumption quota, a 50% or 75% funding alert, incremental funding of a meter, a departmental chargeback rate, or human approval before a user exceeds a quota. Those controls appear in GSA's federal cloud procedure, discussed below as an analogy [15]. Using them in Texas is a recommended management practice unless an agency's own contract makes them a duty. They are not Texas legal requirements.

Vendor lock-in is addressed directly in M-25-22 for federal agencies: solicitations should seek knowledge transfer, data and model portability, clear licensing, and pricing transparency, and contracts should keep switching from becoming cost-prohibitive [14]. Texas § 2157.003 includes compatibility and the capacity to expand as best-value factors [1]. Section 219.11 requires governmental entities to demand enough transparency to understand how a system operates, the source of training data, and intended uses [4]. *Interpretation.* Those factors support a portability evaluation. They do not require a portability clause, a price-index clause, or an exit schedule.

VI. Analysis of Fit and Ambiguity for Metered AI Pricing

A. Where the Rules Fit, and Where They Do Not

The commodity definition already covers a technology service regardless of fee method, so a meter is not outside Chapter 2157 by invoice design alone [1]. Covered cloud still needs TX-RAMP before award and renewal [29], [31]. Covered AI still needs the Chapter 219 ethics duties, and heightened-scrutiny systems still need the risk-officer decision [2], [4]. Chapter 552 still screens prohibited uses for the entities inside its definition [3]. Encumbrance law still distinguishes a present legal obligation from a forecast [8], [9]. None of those texts converts a token rate and an unknown quantity into a contract value or an encumbrance amount. Table II is that map. The open cells are not implied duties.

B. Where AI Introduces Friction or Ambiguity

Table II pairs each instrument with what it settles and what it leaves open. The open cells are not implied duties.

The interaction that matters for budgeting is sequential. An agency must know enough about value to choose a legal vehicle [1]. It must know enough about obligation to encumber lawfully and to let the remainder lapse [7], [8], [9]. It must complete ethics, and sometimes risk and cloud-authorization, steps before deployment [4], [31]. None of those texts states a formula that converts a token rate and an unknown quantity into the value and the obligation. Federal agencies have reported the estimating difficulty in AI procurements [11]. Texas official sources reviewed here do not contain a parallel audit of Texas AI invoices.

VII. Federal Comparison

A. Evidence from GAO, OMB, GSA, and Fiscal Law

GAO-26-107859 reviewed AI acquisition approaches and challenges at the Departments of Commerce, Defense, Homeland Security, and Veterans Affairs, and at GSA, using a nongeneralizable sample. Officials at all five agencies told GAO that understanding AI cost is difficult, particularly because vendor license and service prices are uncertain and markets are moving [11]. GAO also reported requirements definition, intellectual-property and data rights, early and continuous testing, access to technical experts, and traditional timelines as challenge areas [11]. One Army example in the report was a rejected license proposal on the order of $300,000 per vehicle per year. GAO presents that figure as an illustration of license-cost risk in a weapons program, not as a unit price for administrative AI [11]. GAO's recommendations in that report address collection and sharing of acquisition lessons, not a new funding statute [11].

M-25-22 aligns with those challenge areas. It requires or directs, for covered federal agencies and covered AI, cross-functional acquisition teams that include budgeting expertise, pricing-transparency and lock-in provisions, pre-award testing to the extent practicable, contract terms that allow regular performance and risk evaluation, and periodic value evaluation that includes operation and maintenance cost [14]. Sunset criteria are recommended where practicable when cost, need, or performance changes [14]. GAO found that several agencies were not yet organized to feed the GSA lessons repository that M-25-22 contemplates [11].

GSA's consumption procedure is an analogy and a transferable management pattern. It is written for cloud computing on SIN 518210C, not for every AI license, and it is not Texas law. The acquisition letter remains posted as in effect until rescinded or incorporated. A recheck of GSA's page on September 23, 2026, found no rescission. The page's own update marker reads April 16, 2025 [15]. The letter authorizes a requirements task order that is not a FAR 16.503 requirements contract. The ordering officer must state a realistic estimate of consumption and of dollars. That estimate is not a representation that the estimated quantity will be required or ordered. The officer must set a task-order ceiling. All services must be priced at award. The ordering activity may fund incrementally as a bona fide need arises, and it must allot, at award, enough funds to cover the total amount payable on the contract line items actually awarded, and enough funds

TABLE II
TEXAS INSTRUMENTS APPLIED TO VARIABLE AI COST. SECTIONS 2157.068 AND 2157.0685 ARE REQUIREMENTS ONLY FOR A BUYER THOSE SECTIONS COVER.

Instrument	What the text settles	What remains open for metered AI
§ 2157.068(a)(2), fee method	Technology services qualify without regard to fee method [1].	Whether a token rate must be negotiated as a "specific price" or a list discount.
§ 2157.068(e-1)–(e-4); 1 TAC § 212.1	For a covered state-agency purchase, value selects the solicitation count and a $10 million limit. An institution acquiring under Education Code § 51.9335 need not use these bands while that authority is in force [1], [10], [16], [17]. The binding purchase order is the contract for dollar value. A nonbinding encumbrance document is not.	Which of the expected-use figure, the ceiling, or a later invoice total is the amount the agency "may be obligated to pay" when quantity is unknown.
§ 2157.0685	For a covered state agency, an SOW that the contract requires, above $50,000, is unpayable until DIR signs [1], [10], [20].	Which usage changes require a new SOW. The signature rule is not a § 51.9335 duty while that authority is in force.
§§ 403.021, 403.071; eXpendit	Only a timely, encumbered obligation is payable from that appropriation; earmarks lapse [7], [9].	Whether a usage ceiling is a binding encumbrance for its full amount.
APS 018	Over-encumbrance is to be avoided; a binding encumbrance legally obligates funds [8].	The USAS amount for pay-for-what-you-use contracts.
APS 026	Capital riders cap categories, including specified cloud and computer time [24].	Whether a given AI service is in the rider, and at what estimated cost.
§ 2054.008	LBB notice follows a major-information-system contract above a cost figure [6].	How to compute that cost for a meter, including later growth.
§ 2054.003 project definition	Major-project status uses development cost and project form [6].	Whether ongoing consumption is a project at all.
1 TAC §§ 219.11, 219.22, 219.24	Ethics for all covered AI; risk approval before heightened-scrutiny procurement [4].	No pricing, quota, or encumbrance method.
§ 2063.408 (formerly § 2054.0593); 1 TAC § 477.27 (formerly § 202.27)	No covered cloud award or renewal without demonstrated TX-RAMP compliance [29], [31].	Does not cap spend or forecast tokens.
Ch. 552, Bus. & Com. Code	Specified uses are prohibited for covered governmental entities [3].	Does not apply, by its definition, to institutions of higher education; does not regulate price.

when an optional line item is exercised. The contractor must provide usage-tracking tools at no additional cost or notify the ordering officer at 50 percent and 75 percent of each consumption line-item ceiling and of the task-order ceiling, or at other milestones the parties set. The government is not obligated to pay above the ceiling, and the contractor is not obligated to continue performance above it. The ordering officer may use a limitation-of-funds clause for incrementally funded line items. The ordering activity's own funding policies remain in effect under the clause. The procedure cannot be used to acquire, maintain, or modify capital assets [15]. Continuation above the ceiling is a further funded contract action, such as exercising an optional line item, rather than an automatic draw on unobligated funds.

The transferable management pattern, not a Texas duty, is therefore: a realistic estimate that is not a promise to buy the estimated quantity; a contractual ceiling; incremental funding of priced line items; usage monitoring, including alerts; and a funded contract action, with human approval, before payment or performance above the ceiling. A user quota is part of that same recommended pattern. Texas sources reviewed here do not require it. GSA's ordering material describes contract line items that track monthly consumption against funded amounts [15]. The letter's purpose is to pay for actual consumption without obligating every available dollar at the start, while keeping a ceiling and monitoring so fiscal law is observed [15]. This pattern is a federal acquisition design for one GSA schedule rather than Texas law.

B. Similarities, Differences, and Limits on Transfer

Texas and federal buyers share three conditions. Both need an appropriation before they can obligate treasury or federal funds [9], [37]. Both lose the ability to charge an expired appropriation for a cost that was never obligated [7], [38]. Both now have AI governance duties that sit on top of ordinary buying: Texas through Chapter 2054, Subchapter S, and Chapter 219 [2], [4]; federal civilian agencies through M-25-22 and the memoranda it incorporates [14].

The instruments differ. Federal buyers have a published GSA path that defines consumption, requires an estimate and a ceiling, permits incremental funding, and rejects the idea that the estimate is a promise to buy the estimated quantity [15]. Texas buyers have a commodity statute that accepts any fee method and then sorts the purchase by

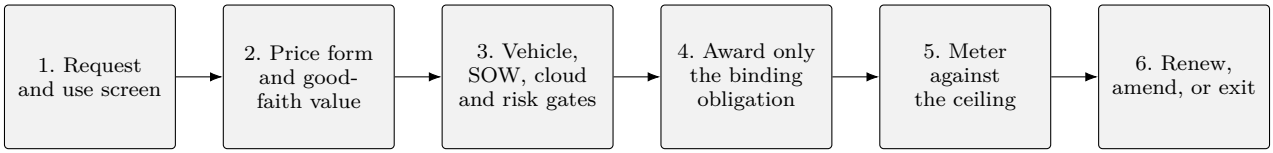


Fig. 2. Lifecycle from request to renewal. Steps 1, 3, and 4 restate Texas legal or administrative gates for covered entities and covered purchases. Steps 2, 5, and 6 add management controls where statewide rules do not specify a metered-price method.

an undefined "value," plus encumbrance rules that require a present legal obligation [1], [8], [9]. Texas buyers also have DIR signature control for covered SOWs and TX-RAMP control for covered cloud [1], [31]. Federal lock-in and pricing-transparency duties in M-25-22 are not Texas duties [14].

Interpretation. A Texas institution may copy the GSA sequence as contract design. That use is a recommended management practice. It is not compliance with GSA's letter and it does not make the letter Texas law. The fiscal step is adapted. GSA requires enough funds at award to cover the amount payable on the line items actually awarded, and it permits further funding as bona fide needs arise [15]. Texas encumbrance rules support recording the present legal obligation and leaving a nonbinding ceiling unencumbered [8], [9], [17]. Incremental funding, in this paper, means a new funded obligation before more is owed. It is not an automatic draw and it is not a Comptroller method for tokens. Section VIII labels the estimate, the quota, the alerts, the human approval, and the renewal memo as recommendations where Texas text is silent. A Texas contract still has to satisfy Chapter 2157, Chapter 403, Chapter 219, and, where applicable, TX-RAMP and Chapter 552. GSA's statement that its procedure does not apply to capital assets is a reminder that a Texas capital-budget rider is a separate limit [15], [24].

What travels is the variable-cost problem itself: a known rate, an unknown quantity, a ceiling that must not be confused with the amount legally owed, and a human decision before spend or performance passes the funded amount. Quotas, alerts, and incremental funding travel only as local management design. What does not travel is Texas-specific: the § 2157.068 value bands and the $10 million commodity limit, the DIR statement-of-work signature, TX-RAMP, Chapter 219's ethics and heightened-scrutiny duties, Chapter 552's coverage, and the encumbrance and lapse rules in Chapter 403 and Comptroller fiscal policy. Another public institution can keep the decision sequence and replace those stops with its own binding rules. That substitution has not been tested. Figure 2 is that sequence. Steps that restate a cited gate are distinguished in the caption from steps that are recommended controls.

VIII. Operational Interim Control Procedure for Metered AI Procurement

The procedure in this section is the paper's operational contribution. It gathers the valuation, vehicle, encumbrance, governance, monitoring, and renewal steps into one sequence that a procurement officer, budget officer, information-technology office, or AI-governance officer can execute. A step is a Texas requirement only when a cited statute, rule, or Comptroller or DIR instruction already imposes it on that buyer and that purchase. Where the reviewed statewide sources do not answer the question, the text says so and then states a recommended interim control. The recommended controls apply duties the cited sources already impose. They are not implied terms of Chapter 2157 or Chapter 403. No Texas invoice sample was available.

Coverage differs by institution. The value bands in § 2157.068(e-1) and the statement-of-work signature duty in § 2157.0685 bind the state agencies those provisions address [1]. Title 1, § 212.1(7) defines "state agency" for the commodity-item rules to exclude an institution of higher education unless the institution has been found out of compliance with Education Code § 51.9337, and the Version 4.0 guide states the DIR threshold procedure for a state agency not including a state institution of higher education [16], [17]. Local governments are not the agencies commanded by § 2157.068(e-1) [1]. Public junior colleges are largely outside Chapter 2054 except for information-security standards and specified shared technology services [6]. Chapter 219's ethics duties, and its heightened-scrutiny duties, use their own coverage [4]. The appropriation and encumbrance steps apply to the spending of appropriated state funds [7], [8], [9]. Section 403.021 applies to a "state agency" as § 403.013 defines it, and that definition lists a university system or an institution of higher education as Education Code § 61.003 defines it [7]. eXpendit likewise states that state agencies and institutions of higher education must report binding encumbrances and payables [9]. Whether a given AI purchase is paid from an appropriation, or from institutional funds outside the General Appropriations Act, is a fact for each buyer. This paper does not decide it, and it does not decide how the lapse rule applies to funds that are not appropriations. The estimate, ceiling, alert, and renewal steps are recommended for each of these buyers.

A. What the September 2026 Sources Settle

Seven questions were checked against the statute, 1 Tex. Admin. Code Chapter 212, 34 Tex. Admin. Code §§ 20.25 and 20.82, Guide Version 4.0, eXpendit, APS 018, general-ledger account 9001, DIR's statement-of-work rule, and DIR's public-cloud description.

Contract value of a meter. Requirement. The Version 4.0 guide tells an agency to use its best judgment of the monetary value of an anticipated information-technology

commodity contract and states that the dollar value is the total over the term, including modifications, renewals, and extensions [16]. Section 20.25(b)(13) defines contract value, for that administrative chapter, as the estimated dollar amount a state agency may be obligated to pay under the contract and all executed and proposed amendments, extensions, and renewals [43]. Section 212.1 counts a fully executed purchase order and a purchase-order change notice as the contract, and it excludes from dollar value a document executed only to encumber funds if that document is not a binding transaction, and a statement of work that has no purchase order [17]. *Unresolved.* No statewide Texas authority located in the reviewed sources specifies this method for converting a token rate, a credit balance, or another meter into that amount. Chapter 219, House Bill 149, and Senate Bill 1964 do not fill the gap [2], [3], [4].

Whether the ceiling is fully encumberable. Requirement. An encumbrance is a present legal obligation under an awarded contract. An earmark is not an encumbrance, and over-encumbrance is to be avoided [8], [9], [25]. *Unresolved.* No statewide Texas authority located in the reviewed sources specifies whether a not-to-exceed ceiling is itself fully encumberable when the contract makes payment due only for measured use.

USAS recording. APS 018 identifies the transaction codes agencies use to record, adjust, and liquidate encumbrances generally [8]. *Unresolved.* No statewide Texas authority located in the reviewed sources specifies this method for a usage-based technology obligation. Absence of that method is not a prohibition on recording the amount the contract legally obligates.

Usage alerts. Unresolved. No statewide Texas authority located in the reviewed sources requires a usage alert or a spend threshold. The 50 percent and 75 percent notices in GSA clause 552.238-199 are federal and apply to consumption cloud ordered under that clause [15].

DIR examples of variable fees. Requirement. Technology services may be commodity items regardless of the fee method [1]. The Version 4.0 guide describes DIR cloud computing services and states that Data Center Services customers enter agreements with DIR to consume infrastructure services. The cooperative-contract value bands and statement-of-work signature process do not apply to services procured through that program under contracts signed by the contracted vendors [16]. DIR's public-cloud page describes variable-priced capacity and reserved instances among those services [44]. *Unresolved.* No statewide Texas authority located in the reviewed sources specifies a token, credit, or generative-AI valuation example.

Amendments. Requirement. When the DIR contract requires a statement of work to initiate services above $50,000, the work is not valid, and money may not be paid, until DIR signs [1], [42]. DIR's signature does not verify that funding is appropriate or sufficient [42]. A purchase-order change notice is itself a contract for dollar-value purposes [17]. The guide calls a written change-control process a best practice and gives a purchase-order change notice or a signed amendment as the ordinary instruments [16]. DIR's 2022 statement-of-work application guide tells users to submit amendments or change orders, including scope changes regardless of cost, to DIR [45]. *Unresolved.* No statewide Texas authority located in the reviewed sources specifies which changes in consumption, rate, model, or scope meet that test. The portal guide does not define ordinary use inside an already signed ceiling as a scope change.

Later guidance. The 2025–2026 instruments add ethics, use limits, risk assessment, and a restatement of procurement practice [2], [3], [4], [16], [32]. They do not adopt the GSA consumption design as Texas law, and they do not state a USAS treatment for a usage ceiling.

B. Steps

The officer completes the steps in order. Figure 3 is the branch rule. Table III is the authority map.

1. Classify the price. Recommended control. Record one of four forms: subscription, in which a term, seat, or bundle is priced; consumption or tokens, in which a rate is priced and quantity is open until used; prepaid credits, in which a pool is purchased; or hybrid, in which a committed base is paired with metered overage. A subscription is placed in a value band by its stated price if that price is the contract value [1]. The meter, the credit pool, and the hybrid overage use the value rules below.

2. Screen the use before solicitation. Requirement. If Chapter 552 covers the buyer, reject a prohibited governmental use [3]. *Requirement.* If the system is heightened-scrutiny, complete the written risk assessment and obtain the AI risk officer's decision before procurement or deployment, and again at a material change [2], [4]. *Requirement.* If the service is covered cloud, obtain evidence of TX-RAMP compliance before award or renewal, and require the vendor to maintain it [29], [31]. *Requirement.* Chapter 219's ethical duties, including human review, accuracy monitoring, disclosure, security, and a contractual vendor duty, apply when that chapter covers the buyer, including systems that are not heightened-scrutiny [4]. These gates do not set a price.

3. Build two quantities. Recommended control. Prepare a good-faith expected-usage estimate: unit rate, expected quantity, term, DIR administrative fee if the buy is through DIR, and the resulting dollars. The Version 4.0 guide tells agencies to develop a cost estimate in procurement planning, in good faith, from price lists, research, standardized methods, or historical spend, and it cites the delegated-purchase estimate rule for that planning instruction [16], [21]. The rule's own text governs delegated purchases. It does not, by that text, govern a DIR cooperative order [21]. For a DIR order, the guide's best-judgment instruction is the closer administrative text [16]. *Recommended control.* Also prepare a higher-usage planning estimate, or stress case, using a stated driver such as peak historical use, a staffing surge, or a higher-priced

model. *Requirement.* Keep the arithmetic in the procurement file for a covered procurement [22]. If expected use is known reasonably well, the expected-usage estimate is the working forecast and the stress case is the planning ceiling's input. If use is highly uncertain, do not treat the expected-usage figure as a stable contract value. Use the stress case as the ceiling input, shorten the term, or use the statutory sandbox for a time-limited pre-deployment test if DIR has placed the buy in that program [2]. A short order with a low ceiling, outside the sandbox, is a management substitute. It does not carry the sandbox statute's regulatory-flexibility language [2].

4. Place the purchase on a vehicle. Requirement. For a covered state-agency commodity purchase, compare the amount the agency may be obligated to pay with the bands in § 2157.068: direct award at or below $50,000; at least three vendors above $50,000 through $1 million; at least six, or all if fewer are listed, above $1 million through $10 million; and no cooperative-contract commodity award above $10 million except a coordinated bulk purchase [1], [16]. Above $5 million through $10 million, the Comptroller-designated request-for-offers method remains available [1]. *Interpretation.* An institution of higher education acquiring under Education Code § 51.9335 is not required to use these bands while that authority remains in force [10], [23]. The bands remain a requirement for a state-agency purchase that § 2157.068 covers, and they apply again if § 51.9337 suspends the institution's authority. *Requirement.* If the service is DIR Data Center Services under a contract signed by the contracted vendor, the guide states that those cooperative-contract bands and the statement-of-work signature process do not apply [16]. *Requirement.* A delegated open-market purchase uses the thresholds and the good-faith estimate rule in 34 Tex. Admin. Code § 20.82, and an item on a cooperative contract is outside that delegation except below the contract's minimum order [21]. *Recommended control.* If the expected-usage dollars and the stress-case ceiling fall in different bands, or if either figure crosses $50,000, $1 million, $5 million, or $10 million, stop solicitation and write which figure will be used as the amount the agency may be obligated to pay, and why. *Unresolved.* No statewide Texas authority located in the reviewed sources specifies this method of choosing between those figures. Using the higher figure for the solicitation count is a conservative management control. It is not a holding that the statute requires the ceiling. If the figure used would exceed $10 million on a cooperative commodity contract, use a lawful vehicle that can carry that value, or reduce the term or ceiling until the obligated amount fits.

5. Separate the ceiling from the encumbrance. Recommended control. State a not-to-exceed ceiling in the purchase order. The guide's general policy is that indefinite contracts are generally prohibited [16]. A ceiling gives the order a number. Set it at or below the vehicle limit and at or below funds the buyer can lawfully use. *Requirement.* Encumber the amount the awarded order legally obligates, in the appropriation year of that obligation [7], [8], [9]. *Recommended control.* Write on the order which amount is the present payment duty. If counsel and the appropriation-control officer conclude that the state owes only measured use, encumber that present duty, or the minimum the order makes payable now, and do not certify the unused ceiling as a binding encumbrance. *Unresolved.* No statewide Texas authority located in the reviewed sources specifies this method. The written conclusion belongs in the contract file because APS 018 and eXpendit do not make it for cloud or AI. *Recommended control.* If they conclude that the order obligates the full ceiling at award, encumber that ceiling and do not also reserve it a second time as an earmark. Account 9001 is for commitments evidenced by a contract, purchase order, or administrative action [25]. A forecast with no present duty does not meet that description on the text reviewed here.

6. Test the appropriation year. Requirement. Spend treasury funds only under an appropriation. The Comptroller states that appropriations drawn from the treasury are constitutionally limited to two-year terms and that General Appropriations Act items are commonly available for one fiscal year [9]. A claim must be presented within the statutory window, ordinarily two years after the fiscal year of the appropriation [7]. Unencumbered nonconstruction balances lapse, and a prior-year claim is payable from that appropriation only if it was encumbered for the claim before the appropriation expired [7]. *Recommended control.* Match the order term to the appropriation that will pay it. If the term will cross a fiscal year, the guide states that the solicitation must include a termination-for-non-appropriations clause, which the guide treats as an essential provision [16]. *Recommended control.* When the appropriation is near expiration, stop new consumption that would be paid from that appropriation unless a binding encumbrance is already in place or a still-living appropriation is identified and obligated. An earmark will not keep the balance alive [7], [9]. The last-quarter expenditure limit that eXpendit restates is a further timing constraint on General Appropriations Act funds, with an exemption that eXpendit describes for expenditures contracted in earlier quarters [9]. *Unresolved.* No statewide Texas authority located in the reviewed sources specifies whether a later metered invoice under an earlier order is an expenditure contracted for in that earlier quarter.

7. Complete the instruments the vehicle requires. Requirement. Where § 2157.0685 covers the buyer and the master contract requires a statement of work above $50,000, consult DIR before sending it to a vendor, include scope, milestones, deliverables, schedule, and acceptance criteria, and obtain DIR's signature before payment [1], [42]. *Interpretation.* That signature condition is a state-agency condition. DIR addresses statement-of-work review to state agencies [20], [23]. It is not a requirement for an institution acquiring under Education Code § 51.9335 while that authority is in force [10]. Do not treat that signature as a funding certification [42]. *Requirement.* Put Chapter 219's vendor duties in the contract when that chapter applies [4]. *Recommended control.* Also state the

rate card, the expected-usage estimate and the statement that it is not a promise to consume that quantity, the stress case, the ceiling, the appropriation year, the amount that is the present obligation, the named budget owner, and the official who may amend, suspend, or stop the service. The statute requires deliverables and performance specifications when a statement of work is the instrument. It does not require this usage block. The block is how the required value and obligation reviews become auditable.

8. Monitor, and escalate before the ceiling or the appropriation fails. Recommended control. Set a quota and internal alerts at 50 percent and 75 percent of the encumbered, funded amount. The percentages are taken from the GSA analogy. They are not a Texas duty [15]. Name the person who receives the alert and who must approve any further funding. At the first alert, compare the run rate with the expected-usage estimate and with the remaining appropriation. At the second alert, or when actual use otherwise approaches the quota, the ceiling, the encumbered amount, or the appropriation, that official chooses one path before further use is incurred: a purchase-order change notice or other amendment that raises the ceiling and creates a new obligation against funds still available; a new solicitation or vehicle if the new total crosses a band the original award did not cover; suspension of new use; or termination. Incremental funding is that new obligation. It is a recommended practice, not an authorization to spend past the encumbrance. *Requirement.* If the change is a statement of work that the DIR contract requires, DIR must sign before payment on that work [1], [42]. *Recommended control.* Route a change in model, rate, scope, vendor, or funded ceiling through the guide's change-control path, and through DIR's amendment path when the buy uses a DIR statement of work [16], [45]. *Unresolved.* No statewide Texas authority located in the reviewed sources specifies that a fluctuation inside an already signed ceiling is such a change. *Requirement.* A heightened-scrutiny material change still requires a fresh risk assessment whether or not the invoice changes [2], [4]. *Requirement.* A covered cloud renewal still requires current TX-RAMP evidence [31].

9. Renew only through a gate. Requirement. Before renewal or a material change, repeat every requirement that applies: appropriation availability, TX-RAMP for covered cloud, the Chapter 219 evaluation record, and, for heightened-scrutiny systems, the risk assessment [4], [9], [31]. *Recommended control.* The renewal memo also records cost against the expected-usage estimate, performance quality, security exceptions, portability and the cost of leaving the vendor, and whether the next appropriation can carry the next ceiling. Section 2157.003 already makes compatibility and capacity to expand best-value factors when that section applies [1]. Section 219.11 requires enough transparency to understand how a covered system operates [4]. M-25-22's pricing-transparency and lock-in instructions are federal and are used here only as a checklist for the recommended memo [14]. A failed gate stops renewal. It does not, by itself, cancel a current obligation that has already been incurred.

10. If no defensible figure can be stated, do not award. Recommended control. If the organization cannot document a good-faith amount that it may be obligated to pay, or cannot state the amount that is a binding encumbrance, stop the award. Record the missing input, who reviewed it, and the date. *Unresolved.* No statewide Texas authority located in the reviewed sources specifies this stop as a statutory prohibition. The stop is the interim control that keeps the officer from inventing a value band or certifying an encumbrance the file cannot support. A prohibited use, a missing heightened-scrutiny approval, a missing TX-RAMP demonstration, or an unsigned statement of work that the law makes a condition of validity or payment remains a legal stop for the buyers those rules cover [1], [3], [4], [31].

C. Decision Rule

Apply the first matching branch in Fig. 3.

Audit trail. Requirement. Retain the need statement, the value method, solicitation evidence, any signed statement of work, the risk assessment, cloud-authorization evidence, the order, the encumbrance record, invoices, amendments, and the evaluation notes for procurements and systems the cited rules cover [4], [7], [22]. *Recommended control.* Naming a budget owner, and showing divisional use against a central AI account, gives the program that drives consumption a quantity signal. Chargeback is not required by the sources reviewed here.

D. Optimization-Support Layer

The model in Section IX may be called from this procedure. It does not change a step or an authority.

Recommended control. After step 3 has produced an expected-use figure and a stress figure, and after step 5 has separated the ceiling from the amount that is presently payable, an officer may use the model in Section IX to recommend a subscribed quantity, a credit pool, a pre-authorized token budget, and one configuration from the permissible set. *Recommended control.* The same model may recommend, after demand is observed, whether to serve with tokens already funded, to request incremental funding, or to leave demand unmet. *Requirement.* The recommendation is usable only if the legal stops in this section are already satisfied for that buyer and that purchase. A prohibited use, a missing heightened-scrutiny decision, a missing TX-RAMP demonstration, or an unsigned statement of work that the law makes a condition of payment still stops the file [1], [3], [4], [31].

Unresolved. The model has no variable for statutory contract value and no variable for a binding encumbrance. Sections VIII and the budgeting sections leave those amounts unresolved, and those determinations remain outside the solver. *Recommended control.* If counsel and the appropriation-control officer have not written which figure is the present obligation, the officer does not enter the model's outlay as an encumbrance. *Recommended control.*

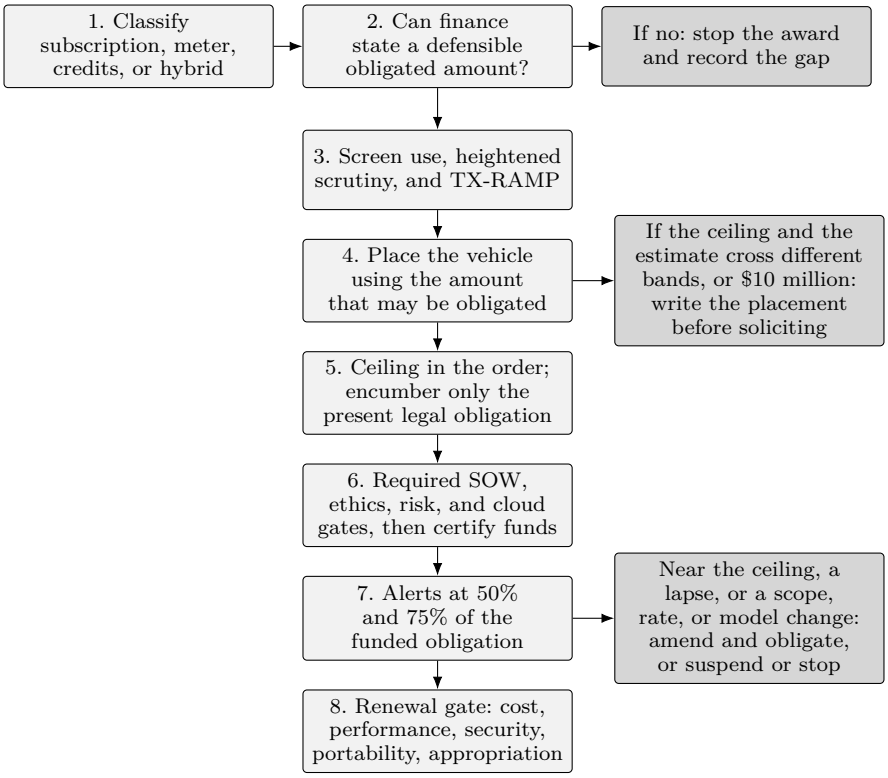


Fig. 3. Decision rule for a metered AI purchase. Shaded boxes are stops or written determinations. The ceiling, quota, 50% and 75% alerts, incremental funding, and human approval are recommended controls taken from a GSA analogy. They are not Texas law. Use, risk, cloud, statement-of-work, and encumbrance stops are requirements only when the cited Texas rule covers the buyer and the purchase.

If the model's quantities and the officer's determination differ, the officer's determination controls. The model cannot authorize a person the institution has not already authorized, and it cannot treat a recommended ceiling as the figure § 2157.068 uses for a value band.

The first model's administrative content is the configuration choice. Hybrid, committed, and meter configurations are three ways to package steps the procedure already contains. Cycle time and staff effort for those packages are not stated in the Texas sources reviewed above. Where Section X uses hours and days, it labels them as synthetic assumptions. Section IX-E then schedules the activities inside a package. That schedule is a simulated file; it creates no approval the statute does not create and preserves every gate the retained precedence marks as mandatory.

IX. Business-Process-Integrated Risk-Aware Optimization

This section writes the integrated planning model that Section VIII may call. The first model minimizes a weighted sum of expected procurement expenditure, the conditional value-at-risk of that expenditure, a penalty on unmet demand, and the assumed staff hours and cycle days of the chosen configuration. It is a two-stage linear recourse program [12] with the scenario linearization of conditional value-at-risk [13], and it embeds the officer-supplied funding cap and human-authorization limit as constraints so that the recommended package and quantities are budgeting-feasible by construction. In that first model, business process optimization (BPO) enters as the first-stage choice among the three configurations the procedure already permits. The extended model in Section IX-E schedules activities and activates the finite-support Wasserstein layer, completing the joint architecture. The program recommends a plan inside an officer-supplied cap; the legal characterization of that cap and of the outlay remains with the responsible officers.

A. Sets, Data, and Decisions

Let S be a finite set of demand scenarios and let $\pi_s = 1/|S|$. Let $K = \{\text{hybrid}, \text{committed}, \text{meter}\}$. For each $k \in K$, the data T_k and E_k are assumed days and person-hours. The prices p^{sub}, p^{cred}, and p^{tok} are positive unit prices. The officer supplies a planning cap $\Gamma > 0$. The runs also impose an exogenous cap of $10 million, the figure in Tex. Gov't Code § 2157.068(e-2), and use the smaller of the two caps. That modeling choice does not decide whether the statute looks to a ceiling, to expected use, or to later invoices. A fraction $\varphi \in (0, 1)$ limits how much of the uncommitted remainder may be pre-authorized for tokens. The fraction represents the incremental-funding control recommended in Section VIII. It is not a statutory ratio. Authorization, if used, adds time t_a and effort e_a.

TABLE III
DECISION POINTS FOR METERED AI PROCUREMENT. A CELL IN THE TEXAS-REQUIREMENT COLUMN APPLIES ONLY WHEN THE CITED PROVISION COVERS THAT BUYER AND THAT PURCHASE.

Decision point	Texas requirement	Unresolved issue	Recommended interim control	Authoritative source
Price form	Technology services may qualify regardless of fee method.	No token or credit definition.	Classify subscription, consumption, prepaid credits, or hybrid before solicitation.	Gov't Code § 2157.068(a) [1]
Expected use is reasonably known	Covered procurements need a documented cost estimate. Delegated purchases need a good-faith estimate.	The delegated-purchase rule does not, by its text, govern a DIR order.	Use the expected-use dollars as the forecast and the stress case only to size the ceiling.	34 TAC § 20.82; Guide v4.0 [16], [21]
Use is highly uncertain	Best judgment of the amount that may be obligated, for a DIR commodity buy.	No statewide method specifies the stress-case formula.	Size the ceiling from the stress case, shorten the term, or use the sandbox if DIR places the buy there.	Guide v4.0; § 2054.706 [2], [16]
Ceiling crosses a threshold	Bands and the $10 million cooperative limit, for the state agencies they cover. Data Center Services under vendor-signed contracts are outside those bands.	Whether value is the ceiling, expected use, or later invoices.	Write the figure used. If it crosses a band or $10 million, place the buy before soliciting.	§ 2157.068; 1 TAC § 212.1; Guide v4.0 [1], [16], [17]
Ceiling versus encumbrance and USAS	Encumbrance requires a present legal obligation. Avoid over-encumbrance. Nonbinding fund documents are excluded from DIR dollar value.	No statewide method specifies the USAS amount for pay-for-use.	Encumber only the amount the order presently obligates. Record the theory. Do not certify a forecast.	APS 018; eXpendit; GL 9001; 1 TAC § 212.1 [8], [9], [17], [25]
Appropriation near expiration	Unencumbered nonconstruction balances lapse. Prior-year payment requires a timely encumbrance.	Whether a later metered invoice was "contracted for" in an earlier quarter.	Stop new use unless a living appropriation is obligated. Use a non-appropriations clause if the term crosses a fiscal year.	Gov't Code §§ 403.021, 403.071; Guide v4.0 [7], [9], [16]

First-stage decisions, chosen before demand is known, are

$$q^{\mathrm{sub}} \ge 0, \quad q^{\mathrm{cred}} \ge 0, \quad R \ge 0, \qquad y_k \in \{0,1\}, \quad \sum_{k\in K} y_k = 1. \tag{9}$$

R is dollars of token service pre-authorized at award. It is not an encumbrance entry. Second-stage decisions, chosen after d_s is known, are

$$s_s^{\mathrm{sub}},\ s_s^{\mathrm{cred}},\ u_s^{\mathrm{tok}},\ m_s,\ z_s \ge 0, \quad a_s \in \{0,1\}. \tag{10}$$

Equations (9) and (10) are the two stages. The committed outlay and the scenario expenditure are

$$W = p^{\mathrm{sub}} q^{\mathrm{sub}} + p^{\mathrm{cred}} q^{\mathrm{cred}}, \qquad L_s = W + p^{\mathrm{tok}} u_s^{\mathrm{tok}}. \tag{11}$$

Equation (11) names the prepaid part W and leaves the token bill inside L_s. W is the amount the unit prices make payable at award because the units are prepaid or subscribed. The program does not relabel W as a binding encumbrance. An officer who concludes that the order obligates a different amount records that conclusion outside the program, as step 5 of Section VIII already requires.

B. Objective

With $\lambda \in [0,1]$, tail level $\beta \in (0,1)$, unmet-demand weight w_m, staff weight r_E, and delay weight r_T, the program minimizes

$$\begin{aligned} &(1-\lambda)\sum_{s\in S} \pi_s L_s \\ &\quad + \lambda\left(\eta + \frac{1}{(1-\beta)|S|}\sum_{s\in S} v_s\right) \\ &\quad + w_m \sum_{s\in S} \pi_s m_s \\ &\quad + \sum_{k\in K} y_k (r_E E_k + r_T T_k) \\ &\quad + \sum_{s\in S} \pi_s a_s (r_E e_a + r_T t_a). \end{aligned} \tag{12}$$

TABLE III
Decision points for metered AI procurement (continued)

Decision point	Texas requirement	Unresolved issue	Recommended interim control	Authoritative source
Use approaches the ceiling	For a buyer covered by § 2157.0685, payment on a required statement of work waits for DIR's signature. A purchase-order change notice is a contract.	No statewide quota, alert percentage, incremental-funding method, or human-approval duty.	Recommended only: a quota, alerts at 50% and 75% of the funded amount, and named human approval before incremental funding. The percentages are a GSA analogy, not Texas law.	§ 2157.0685; 1 TAC §§ 212.1, 212.41; GSA analogy [1], [10], [15], [17], [42]
Model, rate, scope, or vendor change	Material change refreshes heightened-scrutiny review. For a buyer covered by § 2157.0685, DIR must sign a required statement of work.	Portal guidance says scope changes go to DIR and does not define in-ceiling use as scope.	Route the change through change control before payment on the new terms.	1 TAC §§ 212.41, 219.22; Guide v4.0 [4], [10], [16], [42], [45]
Heightened-scrutiny AI	Written risk assessment and risk-officer decision before procurement and at material change.	The assessment does not value tokens.	Finish the assessment before solicitation. Repeat it if the model or use materially changes.	1 TAC § 219.22; S.B. 1964 [2], [4]
TX-RAMP applies	No covered cloud award or renewal without demonstrated compliance.	The program does not cap spend.	Put the certification in the award file and again in the renewal file.	Gov't Code § 2063.408; 1 TAC § 477.27 [29], [31]
No defensible value or encumbrance	The officer still must place a covered buy in a lawful vehicle and encumber only a real obligation.	No statewide source orders a stop in those words.	Do not award until both figures are written and approved.	§ 2157.068; APS 018 [1], [8], [16]
Renewal	New appropriation, TX-RAMP if covered, Chapter 219 evaluation, and heightened-scrutiny reassessment if required.	No Texas duty to exit because a rate moved.	Renew only if cost, performance, security, portability, and appropriation all pass.	§ 2157.003; 1 TAC §§ 477.27, 219.11, 219.22 [1], [4], [31]

The auxiliary constraints are $v_s \geq L_s - \eta$ and $v_s \geq 0$. The term in parentheses is the conditional value-at-risk of procurement expenditure, not of a legal penalty and not of unmet demand. Unmet demand and administrative effort stay outside that tail term so that a reduction in recorded expenditure cannot be mistaken for a reduction in unmet service or in staff work.

C. Constraints

Equation (13) requires service and unmet demand to add up to realized demand. Equation (14) keeps subscription and credit service inside the purchased quantities:

$$s_s^{\mathrm{sub}} + s_s^{\mathrm{cred}} + u_s^{\mathrm{tok}} + m_s = d_s, \tag{13}$$

$$0 \leq s_s^{\mathrm{sub}} \leq q^{\mathrm{sub}}, \qquad 0 \leq s_s^{\mathrm{cred}} \leq q^{\mathrm{cred}}. \tag{14}$$

Equations (15), (16), and (17) are the configuration gates. A configuration that is not selected cannot buy the instruments it excludes:

$$q^{\mathrm{sub}} \leq M(y_{\mathrm{hybrid}} + y_{\mathrm{committed}}), \tag{15}$$

$$q^{\mathrm{cred}} \leq M(y_{\mathrm{hybrid}} + y_{\mathrm{committed}}), \tag{16}$$

$$R \leq \Gamma(y_{\mathrm{hybrid}} + y_{\mathrm{meter}}), \tag{17}$$

with M large enough to admit any purchase that fits under Γ. Write $\Gamma^{\cap} = \min(\Gamma, \Gamma_{\mathrm{stat}})$ for the tighter of the officer-supplied cap and the exogenous cap. The symbol shortens the displays below. It does not change the constraint. Equation (18) keeps the pre-authorized token budget inside the fraction φ of the cap that is not already committed, and it keeps the committed outlay under the tighter cap:

$$R \leq \varphi(\Gamma - W), \qquad W \leq \Gamma^{\cap}. \tag{18}$$

Equations (19) and (20) limit token service to the pre-authorized budget plus incremental funding, and they keep that funding inside the tighter cap. Incremental

funding also requires authorization and a meter-capable configuration:

$$p^{\text{tok}} u_s^{\text{tok}} \le R + z_s, \tag{19}$$
$$z_s + W + R \le \Gamma^{\cap}, \tag{20}$$
$$z_s \le \Gamma^{\cap} a_s, \tag{21}$$
$$a_s \le y_{\text{hybrid}} + y_{\text{meter}}. \tag{22}$$

Equation (21) is the human-authorization restriction. Equation (22) keeps that authorization inside a configuration that still contains the pre-award gates. The gates are not re-derived by the solver. They are a condition of membership in K. If a gate that Texas law requires for the buyer has not been passed, the file never reaches this program.

The stated planning ceiling used in the numerical section is the officer-supplied cap Γ. The program does not search for a smaller ceiling, because under these constraints a lower ceiling only removes feasible service and does not reduce L_s except by cutting service. Reporting Γ as the planning ceiling is not a recommendation to encumber Γ. Step 5 still requires a separate written theory of the present obligation.

D. Scope of the Recommendation

The program does not return a contract value for § 2157.068, a Uniform Statewide Accounting System encumbrance, or a conclusion on whether a not-to-exceed ceiling is fully encumberable when payment is due only for measured use. Those questions remain unresolved in Section VIII. A feasible solution is a recommendation and still requires the human, budgetary, and contractual checks in that section. Mathematical infeasibility is not a statutory prohibition. A preprint on agentic and multi-agent systems treats those systems as bounded and monitored, and it argues that high-impact settings need stronger oversight before deployment [46]. That review informs the oversight design rather than testing this program. The authorization variable remains a recommended control rather than an approval issued by the solver.

E. Extended Activity Network and Finite-Support Wasserstein Layer

The configuration model selects one of three packages and does not schedule the activities inside a package. The schedule below is a resource-constrained project model on a directed acyclic graph. The distributional layer is the linear programming dual of a 1-Wasserstein transportation problem on a fixed finite support [18]. These constructions supply the process-timing and demand-ambiguity blocks of the integrated model. The contribution is their coupling with package choice, CVaR, and the officer-supplied funding and authorization limits in a single optimization problem. The implemented program is the mixed-integer linear program stated below, solved with HiGHS. A Wasserstein ball over every distribution on the real line is a different program: on an unbounded support the inner supremum is a semi-infinite constraint. The finite-support restriction is exact only for distributions that put their mass on the chosen support.

1) Activity Network: The simulated file has eleven activities. Base durations, in days, are classify 2, estimate 3, screen 5, solicit 4, approve 2, credit 4, token 4, encumber 2, award 2, alerts 2, and release 1. A model workday contains six person-hours, so the labor content of an activity is six times its realized duration. Overlap reduces the makespan. It does not remove labor. Credit sizing is present only in the committed and hybrid packages. Token pre-authorization is present only in the meter and hybrid packages. The screen is included because the simulated case is a covered cloud buy, so the ethics, risk, and TX-RAMP checks apply. The approval activity is the recommended human-ceiling control rather than a statutory office.

Let $\mathcal{A}(k)$ be the active set of package k. The AS-IS mode chains that set in the fixed order classify, estimate, screen, solicit, approve, credit, token, encumber, award, alerts, release, skipping inactive activities. With one person, every feasible schedule has makespan equal to the sum of the active durations, because a gap would not finish the file earlier. The TO-BE mode keeps only the partial order

$$\begin{aligned}
&\text{classify} \prec \text{estimate},\\
&\text{classify} \prec \text{screen} \prec \text{solicit},\\
&\text{estimate} \prec \text{approve}, \quad \text{solicit} \prec \text{approve},\\
&\text{approve} \prec \text{credit} \prec \text{encumber},\\
&\text{approve} \prec \text{token} \prec \text{encumber},\\
&\text{approve} \prec \text{alerts},\\
&\text{encumber} \prec \text{award} \prec \text{release},\\
&\text{alerts} \prec \text{release}.
\end{aligned} \tag{23}$$

An arc is dropped when either end is inactive. Equation (23) is a modeling choice among activities whose relative order the cited statutes do not state. It is not a finding that Texas law permits those activities to run together.

Let $d_i(\omega)$ be the duration of activity i in duration scenario ω, drawn independently of demand. Equation (24) is the earliest-start recursion. For each retained arc (i, j), the earliest-start time S_i and the makespan T satisfy

$$\begin{aligned}
S_j &\ge S_i + d_i(\omega),\\
T &\ge S_i + d_i(\omega).
\end{aligned} \tag{24}$$

Staff capacity n^{st}, distinct from the configuration set K, limits how many activities may be in progress. At $n^{\text{st}} = 1$ the TO-BE and AS-IS makespans coincide. At $n^{\text{st}} = 2$ the earliest-start list with a longest-path priority matched the resource-unconstrained critical path on every draw used in the check below, so the critical path was optimal for this graph and this capacity. The case is one file; waiting induced by other files on a shared staff pool is evaluated in Section X-F.

Duration draws are held apart from demand draws. The procurement program uses the mean makespan $\bar{T}_{k,g}$

and the mean labor $\bar{E}_{k,g}$ of package k and mode $g \in \{\text{asis}, \text{tobe}\}$. The 0.90 quantile of the makespan sample is reported as an administrative tail. It is not added to financial CVaR.

2) Package Choice: The joint program selects one pair (k, g) with a binary $y_{k,g}$ and $\sum y_{k,g} = 1$. Quantities remain the first-stage variables in (9), with the package binaries replacing the three configuration binaries. The administrative term in (12) becomes

$$\sum_{k,g} y_{k,g}\big(r_E \bar{E}_{k,g} + r_T \bar{T}_{k,g}\big). \tag{25}$$

The schedule does not depend on $(q^{\text{sub}}, q^{\text{cred}}, R)$ once the package is fixed, so computing the schedule first and then optimizing quantities is an exact block separation, not a Benders cut. The instances below solve in well under a second.

The admin-first sequential rule minimizes (25) first. Ties are broken by the documented order committed before meter before hybrid, and TO-BE before AS-IS. It then optimizes quantities with that package fixed. A finance-first sequential rule, used only in the comparison of Section X-F2, ranks packages by the financial block and then applies the TO-BE mode. The joint (integrated) rule optimizes (25) and the financial terms together.

3) Finite-Support Wasserstein Layer: Let $\hat{\mathbb{P}} = N^{-1}\sum_{i=1}^{N} \delta_{\xi_i}$ be the empirical measure on the training demands ξ_i, and let $\Xi = \{\xi_1, \ldots, \xi_N\}$. For $\varepsilon \geq 0$, Equation (26) is the ambiguity set:

$$\mathcal{P}_\varepsilon = \Big\{P : \ \text{supp}(P) \subseteq \Xi, \ W_1(P, \hat{\mathbb{P}}) \leq \varepsilon\Big\}, \tag{26}$$

where W_1 is the 1-Wasserstein distance with ground cost $|\xi - \xi'|$. This is the atom-restricted specialization of the ball of [18]: mass moves among the training atoms. The radius used in the base robust run is half the mean absolute deviation (MAD) of the training sample from its median, and it is treated as a sensitivity parameter. A radius calibrated to the finite-sample guarantee of the cited work would require a moment constant that this study does not estimate from external data. The grid $\varepsilon \in \{0, 0.25, 0.50, 1\} \times \text{MAD}$ is the sensitivity analysis. At $\varepsilon = 0$ the only feasible coupling is the identity, and the robust model reproduces the nominal model.

Let L_ℓ be the recourse expenditure at atom ξ_ℓ, and let $v_\ell \geq (L_\ell - \eta)_+$. For a fixed recourse, the worst-case expectation and the worst-case CVaR on $\mathcal{P}_\varepsilon$ are the duals

$$\sup_{P \in \mathcal{P}_\varepsilon} \mathbb{E}_P[L] = \min_{\theta_E \geq 0} \ \theta_E \varepsilon + \frac{1}{N}\sum_i s_i^E, \tag{27}$$

$$\sup_{P \in \mathcal{P}_\varepsilon} \text{CVaR}_\beta(L) = \min_{\eta, \theta_C \geq 0} \ \eta + \frac{1}{1-\beta}\Big(\theta_C \varepsilon + \frac{1}{N}\sum_i s_i^C\Big), \tag{28}$$

subject to $s_i^E \geq L_\ell - \theta_E|\xi_\ell - \xi_i|$ and $s_i^C \geq v_\ell - \theta_C|\xi_\ell - \xi_i|$. Both programs are linear in the recourse variables. Together with the package binaries and the authorization binaries, the extended model is a mixed-integer linear program. The objective uses $(1-\lambda)$ times (27) plus λ times (28), plus the empirical mean of unmet demand and (25). The two suprema are separate. Their sum upper-bounds the supremum of the weighted sum under one distribution. Unmet demand is left on $\hat{\mathbb{P}}$ so that a movement of demand mass cannot be counted both as expenditure risk and as service risk.

When the selected package prepays a credit block and buys no tokens, L_ℓ is the same at every atom. Corollary 1 shows that both expenditure-side suprema then equal that constant for every radius. Corollary 2 defines the threshold at which a nondecreasing worst-case score meets a fixed competing score. The identity of the robust and nominal solutions therefore reflects the constant-loss structure of a prepaid plan and not an inactive layer. The layer remains active whenever expenditure takes two or more values on the training atoms.

4) Instrument-Activated Ambiguity: This subsection analyzes how the finite-support dual of [18] behaves when its loss is the recourse expenditure of a procurement plan. The two corollaries below follow from the nested Wasserstein ball and that dual. They explain the prepaid invariance observed in the validated solves and identify the radius at which the preferred package changes. The analysis is carried out on the atom-restricted ambiguity set of Section IX-E, leaves the validated baseline purchase of 8,399.5 credits unchanged, and uses no new demand draw.

Corollary 1 (Radius geometry): Let $\rho_E(\varepsilon) = \sup_{P \in \mathcal{P}_\varepsilon} \mathbb{E}_P[L]$ for a loss fixed on the training atoms, and let $\rho_C(\varepsilon)$ be the worst-case conditional value-at-risk of that same loss over $\mathcal{P}_\varepsilon$. The empirical masses are $1/N$.

1) If $\varepsilon \leq \varepsilon'$, then $\mathcal{P}_\varepsilon \subseteq \mathcal{P}_{\varepsilon'}$, so $\rho_E(\varepsilon) \leq \rho_E(\varepsilon')$. The same nesting gives $\rho_C(\varepsilon) \leq \rho_C(\varepsilon')$.
2) ρ_E is concave. If every pair of distinct atoms has positive ground cost, then $\rho_E(0) = N^{-1}\sum_\ell L_\ell$, and $\rho_E(\varepsilon) \leq \max_\ell L_\ell$.
3) If $L_\ell = c$ for every atom, then $\rho_E(\varepsilon) = c$ and $\rho_C(\varepsilon) = c$ for every $\varepsilon \geq 0$ and every $\beta \in (0, 1)$. The dual multiplier $\theta_E = 0$ is optimal. Positive separation of the atoms is not required.
4) If L is not constant and $L_\ell \neq L_i$ implies $|\xi_\ell - \xi_i| > 0$, then $\rho_E(\varepsilon) > N^{-1}\sum_\ell L_\ell$ for every $\varepsilon > 0$.

Under that separation, $\rho_E(\varepsilon)$ equals the empirical mean for every $\varepsilon \geq 0$ if and only if L is constant on the atoms. Constancy is a property of the recourse expenditure $L_\ell = W + p^{\text{tok}} u_\ell^{\text{tok}}$. A prepaid plan with zero token use is a sufficient condition.

Proof. Nesting is the definition of (26): a coupling of cost at most ε is a coupling of cost at most ε'. The supremum of a fixed objective over a larger set cannot fall. This holds for the linear objective $\mathbb{E}_P[L]$ and for CVaR_β, which is defined on the law of L.

For concavity, the 1-Wasserstein distance is convex in its first argument. If $P_1 \in \mathcal{P}_{\varepsilon_1}$ and $P_2 \in \mathcal{P}_{\varepsilon_2}$, the mixture $tP_1 + (1-t)P_2$ lies in $\mathcal{P}_{t\varepsilon_1 + (1-t)\varepsilon_2}$. Expectation is linear in the measure, so

$$\rho_E(t\varepsilon_1 + (1-t)\varepsilon_2) \geq t\rho_E(\varepsilon_1) + (1-t)\rho_E(\varepsilon_2).$$

The same concavity for ρ_C uses the minimax representation on the fixed finite support. The set $\mathcal{P}_\varepsilon$ is a nonempty compact convex polytope. The map $(\eta, P) \mapsto \eta + (1-\beta)^{-1}\mathbb{E}_P[(L-\eta)_+]$ is convex in η and linear in P, so

$$\rho_C(\varepsilon) = \min_{\eta} \left(\eta + \frac{1}{1-\beta} \sup_{P \in \mathcal{P}_\varepsilon} \mathbb{E}_P[(L-\eta)_+] \right).$$

Each inner supremum is concave and nondecreasing in ε by the argument already given for a fixed function of the atoms, and a minimum of such functions is concave and nondecreasing. The dual (27) is the finite-support case of the dual in [18]: for each multiplier $\theta_E \geq 0$ the objective is affine and nondecreasing in ε, which is the same concavity.

If every distinct pair has positive ground cost, the only coupling of cost zero is the identity, so $\rho_E(0)$ is the empirical mean. Every feasible P satisfies $\mathbb{E}_P[L] \leq \max_\ell L_\ell$. The edge case excluded by this hypothesis is a pair with ground cost zero and unequal losses: mass inside that pair can be rearranged at radius zero, and $\rho_E(0)$ is then the mean of the within-pair maxima, which can exceed the empirical mean. On the 40-atom training draw the smallest separation is 2.49 demand units, so that edge case does not occur.

If $L_\ell = c$ for every ℓ, every probability on Ξ has mean c. The conditional value-at-risk of a constant is that constant, for every β and every distribution, so the worst case over $\mathcal{P}_\varepsilon$ is c. In (27) the point $\theta_E = 0$, $s_i^E = c$ attains it.

If L is not constant, choose $\ell^\star \in \arg\max L$ and $i^\star$ with $L_{i^\star} < L_{\ell^\star}$. Their ground cost d is positive by hypothesis. Moving mass $\delta = \min(1/N, \varepsilon/d)$ from $i^\star$ to $\ell^\star$ remains inside $\mathcal{P}_\varepsilon$ and raises the expectation by $\delta(L_{\ell^\star} - L_{i^\star}) > 0$. □

A hybrid or meter plan is not automatically exposed. It is exposed only when the optimal token bill takes at least two values on the atoms. Scheduling does not enter the statement. Once the package is fixed, the activity times are computed before the quantities, which is the block separation already recorded for (25).

Corollary 2 (Critical radius of two fixed profiles): Fix $\lambda \in [0,1]$ and two recourse profiles whose losses, unmet demand, and authorization indicators do not depend on ε. Let profile 0 have constant expenditure c, and set $J_0 = c + U_0$, where U_0 collects its empirical unmet-demand penalty, empirical authorization term, and administrative term. Let profile 1 have worst-case score

$$J_1(\varepsilon) = (1-\lambda)\rho_E(\varepsilon) + \lambda\rho_C(\varepsilon) + U_1,$$

with U_1 the corresponding constant. Then J_1 is concave and nondecreasing, and

$$\varepsilon^\star = \inf\{\varepsilon \geq 0 : J_1(\varepsilon) \geq J_0\}$$

lies in $[0, \infty]$. Profile 1 scores strictly below profile 0 for every $\varepsilon < \varepsilon^\star$. Profile 0 scores no higher than profile 1 for every $\varepsilon \geq \varepsilon^\star$.

Proof. Corollary 1 gives that both ρ_E and ρ_C are concave and nondecreasing, so J_1 is as well. A nondecreasing function has superlevel sets that are empty or upper rays. The stated infimum is the left endpoint of that ray. If $J_1(0) \geq J_0$, then $\varepsilon^\star = 0$. If $J_1(\varepsilon) < J_0$ for every finite ε, then $\varepsilon^\star = \infty$ and profile 1 remains strictly better at every finite radius. A flat stretch on which $J_1 = J_0$ is a tie: both profiles are optimal throughout that stretch. Monotonicity is the property used for the threshold. Concavity is inherited from the expectation and is not required to define $\varepsilon^\star$. □

TABLE IV
CRITICAL RADIUS FOR THE PUBLISHED $\lambda = 0$ PROFILES ON THE 40-ATOM TRAINING DRAW. SCORES ARE THE TRAINING ROBUST OBJECTIVE. THE COMMITTED SCORE IS \$16,224.6 AT EVERY RADIUS. THE SOLVED PACKAGE IS A FRESH OPTIMUM OF THE SAME MIXED-INTEGER PROGRAM AT THAT RADIUS.

Factor	ε	Hybrid ρ_E	Hybrid score	Solved package
0.000	0.0	9,183.6	15,361.2	hybrid
0.250	561.7	9,942.0	16,119.5	hybrid
0.285	639.6	10,047.1	16,224.6	hybrid (tie)
0.300	674.1	10,093.6	16,271.2	committed
0.500	1,123.5	10,700.3	16,877.9	committed

The threshold is the decision boundary between two frozen profiles. The mixed-integer optimizer ranges over more profiles and may select a different exposed profile before $\varepsilon^\star$. On the published $\lambda = 0$ profiles the nominal hybrid purchase has $W = \$5{,}785.6$ and training objective \$15,361.2, and the committed purchase has constant expenditure \$8,223.1 and training objective \$16,224.6. Those are the objectives already reported for the joint and admin-first sequential rules. For these two frozen profiles the threshold defined above is 639.6 demand units, or 0.285 times the training mean absolute deviation of 2,247.0. Proposition 1 below derives that number in closed form from the loss structure, and an independent primal transportation linear program on the same atoms and the same published losses reproduces it to 10^{-9}. Table IV evaluates that frozen hybrid profile at the published radius factors. The package column is the optimum already computed for this draw; it is not a new sample.

At radius factor 0.25, below 639.6, the frozen hybrid score \$16,119.5 remains below the committed score \$16,224.6, and the solved package is hybrid. At 639.6 the two scores meet, so either profile is optimal; the solver returns the hybrid purchase. At factors 0.30 and 0.50 the frozen hybrid score is higher and the solved package is the committed purchase of \$8,223.1. The out-of-sample conditional-value-at-risk figures already reported for those packages are unchanged.

5) Closed Form of the Critical Radius: Corollary 2 defines $\varepsilon^\star$ but does not compute it. For the frozen hybrid profile the recourse is piecewise linear in demand, and the transport problem then has a closed-form value. Write q for the prepaid credit quantity, $\kappa = p^{\mathrm{tok}}$ for the token price, and Γ for the budget cap. In the published solve with $\lambda = 0$ the token stage buys each unit of demand above q at price κ until the cap binds, so

$$L_\ell = \min\{\Gamma,\ W + \kappa(\xi_\ell - q)_+\}.$$

This is checked exactly on all forty atoms: the maximum absolute deviation of the formula from the stored losses is 0.0, and the largest pairwise slope $|L_\ell - L_i|/|\xi_\ell - \xi_i|$ is 1.35.

Proposition 1 (Closed-form critical radius): Let $\lambda = 0$. Let L be κ-Lipschitz on the atoms, $|L_\ell - L_i| \le \kappa|\xi_\ell - \xi_i|$. Let S be a set of atoms on which $L_\ell = \alpha + \kappa\xi_\ell$ for one constant α, let $\bar{\ell}$ be the atom of S with the largest demand, and set

$$\varepsilon_S = \frac{1}{N}\sum_{i\in S}(\xi_{\bar{\ell}} - \xi_i).$$

Then, with $\bar{L} = N^{-1}\sum_\ell L_\ell$:

1) $\rho_E(\varepsilon) \le \bar{L} + \kappa\varepsilon$ for every $\varepsilon \ge 0$.
2) $\rho_E(\varepsilon) = \bar{L} + \kappa\varepsilon$ for $0 \le \varepsilon \le \varepsilon_S$.
3) Define $G = J_0 - J_1(0) = (c - \bar{L}) + (U_0 - U_1)$. If $G \ge 0$, then $\varepsilon^\star \ge G/\kappa$, and $\varepsilon^\star = G/\kappa$ whenever $G/\kappa \le \varepsilon_S$.

Proof. (1) Let $P \in \mathcal{P}_\varepsilon$ and let π be a coupling of $\hat{P}$ and P with $\sum_{i,\ell}\pi_{i\ell}|\xi_i - \xi_\ell| \le \varepsilon$. Because π has first marginal $\hat{P}$ and second marginal P,

$$\mathbb{E}_P[L] - \bar{L} = \sum_{i,\ell}\pi_{i\ell}(L_\ell - L_i) \le \kappa\sum_{i,\ell}\pi_{i\ell}|\xi_\ell - \xi_i| \le \kappa\varepsilon.$$

This is the Kantorovich–Rubinstein inequality on a finite support.

(2) Fix $\varepsilon \le \varepsilon_S$. Choose $\delta_i \in [0, 1/N]$ for $i \in S$ with $\sum_{i\in S}\delta_i(\xi_{\bar{\ell}} - \xi_i) = \varepsilon$. This is possible because the left side is continuous in δ and equals 0 at $\delta = 0$ and ε_S at $\delta_i \equiv 1/N$. Move mass δ_i from atom i to atom $\bar{\ell}$. The transport cost is ε, so the new law lies in $\mathcal{P}_\varepsilon$. Its expectation exceeds $\bar{L}$ by $\sum_{i\in S}\delta_i(L_{\bar{\ell}} - L_i) = \kappa\sum_{i\in S}\delta_i(\xi_{\bar{\ell}} - \xi_i) = \kappa\varepsilon$, since $\bar{\ell}, i \in S$. With (1) this gives equality.

(3) For $\lambda = 0$, $J_1(\varepsilon) = \rho_E(\varepsilon) + U_1$ and $J_0 = c + U_0$. By (1), $J_1(\varepsilon) \le J_1(0) + \kappa\varepsilon < J_0$ for every $\varepsilon < G/\kappa$, so $\varepsilon^\star \ge G/\kappa$. By (2), if $G/\kappa \le \varepsilon_S$ then $J_1(G/\kappa) = J_1(0) + G = J_0$, so $\varepsilon^\star \le G/\kappa$. □

a) Numbers.: The frozen hybrid profile has $W = q =$ \$5,785.58 and $\Gamma =$ \$16,000. Of the forty atoms, eleven lie at or below q and have $L = W$, twenty-seven lie on the linear branch with slope $\kappa = 1.35$, and two exceed the demand at which the cap binds and have $L = \Gamma$. Take S to be the twenty-seven branch atoms. Their largest demand is 13,244.7, and $\varepsilon_S = 2{,}896.1$. The ingredients of G are

$$c - \bar{L} = 8{,}223.07 - 9{,}183.61 = -960.54,$$
$$U_0 - U_1 = 8{,}001.57 - 6{,}177.55 = 1{,}824.01,$$

so $G = 863.48$, and $G/\kappa = 863.4755/1.35 = 639.6115$. Since $639.6 \le 2{,}896.1$, part (3) applies with equality: $\varepsilon^\star = 639.6115$, which is 0.28465 times the training mean absolute deviation 2,247.0. The transportation linear program returns 639.6115087, and the closed form returns 639.6115087. The same slope identity gives $\rho_E(\varepsilon) = 9{,}183.6 + 1.35\varepsilon$ at every factor in Table IV, for example $9{,}183.6 + 1.35 \times 561.7 = 9{,}942.0$. Factor 1.0 ($\varepsilon = 2{,}247.0$) is also below ε_S, and the same formula gives 12,217.1.

The unmet-demand and administrative terms enter only through G. The threshold is therefore the dollar advantage of the exposed profile at radius zero, divided by the price at which one unit of transported demand raises the token bill. A cheaper token, a larger prepaid-profile cost, or a larger unmet-demand penalty on the prepaid profile raises $\varepsilon^\star$. For $\lambda > 0$ the threshold is computed from the linear programs, because the worst-case conditional value-at-risk is not linear in ε on this profile (Table V shows it reaching the cap). The 639.6 value is specific to $\lambda = 0$ and to this profile.

6) When the Implemented Risk Charge Is Exact and When It Is Conservative: The implemented objective adds ρ_E and ρ_C, each a supremum over $\mathcal{P}_\varepsilon$ taken separately. A decision maker who faces one adversarial law faces the single supremum

$$J_{\text{joint}}(\varepsilon) = \sup_{P\in\mathcal{P}_\varepsilon}\Big\{(1-\lambda)\mathbb{E}_P[L] + \lambda\,\mathrm{CVaR}_\beta^P(L)\Big\},$$
$$J_{\text{sep}}(\varepsilon) = (1-\lambda)\rho_E(\varepsilon) + \lambda\rho_C(\varepsilon).$$

Let $\Delta(\varepsilon) = J_{\text{sep}}(\varepsilon) - J_{\text{joint}}(\varepsilon)$.

Proposition 2 (Exactness and conservatism of the separate charge): Fix a loss L on the atoms.

1) $\Delta(\varepsilon) \ge 0$. For any maximizer P_E of $\mathbb{E}_P[L]$ and any maximizer P_C of $\mathrm{CVaR}_\beta^P(L)$ over $\mathcal{P}_\varepsilon$,
$$\Delta(\varepsilon) \le \min\Big\{\lambda\big(\rho_C - \mathrm{CVaR}_\beta^{P_E}(L)\big), (1-\lambda)\big(\rho_E - \mathbb{E}_{P_C}[L]\big)\Big\}.$$
2) $\Delta(\varepsilon) = 0$ in each of these cases: $\lambda \in \{0, 1\}$; L constant on the atoms; $\varepsilon = 0$ under the separation hypothesis of Corollary 1; or one law $P \in \mathcal{P}_\varepsilon$ attains both suprema.
3) If the unknown demand law P^* lies in $\mathcal{P}_\varepsilon$, then $(1-\lambda)\mathbb{E}_{P^*}[L] + \lambda\mathrm{CVaR}_\beta^{P^*}(L) \le J_{\text{joint}}(\varepsilon) \le J_{\text{sep}}(\varepsilon)$, so the implemented charge never understates the true weighted risk, and its excess is at least $\Delta(\varepsilon)$. If $P^* \notin \mathcal{P}_\varepsilon$, neither inequality is implied.
4) $J_{\text{joint}}(\varepsilon)$ is the value of one linear program,
$$\min_{\eta,\theta\ge0,s,v\ge0}\ \lambda\eta + \theta\varepsilon + \frac{1}{N}\sum_i s_i$$
subject to $s_i \ge (1-\lambda)L_\ell + \frac{\lambda}{1-\beta}v_\ell - \theta|\xi_i - \xi_\ell|$ and $v_\ell \ge L_\ell - \eta$ for all i, ℓ.

Proof. (1) The supremum of a sum is at most the sum of the suprema, so $\Delta \ge 0$. The law P_E is feasible for J_{joint}, so $J_{\text{joint}} \ge (1-\lambda)\rho_E + \lambda\mathrm{CVaR}_\beta^{P_E}(L)$. Subtracting this from J_{sep} gives the first bound. The law P_C gives the second.

(2) At $\lambda = 0$ or $\lambda = 1$ only one term is present, so the two suprema coincide. If L is constant, every term is constant by Corollary 1. At $\varepsilon = 0$ with separated atoms the ball is the single law $\hat{P}$. If one law attains both suprema, it attains their weighted sum, and $J_{\text{joint}} \ge J_{\text{sep}}$.

(3) The first inequality holds because P^* is feasible. The second is part (1).

(4) By Rockafellar–Uryasev, $\mathrm{CVaR}_\beta^P(L) = \min_\eta\{\eta + (1-\beta)^{-1}\mathbb{E}_P[(L-\eta)_+]\}$. The map $(\eta, P) \mapsto \lambda\eta + \mathbb{E}_P\big[(1-\lambda)L + \frac{\lambda}{1-\beta}(L-\eta)_+\big]$ is convex in η and linear in P, and $\mathcal{P}_\varepsilon$ is a compact convex polytope. By Sion's minimax theorem the supremum and the minimum over η can be

TABLE V
SEPARATE-SUPREMA CHARGE AGAINST THE SINGLE-DISTRIBUTION WORST CASE, FROZEN HYBRID PROFILE, $\lambda = 0.50$, $\beta = 0.90$, SAME 40 TRAINING ATOMS. THE GAP IS NONNEGATIVE. DOLLARS.

Factor	ρ_E	ρ_C	Separate	Joint	Gap	Bound
0.00	9,183.6	15,748.0	12,465.8	12,465.8	0.0	0.0
0.25	9,942.0	16,000.0	12,971.0	12,963.8	7.2	36.2
0.50	10,700.3	16,000.0	13,350.2	13,342.9	7.2	36.2
1.00	12,217.1	16,000.0	14,108.5	14,101.3	7.2	36.2

exchanged. For fixed η the inner supremum is a finite transport linear program. Its dual, with multiplier θ on the radius constraint and s_i on the mass of atom i, is the one stated, after writing $(L_\ell - \eta)_+ = v_\ell$. □

Table V applies the proposition to the frozen hybrid loss vector, with $\lambda = 0.50$ and $\beta = 0.90$ as published, on the same forty atoms and the published radius factors. The joint value from the linear program in part (4) is confirmed by an independent computation: a one-dimensional convex minimization over η with a primal transport linear program inside. The two agree to 10^{-4} at every row.

At radius zero the charge is exact, as part (2) says. At every positive tested radius the worst-case conditional value-at-risk is already at the cap, $\rho_C = \$16{,}000 = \Gamma$, because the ball can place the 10% tail mass entirely on atoms whose loss equals the cap. The gap is then \$7.2, which is 0.06% of the joint value, and it lies below the computable bound of \$36.2. For this profile, the sum of separate suprema charges essentially the same as the single-law worst case. The modeling choice that governs interpretation is therefore the ambiguity set itself, and not the splitting of the two suprema. By part (3), the charge bounds the true weighted risk whenever the true law lies in the ball. The radius factors in Table V are sensitivity settings on a forty-atom sample, and the proposition is stated relative to that same ambiguity set. The out-of-sample figures already reported are unchanged.

This geometry rests on the dual of [18], which represents the worst-case expectation by an objective that is linear in the radius for each multiplier and recovers the nominal expectation at zero radius. A constant loss has the same expectation under every distribution, so the radius cannot move it. The procurement content is the identification of a prepaid credit plan as such a constant loss. That identification explains why the validated robust and nominal prepaid solutions coincide. Together with the closed form above, it locates the radius, 639.6, at which the published zero-risk-weight hybrid profile yields to that prepaid profile.

Remark 1 (Governance separation): Let X_F be the financially feasible set of Sections IX–IX-E, and let X_G be the subset that also satisfies the residual-risk inequality (29) and the exogenous TX-RAMP, security-screen, and architecture gates. The governed feasible set is the intersection $X_F \cap X_G$. A financially optimal plan that fails a gate is infeasible for the governed problem. If the intersection is empty, every financially optimal plan is rejected. The remark ensures that a conditional-value-at-risk comparison cannot waive a failed security gate. The numerical instances remain the solves already reported on the 80/500 seeds in Section XII.

TABLE VI
SYNTHETIC BASE INPUTS. NONE OF THESE VALUES IS A TEXAS RATE, A SALARY, OR AN OBSERVED PROCESSING TIME.

Input	Base value
Mean demand and CV	8,000 units; CV = 0.35
Prices	subscription \$1.60; credit \$1.00; token \$1.35
Officer-supplied cap Γ	\$16,000
Exogenous cap Γ_{stat}	\$10,000,000
Pre-authorization fraction φ	0.70 of the uncommitted remainder
Risk parameters	$\beta = 0.90$, $\lambda = 0.50$, $w_m = \$2.50$ per unmet unit
Administrative weights	\$30 per hour; \$15 per day
Hybrid configuration	45 days; 80 hours
Committed configuration	30 days; 55 hours
Meter configuration	40 days; 70 hours
Authorization adder	5 days; 6 hours

X. COMPUTATIONAL EVALUATION

The solves use `experiments/run_bpo_cvar.py`, Python 3.12.7, NumPy 2.2.6, and SciPy 1.17.1, with HiGHS called through `scipy.optimize.milp` and a relative gap tolerance of 10^{-8}. Every reported solve returned an optimal status. The training sample has 80 equiprobable scenarios drawn from standard normal shocks with seed 20261105. The test sample has 500 scenarios drawn with seed 20261106. The same shocks are transformed to a lognormal demand for each coefficient of variation, so the volatility cases are paired. Price, cap, λ, and β cases reuse the base demand sample. The base test sample has mean 7,897.0 units and standard deviation 2,763.2 units. The theoretical mean is 8,000. The difference is sampling variation, and the expected-demand rule still plans for 8,000 because that rule uses the theoretical mean.

Hours, days, prices, the unmet-demand weight, and the cap are synthetic. They were chosen so that a credit, which is paid whether it is used, and a token, which is paid only when it is used, are not ordered by strict dominance at every price and every risk weight. That is a design choice rather than a finding about vendor prices.

The four comparison rules use the same cap, the same prices, and the same test sample. The fixed-budget rule buys Γ / p^{cred} credits at award and uses the committed configuration. The expected-demand rule solves the program at a single scenario equal to the theoretical mean, with $\lambda = 0$, freezes that first stage, and then serves as much test demand as the frozen plan and the cap allow. The subscription-only rule buys $\min(\mu, \Gamma / p^{\text{sub}})$ seats and nothing else. The pay-as-you-go rule buys no committed units, pre-authorizes $R = \varphi\Gamma$, uses the meter configuration, and serves demand up to the cap, including authorized incremental funding when the pre-authorized budget is

TABLE VII
OUT-OF-SAMPLE COMPARISON ON THE SYNTHETIC TEST SAMPLE. CAP-FEASIBLE IS THE SHARE OF SCENARIOS THAT MEET THE FUNDING CAP, THE AUTHORIZATION RULE, AND THE SELECTED CONFIGURATION. IT IS NOT A LEGAL-COMPLIANCE RATE.

Policy	E[cost] ($)	$\text{CVaR}_{0.90}$ ($)	E[unmet] (units)	Service	E[days]	E[hours]	Auth.	Cap-feasible
Proposed plan	8,399.5	8,399.5	863.1	89.1%	30.0	55.0	0.0%	100.0%
Expected-demand	8,000.0	8,000.0	1,021.4	87.1%	30.0	55.0	0.0%	100.0%
Fixed budget	16,000.0	16,000.0	31.3	99.6%	30.0	55.0	0.0%	100.0%
Subscription only	12,800.0	12,800.0	1,021.4	87.1%	30.0	55.0	0.0%	100.0%
Pay as you go	10,407.4	15,945.0	187.8	97.6%	41.9	72.3	38.4%	100.0%

exhausted. The proposed rule solves the full program on the training sample and then re-solves only the recourse on the test sample, so the reported tail risk can include a decision to leave a unit unmet when the objective charges that unit's token cost inside the tail.

Cap feasibility in the tables is the share of test scenarios in which expenditure stays within the cap, incremental token funding is either zero or covered by an authorization on a meter-capable configuration, and the pre-authorized token budget respects φ. Every rule in the base comparison met that check in every test scenario. The check is not a legal audit.

A. Base Comparison

Table VII is the out-of-sample comparison at the inputs in Table VI. The proposed plan selects the committed configuration, buys 8,399.5 credit units, and pre-authorizes no tokens. Because the credit outlay does not depend on the scenario, expected expenditure and $\text{CVaR}_{0.90}$ are the same number, $8,399.5. Expected unmet demand is 863.1 units, which is 10.9 percent of test demand left unserved. Cycle time is 30.0 days and staff effort is 55.0 hours, with no authorization events. On this sample the risk-weighted token is not worth opening. The in-sample training solution has the same credit quantity, and the training CVaR auxiliary matches the sample CVaR of that constant outlay.

The expected-demand rule buys 8,000.0 credits, spends $8,000.0 with the same CVaR, and leaves 1,021.4 units unmet. It uses the same configuration and the same administrative time. The fixed-budget rule spends the entire $16,000 cap on credits, leaves 31.3 units unmet, and serves 99.6 percent of test demand. The subscription rule spends $12,800.0 to obtain the same 8,000.0 units of capacity as the expected-demand rule, because the assumed seat price is $1.60 rather than $1.00. Pay-as-you-go spends $10,407.4 in expectation, has $\text{CVaR}_{0.90}$ of $15,945.0, leaves 187.8 units unmet, and authorizes incremental funding in 38.4 percent of scenarios. Its expected cycle time is 41.9 days and its expected effort is 72.3 hours.

These rules do not line up in the same order on every metric.

Figure 4 plots that disagreement. The proposed plan has lower expected expenditure and a lower expenditure tail than pay-as-you-go and the fixed-budget rule, and it leaves more demand unmet than either. The fixed-budget rule buys service by paying for credits that many scenarios do not use. The subscription rule pays more than the credit rules for no gain in service against the expected-demand rule. Evaluated on the scalar in (12), using each rule's out-of-sample means, the scores are $12,657.3 for the proposed plan, $12,653.5 for expected demand, $16,443.6 for pay-as-you-go, $17,453.5 for subscription only, and $18,178.3 for the fixed budget. The expected-demand score is lower by $3.8. That difference is small beside the $16,000 cap. It is a property of this sample and these weights. It is not a finding that either rule performs better in an agency.

B. Sensitivity

Table VIII varies the coefficient of variation and holds the other base inputs fixed. The selected configuration remains the committed credit purchase. The credit outlay is $8,290.0 at CV 0.15, $8,399.5 at the base CV 0.35, $8,161.6 at CV 0.60, and $7,342.8 at CV 1.00. Expected unmet demand rises at each step, from 332.6 units to 2,637.2 units. Because the outlay is prepaid, every scenario has the same expenditure, and the average of the upper tail is that constant. CVaR therefore equals expected expenditure in every row of Fig. 5. Volatility changes the size of the block and the unmet demand. It does not, at these weights, make the expenditure random.

Table IX varies the token price. At $0.90, which is below the $1.00 credit price, the plan switches to the meter configuration. Expected expenditure is $6,680.1 and $\text{CVaR}_{0.90}$ is $8,728.4, so the tail now lies above the mean. Expected unmet demand is 474.6 units. The pre-authorized token budget is $9,973.8, below the maximum $\varphi\Gamma = \$11{,}200$, and the authorization rate is zero. Relative to a rule that would serve every additional unit the cap can fund, the optimized recourse leaves 459.5 more units unmet on average. That gap is the CVaR term declining token expenditure in the upper tail. At token prices of $1.10, $1.35, $1.80, and $2.50, the plan returns to the same committed credit block as the base case. On this grid, a token must be cheaper than a credit before the base risk weight will use it.

Table X varies the officer-supplied cap. At $8,000 the cap binds: the plan spends $8,000 and leaves 1,021.4 units unmet, the same pattern as buying the mean with credits.

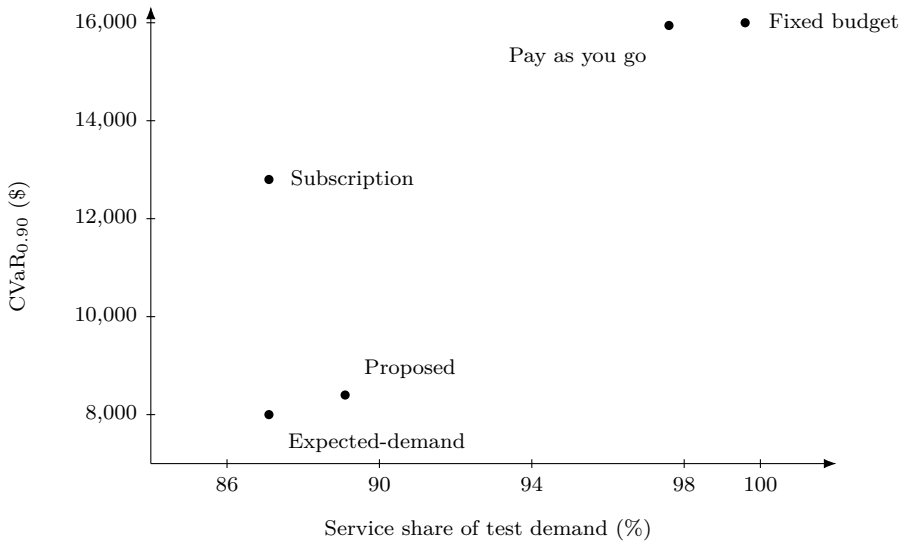


Fig. 4. Out-of-sample service share and expenditure tail for the five synthetic rules at the base inputs. Marker height is $\mathrm{CVaR}_{0.90}$. Expected expenditure equals that tail for every prepaid rule and is lower than the tail for pay-as-you-go (10,407.4 against 15,945.0).

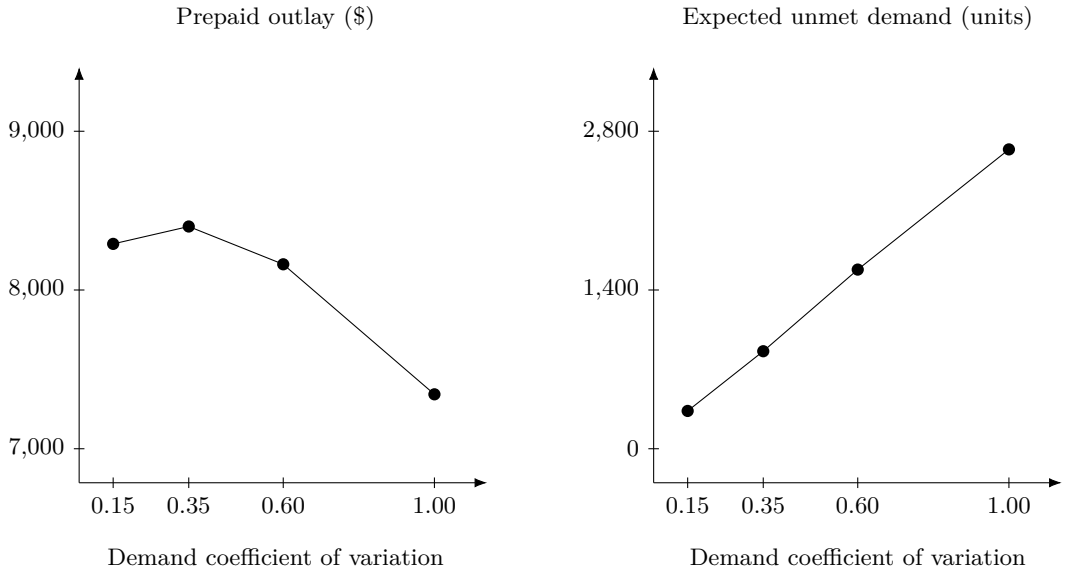


Fig. 5. Sensitivity of the proposed plan to the demand coefficient of variation. The selected configuration remains the committed credit purchase, so conditional value-at-risk equals the prepaid outlay in the left panel. Unmet demand in the right panel rises with volatility. Both panels use the paired shock sample.

TABLE VIII
Sensitivity of the proposed plan to demand volatility. Other base inputs are unchanged, and $\beta = 0.90$.

Demand CV	E[cost] ($)	CVaR_{β} ($)	E[unmet] (units)	Service	E[days]	E[hours]	Process
0.15	8,290.0	8,290.0	332.6	95.8%	30.0	55.0	committed
0.35	8,399.5	8,399.5	863.1	89.1%	30.0	55.0	committed
0.60	8,161.6	8,161.6	1,580.7	79.8%	30.0	55.0	committed
1.00	7,342.8	7,342.8	2,637.2	66.0%	30.0	55.0	committed

At \$16,000, \$22,000, and \$30,000 the credit block stays at \$8,399.5. Extra cap room is not used. The exogenous \$10 million cap is slack in every run. The largest officer-supplied cap in the grid is \$30,000.

Table XI varies λ at $\beta = 0.90$. At $\lambda = 0$ the plan switches to the hybrid configuration. It commits \$5,830.2 to credits, pre-authorizes \$7,118.9 of token budget, and uses authorization in 10.6 percent of test scenarios. Expected expenditure is \$8,803.0, $\mathrm{CVaR}_{0.90}$ is \$15,149.4, expected unmet demand falls to 96.4 units, expected days

TABLE IX
SENSITIVITY OF THE PROPOSED PLAN TO THE SYNTHETIC TOKEN PRICE.

Token price ($/unit)	E[cost] ($)	CVaR$_\beta$ ($)	E[unmet] (units)	Service	E[days]	E[hours]	Process
0.90	6,680.1	8,728.4	474.6	94.0%	40.0	70.0	meter
1.10	8,399.5	8,399.5	863.1	89.1%	30.0	55.0	committed
1.35	8,399.5	8,399.5	863.1	89.1%	30.0	55.0	committed
1.80	8,399.5	8,399.5	863.1	89.1%	30.0	55.0	committed
2.50	8,399.5	8,399.5	863.1	89.1%	30.0	55.0	committed

TABLE X
SENSITIVITY OF THE PROPOSED PLAN TO THE OFFICER-SUPPLIED PLANNING CAP.

Planning cap ($)	E[cost] ($)	CVaR$_\beta$ ($)	E[unmet] (units)	Service	E[days]	E[hours]	Process
8,000	8,000.0	8,000.0	1,021.4	87.1%	30.0	55.0	committed
16,000	8,399.5	8,399.5	863.1	89.1%	30.0	55.0	committed
22,000	8,399.5	8,399.5	863.1	89.1%	30.0	55.0	committed
30,000	8,399.5	8,399.5	863.1	89.1%	30.0	55.0	committed

TABLE XI
SENSITIVITY OF THE PROPOSED PLAN TO THE RISK WEIGHT λ, AT $\beta = 0.90$.

λ	E[cost] ($)	CVaR$_\beta$ ($)	E[unmet] (units)	Service	E[days]	E[hours]	Process
0	8,803.0	15,149.4	96.4	98.8%	45.5	80.6	hybrid
0.25	8,399.5	8,399.5	863.1	89.1%	30.0	55.0	committed
0.50	8,399.5	8,399.5	863.1	89.1%	30.0	55.0	committed
0.75	8,399.5	8,399.5	863.1	89.1%	30.0	55.0	committed
1	8,399.5	8,399.5	863.1	89.1%	30.0	55.0	committed

TABLE XII
SENSITIVITY OF THE PROPOSED PLAN TO THE CVAR LEVEL β, AT $\lambda = 0.50$.

β	E[cost] ($)	CVaR$_\beta$ ($)	E[unmet] (units)	Service	E[days]	E[hours]	Process
0.80	8,399.5	8,399.5	863.1	89.1%	30.0	55.0	committed
0.90	8,399.5	8,399.5	863.1	89.1%	30.0	55.0	committed
0.95	8,399.5	8,399.5	863.1	89.1%	30.0	55.0	committed

are 45.5, and expected hours are 80.6. The average amount withheld relative to full service inside the cap is 0.8 units, so this risk-neutral plan is essentially a full-service plan. The reported CVaR at $\lambda = 0$ is computed from the test expenditure sample. The solver's CVaR auxiliary is not pinned down when its weight is zero, and that auxiliary is not the number in the table. At $\lambda = 0.25$, 0.50, 0.75, and 1, the plan is the same committed block as the base case. Under these weights, a modest weight on expenditure tail risk is enough to give up the token stage and accept the higher unmet demand of a certain credit purchase.

Table XII varies β at $\lambda = 0.50$. The plan does not change. The expenditure of the selected plan is nonrandom, so the tail level has nothing left to reshape.

C. Cost, Tail, and Service

Equation (12) is already a weighted sum of expected procurement expenditure, expenditure tail risk, and unmet demand. At the published unmet-demand weight $w_m = \$2.50$, varying λ produces only the two rows in Table XI. The service axis was therefore traced by solving that same program, on the same 80 training scenarios and the same 500 test scenarios, at $\lambda \in \{0, 0.25, 0.50, 1\}$ and $w_m \in \{0.25, 0.50, 1.00, 1.35, 2.50, 5, 10\}$. Each frozen first stage was scored by the same test recourse used for the base comparison. No objective was added.

A solved point is kept when no other point in this set has expected cost no higher, CVaR$_{0.90}$ no higher, and service no lower, with at least one score strictly better. The comparison set is the weight grid together with the four reference rules. A weighted sum need not recover every nondominated solution of a mixed-integer program, so Table XIII is not a complete Pareto set. Fig. 6 plots the same solved points against service.

On this sample the grid has ten distinct outcomes, and all ten are nondominated against the reference rules. Whenever $w_m \leq \$1.00$, all four risk weights buy no credits and no tokens. Procurement expenditure and CVaR$_{0.90}$ are zero, and service is zero. The solver still selects the committed configuration. At $w_m = \$1.35$, all four risk

TABLE XIII
NONDOMINATED OUTCOMES ON THE BASE TEST SAMPLE. SERVICE IS GIVEN TO TWO DECIMALS SO THE HYBRID ROWS REMAIN DISTINCT. THE BASE CREDIT ROW IS THE 89.1 PERCENT SOLUTION ALREADY REPORTED AT $w_m = \$2.50$. THE SUBSCRIPTION AND PAY-AS-YOU-GO RULES ARE OMITTED HERE BECAUSE BOTH ARE DOMINATED ON THIS SAMPLE.

λ	w_m	Rule	Package	E[cost] ($)	$\text{CVaR}_{0.90}$ ($)	Service (%)
0–1	≤ 1.00	Weighted sum	Committed	0.0	0.0	0.00
0–1	1.35	Weighted sum	Committed	5,774.6	5,774.6	70.35
—	—	Expected-demand	Committed	8,000.0	8,000.0	87.07
0.25–1	2.50	Weighted sum	Committed	8,399.5	8,399.5	89.07
0.50, 1	5.00	Weighted sum	Committed	11,233.7	11,233.7	96.88
0.50, 1	10	Weighted sum	Committed	13,239.1	13,239.1	98.72
0	2.50	Weighted sum	Hybrid	8,803.0	15,149.4	98.78
0	5.00	Weighted sum	Hybrid	8,821.9	15,042.3	98.86
0.25	5.00	Weighted sum	Hybrid	8,853.3	14,939.7	98.87
0	10	Weighted sum	Hybrid	8,857.2	14,978.3	98.90
0.25	10	Weighted sum	Hybrid	8,975.4	14,864.6	98.97
—	—	Fixed budget	Committed	16,000.0	16,000.0	99.60

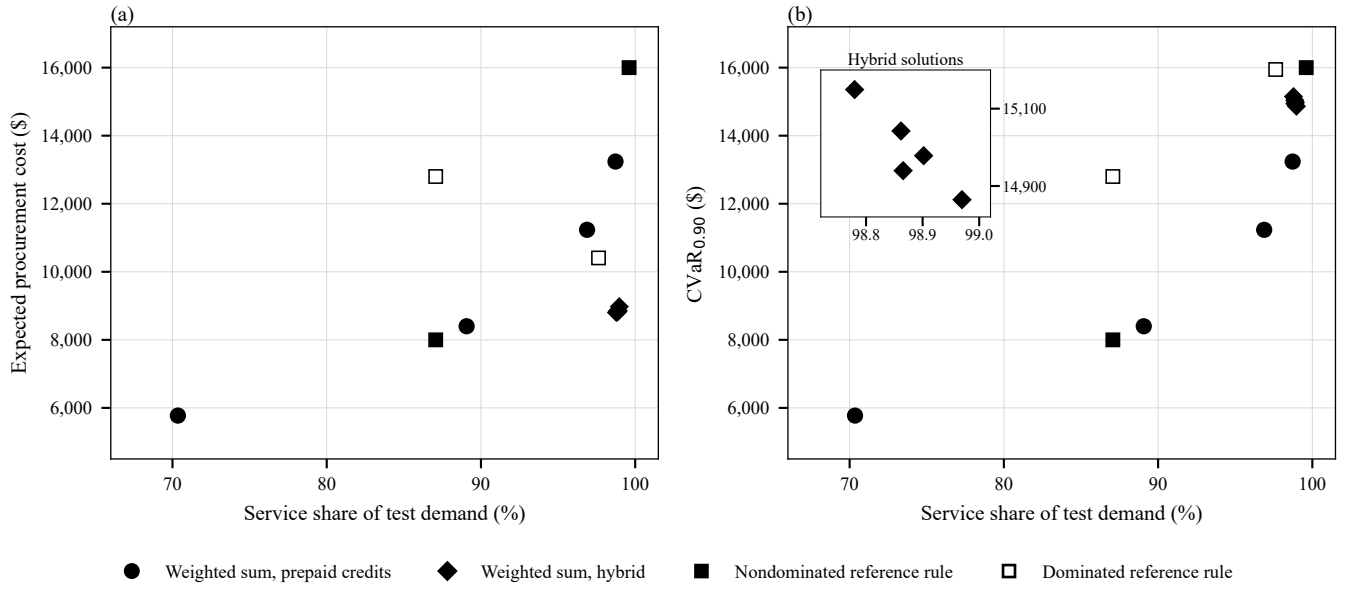


Fig. 6. Out-of-sample expected procurement cost (a) and $\text{CVaR}_{0.90}$ (b) against service on the base 500-scenario test sample. Filled circles and diamonds are solutions of (12) at $\lambda \in \{0, 0.25, 0.50, 1\}$ and $w_m \in \{0.25, 0.50, 1.00, 1.35, 2.50, 5, 10\}$. Filled squares are the expected-demand and fixed-budget rules. Open squares are the subscription and pay-as-you-go rules. The zero-purchase solution, returned whenever $w_m \leq 1.00$, is at zero cost and zero service and is outside the axes. The inset separates the five hybrid solutions. No curve is drawn. These markers are the nondominated subset of the solved grid and the four reference rules, not a complete Pareto set of the mixed-integer program.

weights buy 5,774.6 credits, spend \$5,774.6 with the same tail, and serve 70.35 percent of test demand. At $w_m =$ \$2.50, the solves repeat Table XI: $\lambda = 0$ is the hybrid row, and $\lambda \geq 0.25$ is the base credit purchase of 8,399.5. At $w_m =$ \$5.00 and \$10, with $\lambda \in \{0.50, 1\}$, the plan stays prepaid and buys 11,233.7 or 13,239.1 credits. Service is 96.88 or 98.72 percent, and the tail equals the outlay. The other four cells at those two weights, together with the published $\lambda = 0$, $w_m =$ \$2.50 cell, are hybrid. Their expected costs lie between \$8,803.0 and \$8,975.4. Their tails lie between \$14,864.6 and \$15,149.4. Their service lies between 98.78 and 98.97 percent. For a prepaid row the tail equals expected cost. For a hybrid row the tail is about \$6,000 above the expected cost. That gap is the quantity this grid contributes. It is not a ranking of the two families.

The expected-demand rule, at \$8,000 and 87.07 percent service, is nondominated. It sits between the \$5,774.6 block and the base block, neither of which the weight grid hits at exactly 8,000 credits. The fixed-budget rule, at \$16,000 and 99.60 percent service, is the only computed policy above 99 percent service, and it is nondominated. The subscription rule serves the same 87.07 percent as the expected-demand rule and costs \$12,800.0 rather than \$8,000.0, so it is dominated. The pay-as-you-go rule costs \$10,407.4, has $\text{CVaR}_{0.90}$ of \$15,945.0, and serves 97.62 percent. Each of the five hybrid solutions has a lower expected cost, a lower tail, and higher service than that rule on this test sample. The base credit plan does not dominate it: the base tail is lower, and the base service is lower. Dominance here is a comparison of these computed

triples. A different price or a different draw can change it.

An officer uses the table by stating three limits that the calculation does not choose. The funding cap in this grid is the published $16,000, and every solved policy is feasible for that cap. A service floor removes the zero-purchase row and any block below the floor. A ceiling on the expenditure tail removes every hybrid once the ceiling is below about $14,865, and it removes a prepaid block once the block itself exceeds the ceiling. If more than one row remains, the rows that remain trade expected cost, tail risk, and service. Prepaid rows are the computed policies for an officer who requires the tail to equal expected expenditure. Hybrid rows are the computed policies for an officer who will accept a tail near $15,000 in return for service near 99 percent. The fixed-budget row is the computed policy for an officer who wants the highest service in this set and will spend the whole cap. The weights λ and w_m are one way to name that choice; no office is assumed to hold them.

D. Reading the Calculation

The base result shows that the integrated program can return a coherent one-instrument recommendation that is simultaneously feasible for the funding cap, the authorization limit, and the selected administrative configuration. Two-stage recourse and CVaR remain active in the formulation in every run. They change the recommendation when the token is cheap relative to the credit and when the objective puts no weight on the expenditure tail—the settings in which risk sensitivity and meter recourse are material. Across the base grid of β, across token prices at or above $1.10, and across caps at or above $16,000, the same prepaid credit optimum remains selected. That identity is the expected mathematical consequence of a nonrandom prepaid outlay: once every scenario spends the same amount, CVaR_β collapses to that constant, the tail level β has no further decision content, and additional cap room above the block is slack. Administrative time and staff hours change when the configuration changes, which is why the business-process term is decisive in the $\lambda = 0$ row and in the $0.90 token row, and why it is compatible with an unchanged financial purchase when the committed configuration is already optimal.

This section reports planning comparisons rather than predicted savings for a Texas institution. The ordering of the five rules would move if the unmet-demand weight, the staff weight, or the prices moved. The $3.8 planning-score gap between the proposed plan and the expected-demand rule is smaller than the precision with which an agency could know its own demand. A different test sample could reverse it. The demonstrated improvement is the joint recommendation itself—a risk-aware, authorization-feasible funded amount and package—together with the cases in which risk weight, token price, scheduling, or Wasserstein radius changes that recommendation. Theoretical capability of the layers is broader than those demonstrated switches; the experiments report both.

E. Extended Numerical Experiment

The extended solves use `experiments/run_joint_dro.py` and the Python, NumPy, and SciPy versions stated above, with HiGHS called through `scipy.optimize.milp` and a relative gap tolerance of 10^{-6}. Every solve recorded in `joint_dro_results.json` returned an optimal status and a gap of zero. The largest first-stage solve at $N = 40$ took 0.29 seconds on one simulated file.

The extended sample is not the sample in Section X. Training uses 40 equiprobable demands from seed 20261105. Testing uses 200 demands from seed 20261106. Duration noise uses seed 20261107, with 64 draws in the base schedule, each base duration multiplied by $\exp(0.20\,Z)$ and rounded up to at least one day. The training mean is 8,099.6 units and the test mean is 7,818.6 units. The training mean absolute deviation from the median is 2,247.0 units, so the base radius is 1,123.5 units. Because the samples differ, the extended credit quantity is not comparable with the 8,399.5 credit quantity above.

Table XIV is the base schedule at staff capacity two. TO-BE shortens the committed file by 5.02 days and leaves the 165.0 hours unchanged. On 100 fresh duration draws, the list schedule and the critical path differed by 0 days in every draw. A discrete-event queue of several files was not run. At capacity one, a separate 24-draw check made TO-BE serial: the committed file took 27.38 mean days, against 22.33 days at capacity two. The financial purchase did not change.

Table XV compares the five rules and preserves the verified out-of-sample numbers. Model A buys the mean, 8,000 credits, and leaves 973.0 units unmet. Model B buys 8,223.1 credits. Expenditure does not depend on the test draw, so CVaR equals the mean by the Rockafellar–Uryasev identity for a constant loss. Model C, at the base radius, returns the same purchase and the same out-of-sample numbers. The in-sample worst-case CVaR is also 8,223.1, for every radius in $\{0, 0.25, 0.50, 1\} \times \mathrm{MAD}$. That agreement is the prepaid-expenditure case of Section IX-E: every atom has the same expenditure, so the Wasserstein dual has no worse expenditure mass to transport and the ambiguity constraints are inactive on the financial block. An $\varepsilon = 0$ dual solve matched model B to a numerical gap of zero on cost, CVaR, and package, confirming consistency of the robust layer with the nominal model when the radius collapses. Model D forces the serial mode and adds 5.02 days at the same cost, the same unmet demand, and the same 165.0 hours, which isolates the demonstrated schedule value of TO-BE without altering the financial optimum. Model E's administrative first step selects committed TO-BE with no tie, and the joint rule selects the same package. Under this base weight the admin-first and integrated rules therefore agree on a stable prepaid optimum; agreement is solution stability of the integrated architecture, not evidence that joint optimization is unnecessary. Section X-F2 reports the distinct finance-first comparison. The demonstrated value of the joint model

TABLE XIV
MEAN SCHEDULE OF THE SYNTHETIC FILE. LABOR IS SIX PERSON-HOURS TIMES REALIZED DURATION AND DOES NOT FALL WHEN ACTIVITIES OVERLAP. THE 0.90 MAKESPAN QUANTILE IS AN ADMINISTRATIVE TAIL, NOT A FINANCIAL CVAR.

Package	Mode	Mean days	Day quantile 0.90	Mean hours
Committed	AS-IS	27.50	30.00	165.0
Committed	TO-BE	22.48	25.70	165.0
Meter	AS-IS	27.69	30.70	166.1
Meter	TO-BE	22.67	25.70	166.1
Hybrid	AS-IS	31.77	34.70	190.6
Hybrid	TO-BE	23.06	26.00	190.6

TABLE XV
OUT-OF-SAMPLE COMPARISON OF THE EXTENDED MODELS ON THE 200-SCENARIO TEST DRAW. FINANCIAL CVAR IS NOMINAL ON THAT DRAW. IDENTICAL PREPAID ROWS INDICATE INACTIVE EXPENDITURE AMBIGUITY AND SOLUTION STABILITY UNDER THE BASE RISK WEIGHT; CYCLE TIME IS THE MEAN SCHEDULE OF THE SELECTED PACKAGE, PLUS AUTHORIZATION DELAY WHEN IT OCCURS.

Model	Package	E[cost]	$\text{CVaR}_{0.90}$	E[unmet]	Service	Mean days
A. Deterministic joint	Committed, TO-BE	8,000.0	8,000.0	973.0	87.6%	22.48
B. Two-stage CVaR, joint	Committed, TO-BE	8,223.1	8,223.1	886.9	88.7%	22.48
C. Finite-support Wasserstein	Committed, TO-BE	8,223.1	8,223.1	886.9	88.7%	22.48
D. AS-IS only, model B	Committed, AS-IS	8,223.1	8,223.1	886.9	88.7%	27.50
E. Sequential, then model B	Committed, TO-BE	8,223.1	8,223.1	886.9	88.7%	22.48

on this table is decision feasibility itself: one package, one funded amount, and one schedule that satisfy the financial, authorization, and timing terms together.

The identity of Models B and C breaks when expenditure is random. At $\lambda = 0$, model B selects hybrid TO-BE: prepaid credits $W = \$5{,}785.6$, token budget $R = \$7{,}150.1$, out-of-sample expected expenditure \$8,733.5, $\text{CVaR}_{0.90}$ \$15,239.8, unmet demand 56.8, service 99.3%, and authorization in 11% of test scenarios. The same weight with the base Wasserstein radius selects the committed credit purchase instead. Its out-of-sample cost and CVaR are \$8,223.1, and unmet demand rises to 886.9. The robust plan has the lower expenditure tail and the lower service. Neither plan dominates the other on every metric; the switch is the demonstrated setting in which demand ambiguity changes the optimal package. The auxiliary CVaR variables are not read at $\lambda = 0$, because their objective weight is zero.

At a token price of \$0.90 and $\lambda = 0.50$, both the nominal model and the robust model select meter TO-BE with $R = \$9{,}942.6$ and no prepaid credits. Out-of-sample expected expenditure is \$6,665.3, nominal $\text{CVaR}_{0.90}$ is \$8,866.2, unmet demand is 412.8, and service is 94.7%. The first-stage decision does not move. The in-sample robust expectation is \$8,033.2, above the in-sample nominal expectation of \$7,022.0, and the in-sample worst-case CVaR equals the pre-authorized budget, \$9,942.6. The ambiguity set therefore changes the objective value while leaving the selected plan stable—another form of robustness of the meter recommendation under the tested radius.

Across demand coefficients of variation 0.15, 0.35, and 0.60, models B and C stay on the committed TO-BE credit. Out-of-sample costs are \$8,213.2, \$8,223.1, and \$7,883.8. Unmet demand is 348.1, 886.9, and 1,584.2. Service is 95.6%, 88.7%, and 79.4%. Caps of \$8,000, \$16,000, and \$30,000 repeat the first experiment's pattern on the new sample: the \$8,000 cap binds at a cost of \$8,000 and unmet demand of 973.0; the larger caps stay at \$8,223.1, so extra authorization room is slack once the prepaid block is chosen. Duration scales of 0.5, 1, and 2, each with its own 16-draw schedule, change days and hours and do not change the credit quantity, which again separates schedule effects from financial selection under prepaid outlays. The delay-weight grid $\{0, 15, 60, 240\}$ is the computed Pareto sweep. At a zero delay weight the serial and parallel committed plans have the same objective, and the solver returned AS-IS. That return is a tie, not a preference. At every positive delay weight the parallel plan is the unique optimum—the demonstrated setting in which scheduling changes the selected mode. No other package enters the sweep, so the nondominated financial outcome on these weights is a single credit plan with a schedule choice that responds to the delay weight.

The finite-support robust solve at $N = 10$, 20, and 40 took 0.05, 0.07, and 0.17 seconds, each with a zero gap and the committed TO-BE package. The integer dimension is the six package binaries plus one authorization binary per atom. That structure explains the short times; larger activity graphs and balls off the sample would require separate timing studies.

F. Further Checks on Integration, Robustness, and Waiting

The checks use `experiments/run_strengthen.py` and HiGHS. The committed outlay is \$8,223.1, matching the

stored extended solution, with a zero gap. Every financial solve finished with a zero gap.

1) Joint Selection and the Admin-First Sequence: Three selection rules are distinguished throughout. The *admin-first* sequence ranks packages by the administrative charge only, then solves the financial block inside that package; it is the sequential rule already reported in this section. The *finance-first* sequence ranks packages by the financial block—expected expenditure, CVaR, unmet-demand penalty, and authorization effort—then applies the TO-BE mode. The *integrated* rule minimizes the sum of those two blocks, combining established components within one objective.

The admin-first rule still selects committed TO-BE, because that package has the lowest administrative cost. The integrated rule selects the same package at $\lambda = 0.25$ and at $\lambda = 0.50$. Both then have the training objective \$16,224.6. That agreement is between admin-first and integrated selection: under positive risk weight on this sample, the administratively cheapest package is already the joint financial optimum, and the integrated model certifies that coincidence rather than overturning it. A finance-first sequence is a different baseline. On the same published training draw at $\lambda = 0.50$ it selects hybrid TO-BE, with training objective \$16,988.6, which is \$764.0 above the integrated objective. The financial blocks of hybrid TO-BE and committed TO-BE differ by \$12.5 (\$10,924.9 versus \$10,937.4); the administrative charges differ by \$776.5. Integration therefore overturns finance-first on this draw because a near-tie in the financial block is smaller than the process-cost gap, and it does not overturn admin-first. At $\lambda = 0$ the integrated rule and the finance-first rule both select hybrid TO-BE and the training objective falls to \$15,361.2. The admin-first rule, forced to stay on the committed package, stays at \$16,224.6. The integrated rule is then better on the objective it optimizes and surfaces the risk–service trade-off that an admin-first sequence hides: hybrid $\mathrm{CVaR}_{0.90}$ is \$15,239.8 and service is 99.3%, while the committed plan has $\mathrm{CVaR}_{0.90}$ \$8,223.1 and service 88.7%. Integration therefore changes the recommendation relative to admin-first when the financial gain from hybrid is allowed to outweigh its extra hours, and it changes the recommendation relative to finance-first when the process-cost gap outweighs a financial near-tie.

2) When Does Integration Matter?: The preceding paragraph uses the published training draw. A separate thirty-seed experiment, reported here for the first time and complementary to the validated tables, evaluates the three rules on identical scenarios, durations, and parameters. The recorded protocol is `run_integration_value.py` with seeds 20262001–20262030, $N_{\mathrm{train}} = 40$, $N_{\mathrm{test}} = 200$, and $N_{\mathrm{dur}} = 32$, calling the published solver `run_joint_dro.py`; the stored output is `integration_value_results.json`. Table XVI reports means of the out-of-sample objective J. At $\lambda = 0.50$ and with no Wasserstein radius, finance-first selects hybrid TO-BE on every seed and the integrated rule selects committed TO-BE on every seed. The paired difference in J is \$685.2, with 95% interval [\$663.7, \$706.6]. With the published finite-support radius of 0.50 times the training mean absolute deviation, the two rules select the same package on 28 of 30 seeds. The paired difference is \$49.4, with 95% interval [\$-20.8, \$119.6] and paired t $p = 0.16$. That difference is not statistically distinguishable from zero, so under the robust design the two rules perform comparably. At $\lambda = 0$ they agree on every seed. The admin-first rule differs from the integrated rule on 13 of 30 seeds in both the nominal and the robust designs, because a near-tie in administrative cost between committed TO-BE and meter TO-BE is resolved by duration noise; that split is not counted here as an integration benefit.

The relationship between the financial-objective gap and the administrative process-cost gap is an empirical finding that holds under the stated additive model, simulated sample, and parameter settings. On these draws, the finance-first recommendation changes when the financial gap between candidate packages is smaller than their process-cost gap, and it does not when that financial gap is large or when a prepaid plan already removes the near-tie. Hard residual-risk caps, applied as filters on the same frontier, shrink or empty the feasible set; they do not create a non-additive coupling between finance and process. The experiment is a single-file sensitivity analysis with an additive staff-and-delay charge; multi-file and shared-resource settings are outside its scope, and the validated \$16,224.6 and \$15,361.2 training objectives are unchanged.

3) When the Wasserstein Decision Changes Out-of-Sample Risk: Table XVII keeps the nominal and robust decisions on the same test draw. At $\lambda = 0.50$, on lognormal demands with coefficients of variation 0.35 and 1.00, on a gamma sample with the same mean and coefficient 0.35, and on a normal sample truncated at zero with the same moments, the nominal plan is already a prepaid credit block. Raising the radius from 0.25 to 1 times the training mean absolute deviation does not change the package or the out-of-sample CVaR. Because expenditure is nonrandom, the Wasserstein dual multipliers on expenditure remain inactive for every tested radius; identity of the robust and nominal decisions is therefore expected mathematical stability, not a missing robust gain.

At $\lambda = 0$ the nominal plan is hybrid and the tail is random—the setting in which the ambiguity layer can change the decision. On the base lognormal sample, a radius of 0.25 times the mean absolute deviation leaves that hybrid plan in place, so out-of-sample CVaR stays \$15,239.8. At 0.50 and at 1 times the mean absolute deviation the robust plan switches to committed TO-BE, out-of-sample CVaR falls to \$8,223.1, and service falls from 99.3% to 88.7%. The gamma sample and the truncated normal switch at the smaller radius 0.25 as well: CVaR falls from \$14,710.8 to \$8,491.2, and from \$14,550.4 to \$8,621.3, while service falls from 99.7% and 99.7% to 89.5% and 89.7%. On the lognormal sample with coefficient 1.00, the switch occurs at radius factor 0.50: CVaR falls from \$16,000.0, the cap, to \$6,970.8, and

TABLE XVI
FINANCE-FIRST, ADMIN-FIRST, AND INTEGRATED SELECTION. THE PUBLISHED-DRAW COLUMN IS THE VALIDATED TRAINING OBJECTIVE ON THE STORED EXTENDED SAMPLE AT $\lambda = 0.50$. THE THIRTY-SEED COLUMNS ARE NEW OUT-OF-SAMPLE MEANS OF J THAT COMPLEMENT THE VALIDATED TABLES. ROBUST USES THE PUBLISHED RADIUS OF 0.50 TIMES THE TRAINING MEAN ABSOLUTE DEVIATION.

Rule	Published draw, package / J_{train}	30-seed J, nominal	30-seed J, robust
Finance-first	hybrid TO-BE / 16,988.6	16,782.5	16,143.8
Admin-first	committed TO-BE / 16,224.6	16,949.6	16,981.0
Integrated	committed TO-BE / 16,224.6	16,097.3	16,094.4

TABLE XVII
FINITE-SUPPORT WASSERSTEIN DECISION AGAINST THE NOMINAL DECISION ON THE SAME TEST DRAW. A CHANGE IN OUT-OF-SAMPLE CVAR OCCURS ONLY WHEN THE PACKAGE CHANGES. SERVICE IS THE SHARE OF TEST DEMAND SERVED.

Distribution	λ	Radius factor	Nominal CVaR	Robust CVaR	Service, nominal / robust
Lognormal, CV 0.35	0	0.25	15,239.8	15,239.8	99.3% / 99.3%
Lognormal, CV 0.35	0	0.50	15,239.8	8,223.1	99.3% / 88.7%
Lognormal, CV 1.00	0	0.25	16,000.0	16,000.0	85.2% / 85.2%
Lognormal, CV 1.00	0	0.50	16,000.0	6,970.8	85.2% / 65.9%
Gamma, CV 0.35	0	0.25	14,710.8	8,491.2	99.7% / 89.5%
Truncated normal, CV 0.35	0	0.25	14,550.4	8,621.3	99.7% / 89.7%
All four distributions	0.50	0.25 to 1	unchanged credit plan	same	same

service falls from 85.2% to 65.9%. Across those demonstrated cases the robust rule lowers the expenditure tail by refusing the hybrid plan and accepts lower service in return. That trade-off is the theoretical capability of the Wasserstein layer made visible; it is not claimed as a uniform dominance on service.

Three files were also given separate \$16,000 caps and then equal shares of one \$16,000 cap. The separate committed outlays are \$8,223.1, \$8,294.0, and \$8,759.9, which sum to \$25,277.0 and cannot all be prepaid inside one \$16,000 cap. The equal split buys \$5,333.3 of credits for each file. Out-of-sample unmet demand on the three test draws is 3,110.4, 3,178.7, and 2,988.6 units. This split is not a coupled mixed-integer optimum. It shows only that three copies of the single-file plan are not jointly feasible under one cap of the same size.

4) Staff, Parallel Work, and Several Files: The duration draws in Table XVIII are paired across staff levels and are not paired across the number of files. The network is the committed TO-BE graph. The large network adds six modeled clerical activities named for the audit trail already in the procedure: a value note, an evidence file, a risk file, a quota log, an invoice rule, and a closeout. Those names are not new statutory steps. The schedule is a longest-path list rule with a shared staff limit. For one small file and two staff, mean waiting is zero and the mean duration is 22.68 days, in line with the 22.48-day critical-path mean of the earlier 64-draw sample. Four staff leaves that file at 22.68 days. One staff stretches it to 27.62 days and the mean activity start waits 1.77 days, because the single person cannot run the parallel arcs together. Deleting the approval activity, which also deletes its labor, shortens the one-file, two-staff case from 22.68 days to 20.77 days. That is a graph deletion, not a legal option, and it is not evidence that the approval gate can be skipped.

Three small files on two staff finish in a portfolio of 41.40 days, against 23.20 days when four staff are shared. Mean start delay falls from 4.00 days to 0.46 days. Five large files remain congested at four staff: mean file duration 52.02 days and mean start delay 4.34 days, against 102.90 days and 11.14 days at two staff. A simulated availability shock, which cuts one staff member on a quarter of the days, moves the three large files at two staff from 62.04 days to 65.08 days. The single-file critical path does not contain these waits.

G. Simulation-Based Validation and Robustness

Thirty independent training seeds, 20262001 through 20262030, provide a second simulated sample, distinct from the 80-scenario training and 500-scenario test above. Each seed draws 24 scenarios. The paired test draw for a seed uses that seed plus 5,000 and contains 80 scenarios. Demand is lognormal with mean 8,000 and coefficient of variation 0.35 unless a row says otherwise. Prices, the cap, the unmet-demand weight, and the administrative times are the synthetic inputs in Table VI. Table XIX reports the thirty-seed means. Table XX reports the paired cost differences. Table XXI reports the stress rows. Every method is scored by the same greedy service rule on its frozen first stage. Intervals are Student-t 95 percent intervals for the mean of these thirty seeds. A paired interval is the interval for the seed-wise difference, method minus the base conditional-value-at-risk plan. A paired t probability and, where the differences are not all zero, a Wilcoxon signed-rank probability summarize this simulator.

Table XIX shows that the base plan buys a prepaid committed block on every seed. Mean expenditure and conditional value-at-risk are both \$8,178.6, and mean service is 86.8 percent. The deterministic mean plan spends

TABLE XVIII
List-schedule durations for synthetic files that share a staff pool. Waiting is the mean delay between the time an activity's predecessors finish and the time it starts. Draws are paired across staff levels.

Files	Network	Staff	Mean file days	Mean wait	Portfolio days
1	Small	1	27.62	1.77	27.62
1	Small	2 or 4	22.68	0.00	22.68
1	Small, approval removed	2	20.77	0.00	20.77
1	Large	2	29.98	0.45	29.98
3	Small	2	40.77	4.00	41.40
3	Small	4	22.82	0.46	23.20
5	Large	2	102.90	11.14	105.20
5	Large	4	52.02	4.34	53.42

TABLE XIX
Thirty-seed out-of-sample means. Expenditure and conditional value-at-risk are in dollars. Unmet demand is in service units. Brackets are Student-t 95 percent intervals. Cap-feasible was 100 percent for every solved row in the source file. That check is mathematical, not statutory.

Method	E[cost]	$\text{CVaR}_{0.90}$	E[unmet]	Service (%)
Base CVaR	8,178.6 [7,897.7, 8,459.4]	8,178.6 [7,897.7, 8,459.4]	1,066.5 [902.7, 1,230.3]	86.8 [84.9, 88.7]
Deterministic mean	8,000.0 [8,000.0, 8,000.0]	8,000.0 [8,000.0, 8,000.0]	1,095.8 [1,021.9, 1,169.7]	86.4 [85.6, 87.1]
Risk-neutral	8,916.2 [8,791.5, 9,040.9]	14,973.4 [14,441.2, 15,505.5]	135.5 [59.0, 212.1]	98.3 [97.4, 99.3]
Fixed budget	16,000.0 [16,000.0, 16,000.0]	16,000.0 [16,000.0, 16,000.0]	22.3 [12.4, 32.3]	99.7 [99.6, 99.8]
Subscription	12,800.0 [12,800.0, 12,800.0]	12,800.0 [12,800.0, 12,800.0]	1,095.8 [1,021.9, 1,169.7]	86.4 [85.6, 87.1]
Pay as you go	10,531.7 [10,427.9, 10,635.5]	15,892.8 [15,825.5, 15,960.1]	214.2 [180.5, 247.9]	97.3 [97.0, 97.7]
Wasserstein, fixed times	8,178.6 [7,897.7, 8,459.4]	8,178.6 [7,897.7, 8,459.4]	1,066.5 [902.7, 1,230.3]	86.8 [84.9, 88.7]
Schedule only	8,178.6 [7,897.7, 8,459.4]	8,178.6 [7,897.7, 8,459.4]	1,066.5 [902.7, 1,230.3]	86.8 [84.9, 88.7]
Schedule and Wasserstein	8,178.6 [7,897.7, 8,459.4]	8,178.6 [7,897.7, 8,459.4]	1,066.5 [902.7, 1,230.3]	86.8 [84.9, 88.7]

TABLE XX
Paired cost differences, in dollars, against the base conditional-value-at-risk plan on the same thirty seeds. A negative difference means the named method spent less. A dash means the seed-wise cost gap is a numerical residual and is treated as zero.

Method	Δ cost	95% interval	t	Wilcoxon
Deterministic mean	-178.6	[-459.4, 102.3]	0.204	0.198
Risk-neutral	737.6	[401.6, 1,073.7]	< 0.001	< 0.001
Fixed budget	7,821.4	[7,540.6, 8,102.3]	< 0.001	< 0.001
Subscription	4,621.4	[4,340.6, 4,902.3]	< 0.001	< 0.001
Pay as you go	2,353.1	[2,022.9, 2,683.3]	< 0.001	< 0.001
Wasserstein, fixed times	0.0	[0.0, 0.0]	—	—
Schedule only	0.0	[0.0, 0.0]	—	—
Schedule and Wasserstein	0.0	[0.0, 0.0]	—	—

TABLE XXI
Stress of the configuration model on the same thirty shock pairs. The Wasserstein rows keep the fixed administrative times and change only the radius factor. Service is the mean share of test demand.

Case	Modal package	E[cost]	$\text{CVaR}_{0.90}$	Service (%)
CV 0.15	committed (30/30)	8,185.2	8,185.2	94.8
CV 1.00	committed (30/30)	6,978.2	6,978.2	62.2
Token \$0.90	meter (30/30)	7,207.9	12,420.1	99.9
Token \$2.50	committed (30/30)	8,178.6	8,178.6	86.8
Cap \$8,000	committed (30/30)	7,778.9	7,778.9	85.0
Cap \$30,000	committed (30/30)	8,178.6	8,178.6	86.8
$\beta = 0.80$	committed (30/30)	8,178.6	8,178.6	86.8
$\beta = 0.95$	committed (30/30)	8,178.6	8,178.6	86.8
Radius factor 0	committed (30/30)	8,178.6	8,178.6	86.8
Radius factor 0.25	committed (30/30)	8,178.6	8,178.6	86.8
Radius factor 1	committed (30/30)	8,178.6	8,178.6	86.8

\$8,000. Its paired cost difference is \$178.6 lower, and the 95 percent interval runs from \$459.4 lower to \$102.3 higher, so it contains zero. The risk-neutral plan selects the hybrid package on 29 of 30 seeds. Paired against the base plan, it spends \$737.6 more, raises the expenditure tail by \$6,794.8, and leaves 931.0 fewer units unmet.

Those three intervals exclude zero. The fixed-budget rule spends the $16,000 cap and serves 99.7 percent of test demand. Subscription spends $12,800 and matches the deterministic plan's service. Pay-as-you-go spends $2,353.1 more than the base plan and carries a higher tail while serving 97.3 percent.

Table XX and Table XXI are the comparisons that follow. The largest absolute seed-wise cost difference among the base plan, the fixed-time Wasserstein layer, and both schedule rows is below 10^{-11} dollars. Those rows are the same purchase. A paired probability is not reported for that residual. At a token price of $0.90, every seed selects the meter package: mean expenditure is $7,207.9 and the tail is $12,420.1. At a coefficient of variation of 1.00, mean service falls to 62.2 percent. Setting β to 0.80 or 0.95, raising the cap to $30,000, or moving the radius factor from 0 to 1 leaves the committed purchase and these out-of-sample means unchanged. A cap of $8,000 lowers mean expenditure to $7,778.9 and mean service to 85.0 percent.

On a prepaid credit plan the expenditure tail equals expected expenditure, so a Wasserstein layer on that expenditure has no random tail to reshape. Where a paired interval contains zero, the thirty seeds do not separate that method from the base plan on cost. Where an interval excludes zero, the difference is a property of these draws and these weights. A reimplementation of the published reservation-and-on-demand cloud model, a Wasserstein ball off the training atoms, and a coupled multi-file mixed-integer program are outside this comparison. The scores measure simulator performance rather than return on investment. A review of measured enterprise-AI productivity does not find an independently measured financial return for named enterprise assistants [47]. A governance-priority score whose Monte Carlo check is described as internal coding stability, rather than external validation, has the same limit: the intervals above describe this simulator [48]. A leakage-safe multi-seed design for a digital-twin enterprise-networking benchmark is a different application [49].

H. Component Contribution and Ablation Analysis

Each row removes one component from the scheduled finite-support plan and is scored on the thirty test draws of Section X-G. The solves use `experiments/run_ablation.py` and the solvers in `run_sim30.py`, with the same seeds and the same 24-scenario training and 80-scenario test draws. The complete row keeps the activity-network times, the finite-support Wasserstein layer at radius factor 0.50, and risk weight $\lambda = 0.50$. The no-CVaR row keeps the schedule and the Wasserstein layer and sets $\lambda = 0$. The no-Wasserstein row keeps the schedule and $\lambda = 0.50$ and uses the nominal scenario weights. The no-schedule row keeps the Wasserstein layer and $\lambda = 0.50$ and substitutes the fixed configuration times. The deterministic row solves one scenario at the theoretical mean of 8,000, with $\lambda = 0$, and then serves the same test draw. Intervals are Student-t 95 percent intervals for the mean of the thirty seeds. A paired difference below 10^{-6} is reported as zero. The sample is synthetic.

The complete plan buys a prepaid credit block on the scheduled committed package in all thirty seeds. Mean quantity, mean expenditure, and conditional value-at-risk are the same number, $8,178.6, because the credit price is $1 and the outlay does not depend on the scenario. Mean service is 86.8 percent. Mean cycle time is 22.5 days and mean staff effort is 166.4 hours. No seed pre-authorizes tokens.

Removing the CVaR weight, or removing the Wasserstein layer, does not change that purchase. The paired differences in cost, conditional value-at-risk, unmet demand, service, days, and hours are numerical residuals and are treated as zero. The reason is visible in the selected plan. A prepaid block has no expenditure tail distinct from its mean, so a tail weight and a finite-support ball on that expenditure have no random cost to reshape. On the fixed-time configuration model in Table XIX, the same thirty seeds with $\lambda = 0$ and no Wasserstein layer select the hybrid package on 29 of 30 seeds, raise mean expenditure to $8,916.2, and raise the expenditure tail to $14,973.4. The risk weight changes the package when the fixed administrative times leave the token tail open. It leaves the package unchanged once the scheduled plan has selected the prepaid block.

Removing the activity schedule also leaves the credit quantity and the expenditure tail unchanged. What changes is the time source. The fixed configuration table assigns the committed package 30.0 days and 55.0 hours. The schedule assigns the same purchase a mean of 22.5 days and 166.4 hours. Paired against the complete plan, the fixed-time row is 7.5 days longer, with interval $[7.0, 8.0]$, and 111.4 hours lower, with interval $[-114.8, -108.0]$. Both intervals exclude zero. The 55-hour figure is the configuration-table assumption, not the unscheduled sum of the same activities, so the hour gap is not a labor saving. The within-file comparison holds work content fixed: a parallel schedule shortens cycle time and does not reduce the activity hours. On these seeds, the schedule changes the time attached to the purchase and does not change the purchase. Section X-I then varies only the staff limit on the committed activity set.

The deterministic mean plan buys 8,000.0 credits on every seed, spends $8,000.0, and has the same conditional value-at-risk. Mean unmet demand is 1,095.8 units and mean service is 86.4 percent. Paired against the complete plan, the cost difference is $178.6 lower, with interval $[-459.4, 102.3]$, and the unmet-demand difference is 29.3 units higher, with interval $[-92.6, 151.1]$. Both intervals contain zero. The decision difference is still visible: every deterministic seed buys the theoretical mean, while the complete plan's credit quantity moves with the training sample. Its 30.0 days and 55.0 hours come from the fixed configuration table, not from the activity network. On this design the deterministic baseline is a smaller prepaid block. These thirty seeds do not separate it from the

TABLE XXII

One-factor ablation on the thirty-seed design. Entries are means. Expenditure and conditional value-at-risk are out-of-sample dollars. Unmet demand is in service units. Service is the share of test demand served. Days and hours are the assumed cycle time and staff effort attached to the frozen plan. Student-t intervals are stated in the text. Every solved row was cap-feasible under the model constraint.

Plan	Package (30 seeds)	E[cost]	$\text{CVaR}_{0.90}$	E[unmet]	Service	Days	Hours
Complete	committed, scheduled	8,178.6	8,178.6	1,066.5	86.8%	22.5	166.4
No CVaR weight	committed, scheduled	8,178.6	8,178.6	1,066.5	86.8%	22.5	166.4
No Wasserstein layer	committed, scheduled	8,178.6	8,178.6	1,066.5	86.8%	22.5	166.4
No activity schedule	committed, fixed times	8,178.6	8,178.6	1,066.5	86.8%	30.0	55.0
Deterministic mean	committed, fixed times	8,000.0	8,000.0	1,095.8	86.4%	30.0	55.0

TABLE XXIII

One-factor changes to the scheduled plan on the thirty-seed design. Entries are means. A row marked "same purchase" matches the reference package and the reference quantities on all thirty seeds. Hours belong to the selected activity set. Student-t intervals for the paired differences are in the text.

Case	Package	E[cost]	$\text{CVaR}_{0.90}$	Service	Days	Hours
Reference	committed, 30/30	8,178.6	8,178.6	86.8%	22.5	166.4
$\lambda = 0$, radius 0	hybrid, 30/30	8,955.3	15,226.0	98.8%	23.7	192.4
$\lambda = 0$, radius 0.25	hybrid 16, committed 14	8,748.1	12,047.4	94.0%	23.0	179.3
CV 0.15	committed, 30/30	8,185.2	8,185.2	94.8%	22.5	166.4
CV 1.00	committed, 30/30	6,978.2	6,978.2	62.2%	22.5	166.4
Token \$0.90	meter, 30/30	7,207.9	12,420.1	99.9%	23.9	168.6
Cap \$8,000	committed, 30/30	7,778.9	7,778.9	85.0%	22.5	166.4
One staff	committed, serial, 30/30	8,178.6	8,178.6	86.8%	27.7	166.4
Four staff	same purchase	8,178.6	8,178.6	86.8%	22.5	166.4

complete plan on expected cost or unmet demand.

I. When the Scheduled Plan Changes the Purchase

The rows below use the thirty seeds, the 24-scenario training draws, and the 80-scenario test draws of Section X-G. The reference row is the integrated plan: activity-network times at two staff, radius factor 0.50, and $\lambda = 0.50$. Each other row changes one input, and Table XXIII lists them. The solves are in `experiments/run_decision_grid.py`. A fixed-time risk-neutral anchor, solved on the same draws, reproduced the published mean expenditure \$8,916.2. Scheduled solves returned a relative optimality gap no larger than 4.8×10^{-7}.

With the radius factor fixed at 0.50, the reference prepaid purchase remains stable when the risk weight is removed, and the same purchase is selected at $\lambda = 0.25$ and at $\lambda = 1$, and at radius factors 0, 0.25, and 1 when $\lambda = 0.50$. On each of those settings the package and the quantities agree with the reference on all thirty seeds, and the paired differences are zero. That agreement is solution stability of the prepaid block under risk-weight and radius perturbations that leave expenditure nonrandom.

The package changes when $\lambda = 0$ and the radius is small—the demonstrated regime in which open expenditure risk and weak ambiguity pressure favor token recourse. With no Wasserstein layer, all thirty seeds select hybrid TO-BE. Mean service is 98.8 percent and mean $\text{CVaR}_{0.90}$ is \$15,226.0, against 86.8 percent and \$8,178.6 for the reference. Paired against the reference, expenditure is \$776.8 higher, with interval [443.6, 1,109.9]; the tail is \$7,047.5 higher, with interval [6,637.0, 7,458.0]; unmet demand is 966.6 units lower, with interval [817.0, 1,116.3]; and service is 12.0 percentage points higher, with interval [10.2, 13.7]. All four intervals exclude zero. At radius factor 0.25 the split is 16 hybrid seeds and 14 committed seeds. Mean service is 94.0 percent and mean $\text{CVaR}_{0.90}$ is \$12,047.4. The paired expenditure difference is \$569.6, with interval [224.6, 914.5], and the paired service difference is 7.2 percentage points, with interval [4.2, 10.2]. Both exclude zero. At radius factor 0.50 or 1, with $\lambda = 0$, every seed returns to the reference credit block: a large enough ambiguity radius restores the prepaid optimum by making the hybrid tail unattractive. The higher service is the hybrid plan's token recourse. The higher tail is the expenditure that recourse leaves open. The two rows trade service against expenditure risk; neither is claimed as uniformly better.

The extra hours on the hybrid row are the token activities in that package. They are not a staff effect. On the committed activity set, one, two, and four staff all record 166.4 hours. One staff schedules that set serially: the credit quantity matches the reference on every seed, and cycle time is 5.2 days longer, with interval [5.0, 5.4]. Four staff matches the two-staff purchase and the two-staff cycle time on every seed. The 55.0 hours in Table XXII remain the fixed configuration table. They are not the labor content of this network.

A token price of \$0.90 selects the meter package on all thirty seeds. Paired expenditure is \$970.7 lower, with interval [647.8, 1,293.6], and paired $\text{CVaR}_{0.90}$ is \$4,241.5 higher, with interval [3,780.9, 4,702.2]. Service is 13.1 percentage points higher, with interval [11.2, 15.0]. A token price of \$2.50 leaves the reference purchase unchanged on every seed. A cap of \$30,000 does the same. A cap of

TABLE XXIV
MATHEMATICAL FOUNDATIONS USED IN THE INTEGRATED DECISION MODEL, AND THE SCOPE OF THEIR USE.

Source	What it solves	What this paper uses it for	Scope and open extensions
Two-stage stochastic programming [12]	Decisions before a random outcome is known, and recourse after it is known.	First-stage package and prepaid quantity. Second-stage use, unmet demand, and incremental funding.	Used as formulated; the contribution lies in the coupling with package choice and governance limits.
Coherent risk measures [50]	Axioms that quantile-based value-at-risk fails, including subadditivity.	Motivation for using conditional value-at-risk on expenditure rather than a quantile alone.	Used as the risk measure of the expenditure tail.
Conditional value-at-risk [13]	A tail expectation with a linear representation on a finite scenario set.	The expenditure tail in (12), at a stated confidence level, on an equiprobable sample.	The tail is expenditure only. Unmet demand is a separate expected penalty.
Wasserstein distributionally robust optimization [18]	A tractable robust counterpart for a Wasserstein ball, with performance guarantees for that formulation.	A finite-support 1-Wasserstein ball on the training atoms, with separate worst-case bounds for expected expenditure and for the conditional-value-at-risk auxiliary.	Specified on the training atoms; a ball with mass off the atoms is an open extension. The separate bounds can charge two different worst-case distributions, and Proposition 2 bounds the resulting gap. Finite-sample guarantees of the unrestricted formulation are not asserted for the atom-restricted ball.
Resource-constrained project scheduling [51]	Classification and exact or heuristic methods for scheduling activities under resource limits.	Makespan of one synthetic file, by critical path, a list schedule, or a time-indexed mixed-integer program when the list schedule is not already optimal.	Multi-file sharing of one cap is not formulated as one coupled mixed-integer program. When the list schedule is the method used, optimality is not established.
Cloud reservation and on-demand provisioning [52]	A stochastic program for reserved and on-demand cloud capacity under demand and price uncertainty.	A comparable published pattern for paying ahead versus paying for measured use. The code from that paper was not reimplemented.	No Texas appropriation, encumbrance, ethics, or human-authorization gate. The comparison is conceptual, and no numerical comparison is reported.

$8,000 keeps the committed package and lowers the credit quantity on 19 of 30 seeds. Paired expenditure is $399.7 lower, with interval [225.2, 574.1], and paired service is 1.8 percentage points lower, with interval [1.1, 2.5]. Both exclude zero.

At a coefficient of variation of 0.15 or 1.00 the package remains a prepaid credit block, and the credit quantity changes on every seed. These service rates are computed on the same shocks after the shocks have been rescaled to that coefficient. They are not two policies scored on one demand sample. At coefficient 0.15 the paired expenditure difference is $6.7, with interval [−149.6, 162.9], which contains zero. The quantities still differ on every seed, so the interval is not evidence that the expenditures are equal. Paired service on the rescaled draw is 8.0 percentage points higher, with interval [6.9, 9.0]. At coefficient 1.00, paired expenditure is $1,200.3 lower, with interval [905.4, 1,495.3], and paired service is 24.6 percentage points lower, with interval [22.6, 26.7].

XI. Comparative Methodological Positioning

Table XXIV positions the integrated model against the mathematical foundations it builds on, and states the scope of each use. The contribution stated above is the joint decision formulation for risk-aware AI procurement—package, quantities, funded amount, schedule, CVaR, and finite-support Wasserstein ambiguity under Texas-labeled funding and authorization constraints—assembled from the constituent methods in the table. The Texas labels in Section VIII stay outside the solver. The table is a focused positioning of the methods the formulation uses rather than a systematic review.

The integrated application is therefore the joint choice, under the funding cap and the human-authorization restriction labeled as a recommended control, among subscription, credit, and token quantities and among the permissible administrative packages and schedules. That coupling is what enables a single recommendation that is simultaneously risk-aware, funding-feasible, authorization-feasible, and schedule-coordinated. A feasible mathematical plan is not a statutory approval.

XII. Governance-Constrained Multi-Objective Extension

The validated stochastic-CVaR, Wasserstein, funding-authorization, and scheduling blocks are retained as the scientific foundation. This section reports a computational extension that couples those blocks to an explicit governance layer. The full COBIT 2019 map of all forty objectives, the complete epsilon-constraint tables, and the implementation details appear in the supplementary material. Vendor residual-risk scores used below are synthetic labels rather than institutional assessments, COBIT maturity scores, or compliance outcomes.

A. Coupling and Hard Constraints

Let z_v select a vendor from the simulated catalog and y_k a package. Financial recourse, CVaR, the officer-supplied

funding cap, and human authorization remain as in Sections IX–IX-E. Residual risk is the hard inequality

$$\rho = \sum_v \rho_v z_v + \sum_k \rho_k^{\mathrm{pkg}} y_k \le \bar{\rho}, \tag{29}$$

with the assigned scores $(\rho_v, \rho_k^{\mathrm{pkg}})$. TX-RAMP, security-screen, and architecture clearances are exogenous hard gates: if a required gate fails, the award program is infeasible. The solver cannot accept residual risk above $\bar{\rho}$. Statutory and nondelegable human controls therefore remain constraints, not objective penalties.

B. Why Some Optima Coincide

Three facts explain coincident solutions without removing model components. (i) If $L_s = W$ in every scenario, then $\mathrm{CVaR}_\beta(L) = W$ for every β. (ii) If expenditure is constant on the training atoms, Corollary 1 leaves the expenditure-side dual inactive at every finite-support radius, so the robust and nominal prepaid optima coincide. (iii) If the administratively cheapest package also minimizes the joint administrative-plus-financial objective, the admin-first sequence and the integrated rule agree. Agreement with a finance-first sequence requires a different condition: the financially cheapest package must also minimize that joint objective. Facts (i) and the admin-first case of (iii) are proved in the supplement; fact (ii) is the constant-loss case of Corollary 1. On the base 80-scenario training draw, the admin-first case of (iii) holds for the committed package. On the published extended draw at $\lambda = 0.50$, finance-first and integrated selection do not agree (Section X-F2). These identities are mathematical properties of prepaid plans and of additive process charges. When expenditure is random, the CVaR and Wasserstein layers are active, as the switching experiments show.

C. Executable Comparisons

All new solves reuse the validated 80/500 seeds. The original stochastic-CVaR baseline remains a committed credit purchase with out-of-sample cost and CVaR equal to \$8,399.5 and service 89.1 percent. With only the baseline vendor admitted, the governance layer recovers that same purchase. Raising the residual-risk appetite to $\bar{\rho} = 0.70$ switches the vendor to the higher-risk catalog entry and changes the funded amount to \$8,550.6. Tightening to $\bar{\rho} = 0.15$ forces the low-risk high-price vendor (\$9,619.7). Failed TX-RAMP or security-screen gates make every plan infeasible. An epsilon service floor of 0.90 under cost minimization selects a hybrid plan on the higher-risk vendor (training service 93.2 percent). A 0.97 service floor is infeasible at the base risk weight because no admissible pair reaches that training service. Combining service ≥ 0.90 with CVaR $\le$ \$9,000 is likewise infeasible: high-service pairs violate the CVaR cap, while CVaR-feasible prepaid pairs miss the service floor. Risk weight zero with $\bar{\rho} = 0.70$ selects hybrid and raises test service to 99.5 percent at CVaR \$13,559.7. A token price of \$0.90 on the baseline vendor switches to meter (test CVaR \$8,728.4, service 94.0 percent). A \$8,000 funding cap binds at cost \$8,000.

Epsilon-constraint enumeration over vendor–package pairs and governance caps produced six nondominated training points trading cost, CVaR, service, days, hours, and residual risk. Ablation on identical draws shows: removing vendor choice restores the baseline vendor; minimizing residual risk alone selects the low-risk vendor; minimizing cost alone selects the high-service hybrid; and the service-floor ablation reproduces the infeasibility above. The published scheduling and Wasserstein rows from the distinct 40/200 sample are cited for baseline comparison and are not recomputed here. The demonstrated contribution is the coupling—hard governance gates, vendor choice, and multi-objective epsilon filters—together with measurable decision changes on shared scenarios.

XIII. Scope and Applicability

The procurement, CVaR, Wasserstein, scheduling, and governance-finance blocks above remain the scientific core. This section defines the decision scope within which those results apply.

The framework informs decisions to buy, adopt, expand, or discontinue AI services; it does not make them. A mathematically efficient configuration is a planning input, and purchase recommendations, legal authorizations, and CISO determinations remain with the authorized officials. The computational results characterize model behavior and do not estimate institutional cost savings, productivity gains, or compliance improvements. AI adoption comparisons, activity-level tooling placement, and organizational return-on-investment analysis fall outside the central manuscript. Exploratory material of that kind in code or supplementary notes is labeled as simulation-based and does not modify the validated \$8,399.5 baseline or the conclusions of this paper.

What remains in scope is transparent comparison of subscription, consumption-based, prepaid-credit, and hybrid arrangements under identical assumptions; quantification of expected expenditure, CVaR, Wasserstein sensitivity, budget exposure, and funding requirements; and separation of financial risk measures from hard governance constraints. Officers may use those quantities only after the legal stops in Section VIII, and only as planning estimates.

XIV. Discussion

The framework changes the handoff, not the statute. Each office already has a Texas duty. What changes is that the file separates that duty from the recommended controls for an uncertain meter.

Procurement can refuse to place a rate-priced buy in a value band until the file states which figure is the amount that may be obligated, and why. A GSA or vendor "no minimum" offer is not permission to skip DIR, TX-RAMP, or encumbrance rules [1], [9], [31], [39]. The officer still cannot treat the recommended ceiling as the statutory value. The statute does not say that.

Finance can refuse to certify the unused ceiling as a binding encumbrance unless counsel concludes that the order obligates that amount, and can refuse incremental funding that is only an earmark [8], [9]. The framework asks for a written theory of what is owed today, matched to the appropriation year. That theory is a recommended control rather than a Comptroller ruling on token contracts. Capital-budget riders remain an agency-specific overlay [24].

Information-technology and security staff can keep the price question out of the cloud question. If the service is covered cloud, there is still no award or renewal without TX-RAMP [29], [31]. A model change can be a price event, as GSA officials told GAO [11], and, when heightened scrutiny applies, a fresh Texas risk decision [2], [4]. The framework asks for both records. It does not merge them into one approval, and it does not make a quota or an alert a security finding. A systematic review of enterprise generative-AI security risks treats prompt injection, retrieval-corpus poisoning, and excessive agency as exposures that ordinary boundary controls do not close, and it offers its classification instruments as synthesis tools rather than as predictions [53]. Because that review addresses security exposure rather than token pricing or the contract value left open here, it supports keeping the security file separate from the expenditure file: a cloud-authorization record and a heightened-scrutiny decision are not a forecast of spend, and the planning model does not score those security exposures.

Executives can use the decision rule as a stop on the next commitment. Before award, renewal, or a higher ceiling, the packet shows whether the use screen, the vehicle, the encumbrance, and the funded amount are present, and whether a named person must approve any move past the quota or the alert. Human approval, the quota, and the alert are recommended practices rather than offices created by Texas law. A preprint on cybersecurity controls in information-technology organizations treats access that is not tied to accountable attribution as an internal control failure, and it proposes a competence-based promotion framework drawn from ISACA and Control Objectives for Information and Related Technologies (COBIT) 2019 [54]. The named approval in this procedure is narrower. It identifies who may authorize further funded use of a metered service. It is distinct from that promotion framework, and the preprint does not evaluate the planning model. A failed recommended gate stops the larger commitment. It does not cancel an obligation already incurred, and it does not replace the legal stops that apply when a use is prohibited, a required statement of work is unsigned, TX-RAMP is missing, or heightened-scrutiny approval is absent.

The 89th Legislature regulated use, ethics, and risk more directly than price [2], [3], [4]. Institutions can run Section VIII without new legislation. A ceiling method, a USAS treatment, or a usage alert becomes a statewide duty only if the Comptroller, DIR, or the Legislature states it. Local governments, institutions of higher education, and public junior colleges do not share one set of Chapter 2157 and Chapter 552 duties [1], [3], [4], [6], [10]. The encumbrance logic is the piece that most nearly generalizes across entities that spend appropriated state funds [7], [8], [9]. The variable-cost sequence can transfer further, as stated above, only if the receiving institution substitutes its own legal stops.

The planning layer does not change that account. It supplies numerical planning estimates—mix, configuration or schedule, and funded amount—after the legal stops have been applied. In the base case of the numerical study the selected configuration is a credit purchase on the committed package. Expected expenditure and conditional value-at-risk coincide because the outlay is prepaid, which is exactly the setting in which CVaR, Wasserstein transport on expenditure, and admin-first package selection can share the same optimum without contradiction. A finance-first sequence need not share that optimum when a financial near-tie is smaller than the process-cost gap (Section X-F2). The expected-demand rule is within a few dollars of that plan on the weighted score and leaves more demand unmet. The fixed-budget and subscription rules spend more. Pay-as-you-go spends a random amount, carries a higher tail, covers more of the test demand, and uses more staff time. Those contrasts follow from the assumed prices and weights. They answer the financial-clarity research questions by showing how the joint objective trades expenditure, tail risk, and service under one authorization-feasible cap.

The weight grid in Table XIII is the practical form of that contrast. An officer who requires a nonrandom expenditure stays on a prepaid row. The solved prepaid outlays are \$0, \$5,774.6, \$8,000.0, \$8,399.5, \$11,233.7, \$13,239.1, and the \$16,000 cap. Service coverage rises with that outlay, from zero to 99.60 percent. An officer who can accept a tail near \$15,000 can inspect a hybrid row with service between 98.78 and 98.97 percent at an expected cost between \$8,803.0 and \$8,975.4. No solved hybrid has a tail as low as the \$13,239.1 credit block, and no solved prepaid row reaches 99.60 percent service at less than the cap. The subscription rule and the pay-as-you-go rule remain available configurations; on this sample both are dominated on the stated objective. That statement applies to the solved points on this sample. At the published weight $w_m = \$2.50$, every λ from 0.25 to 1 returns the same credit block: once any positive weight is placed on the expenditure tail, the prepaid planning estimate is stable across that risk-weight range. Stability of that kind is useful when an officer needs a funded-amount figure that does not move with every small change in λ.

Across the first grid, token recourse and a positive conditional-value-at-risk weight change the configuration in two demonstrated cases: a zero risk weight, and a token price of \$0.90. Elsewhere the optimum is the committed credit block. The extended grid has the same pattern. The Wasserstein objective changes the package when expenditure is random and the risk weight is zero. Corollary 1 explains why a prepaid credit plan is unchanged: a constant expenditure makes the dual multiplier on transport

optimal at zero for every radius. Corollary 2 gives the frozen-profile threshold. On the published profiles that threshold is 639.6 demand units, between the reported factors at which the solved package changes, and Proposition 1 derives it as $863.5 divided by the token price $1.35. Proposition 2 bounds how far the sum of separate worst-case suprema can exceed the single-distribution worst case over the same set, and gives $7.2 on the frozen hybrid profile. That identity preserves the nominal optimum, and the layer remains an active part of the model whenever expenditure is random.

Taken together, the analytical and computational results show what the integrated formulation adds to sequential practice, and they bound that addition. Relative to an admin-first sequence, joint selection surfaces decisions that that sequence cannot, such as accepting a higher expenditure tail for higher service when the risk weight is zero. Relative to a finance-first sequence, it surfaces a different class of decision: a process-cheaper prepaid package when the financial blocks are nearly tied. Both contrasts are conditional on the simulated sample and the additive process charge rather than general dominance results. The analysis explains when the risk and ambiguity layers are decision-relevant and when they are not: expenditure must take at least two values on the training atoms. The closed-form critical radius gives an officer a threshold with a direct reading, the advantage of the exposed profile at radius zero divided by the token price. Hard governance gates enter as constraints, so a failed gate cannot be traded against expenditure.

In the configuration model, cycle time and staff hours change when the configuration changes. The extended model schedules one file. Overlap at two staff reduces expected cycle time and does not reduce labor, separating schedule value from labor content. At a zero risk weight the finite-support Wasserstein model selects the credit plan where the nominal model selects hybrid, demonstrating that ambiguity can change the package when the tail is open. At the base prepaid plan the same layer leaves the optimum unchanged, which is the inactive-constraint case above. These objective values are model-generated, and the overlap leaves every mandatory precedence intact.

Section X-F separates three questions that the first tables leave open. Joint selection improves the training objective relative to the admin-first sequence at a zero risk weight, and on that run it accepts a higher out-of-sample expenditure tail in return for higher service—a demonstrated trade-off that an admin-first sequence cannot surface. At risk weights 0.25 and 0.50 on the published training draw, the admin-first and integrated rules select the same package because the administratively cheapest package is also the joint optimum under those risk-aware weights. On that same draw the finance-first rule selects hybrid TO-BE. Section X-F2 reports when that disagreement survives out of sample and when it does not. The finite-support Wasserstein rule lowers out-of-sample CVaR when it replaces a random hybrid plan with a prepaid credit plan. At risk weight 0.50 on the distributions tried here the prepaid plan is already selected, so the robust and nominal decisions coincide. One file with two staff has no resource wait. Three or five files on that same pool do. Raising the shared pool from two staff to four removes most of that wait for three small files and only part of it for five large files.

On the fixed-time model, risk-neutral service is 98.3 percent and base service is 86.8 percent. The paired intervals for expenditure, the expenditure tail, and service exclude zero: the risk-neutral plan spends $737.6 more, carries a tail $6,794.8 higher, and serves 11.5 percentage points more of the same test demand. Section X-I records the same exchange on the activity network when the risk weight and the radius are both removed, with service 98.8 percent and a tail of $15,226.0. Token recourse raises service. The open tail is the cost of that recourse. Neither plan is better on both criteria. At radius factor 0.50 the scheduled plan keeps the prepaid block for every risk weight from 0 to 1. A token price of $2.50 and a cap of $30,000 do the same. A coefficient of variation of 0.15 changes the credit quantity on every seed. The paired expenditure interval contains zero and does not establish equal expenditure. One staff member lengthens the committed file and leaves its 166.4 hours and its credit purchase unchanged. The 55-hour entry in Table XXII is the configuration table, not the labor content of the activity network.

XV. Limitations

This is a source study rather than an audit. It does not use Texas invoice data, interviews, or institutional procurement records. It therefore does not measure how frequently valuation gaps, encumbrance disputes, threshold crossings, or forecast errors occur in Texas practice, and the effect of the interim procedure on cost and compliance remains to be measured. GAO's findings are federal, interview-based, and expressly nongeneralizable [11]. Silence in the searched statewide sources is not proof that no agency manual addresses metered AI. Quotations from the Comptroller's Procurement and Contract Management Guide are from Version 4.0, publication 96-1809, read on September 23, 2026 [16], [22]. The guide is long, and this paper quotes only the passages that bear on value, statements of work, data-center services, change control, indefinite contracts, and cross-year terms. GAA riders differ by agency and biennium; capital-budget and last-quarter conclusions are stated at the level of Comptroller guidance, not as a reading of every rider [9], [24]. Coverage of institutions of higher education follows Chapters 2054 and 2157, Business and Commerce Code Chapter 552, 1 Tex. Admin. Code § 212.1(7), and Education Code §§ 51.9335 and 51.9337. Whether a named institution currently meets § 51.9337 is not determined. The sandbox's operational status was not confirmed beyond the statute and the DIR page reviewed [2], [36]. DIR's statement-of-work application guide is a 2022 portal instruction. It is cited for its amendment sentence and is not treated as a 2026 amendment of § 2157.0685 [45]. Federal consumption

procedures cited here are cloud procedures on a GSA schedule [15]. They are analogues, not Texas authority. The 50 percent and 75 percent alerts in the procedure are taken from that federal clause and are labeled as recommended controls.

The planning calculation adds limitations of its own. Demand is an independent draw from a lognormal distribution with a chosen mean and coefficient of variation. Prices, the unmet-demand weight, the staff weight, the delay weight, the pre-authorization fraction, the hours and days attached to each configuration, and the governance-layer vendor residual-risk scores are synthetic. The finance-first comparison in Section X-F2 uses those same prices and an additive staff-and-delay charge on one file; it is a single-file sensitivity analysis that complements the validated tables rather than a multi-file or shared-resource test. The legal sources do not report those quantities, and this paper does not impute them to any institution. The COBIT 2019 mapping classifies objectives relative to the optimization model; it does not score institutional maturity, and it does not assert that all forty objectives are decision variables. The thirty-seed design in Section X-G is a separate sample from the 80/500 design. It does not estimate an effect for a procurement office. A systematic review of enterprise AI productivity and return on investment did not identify an independently measured financial return for named enterprise assistants [47]. The simulation scores do not estimate that return or organizational savings. The base optimum does not open a token stage. The base table therefore demonstrates that the joint model selects and stabilizes a prepaid, authorization-feasible planning estimate under positive risk weight, and the comparison with a deterministic credit purchase is reported metric by metric. The weight grid is the nondominated subset of the solved weights and the four reference rules on one sample. It is not a complete Pareto set of the mixed-integer program, and it does not rank cycle time. The $3.8 score difference against the expected-demand rule can change sign on another sample. The $10 million figure is imposed as an exogenous cap and is slack; the runs do not test how a buy near that figure should be valued. Cap feasibility is a check on the mathematical constraints rather than a compliance rate, and a solver outcome on these draws is a planning result rather than an institutional outcome. A paired interval that contains zero means these seeds do not separate the two means. It does not show that the means are equal. The out-of-sample scores are not the optimality-gap intervals of Mak, Morton, and Wood [41]. The critical radius in Corollary 2 is computed for the published profiles on the 40-atom training draw. It is the decision threshold between two fixed scores on the atom-restricted ambiguity set, and a tie at that radius means either profile is optimal. Extending the analysis to a ball with mass off the training atoms is a direction for future work.

XVI. Future Research

Empirical validation of the framework is the natural next step, and three designs are available. First, expert review: procurement, finance, information-technology, security, and legal officers review the decision table and mark each recommended control as usable, unusable, or already required by an agency procedure this study did not collect. Second, historical procurement cases: compare the framework's placement, ceiling, and encumbrance branches with completed usage-priced orders, using invoices and encumbrance records rather than a reconstructed narrative. A preprint on auditing operational data records argues that help-desk, infrastructure, and cloud summaries, including cost and security-event reports, can misstate the facts used for budget and compliance decisions when the source records are incomplete or inconsistent, and it treats a system-generated dashboard as information that still needs a reliability check [55]. That is why an exported usage total cannot be entered as demand, or as the legal obligation, until the record itself has been checked. That preprint addresses operational records rather than procurement files of the kind studied here. Third, application to real AI purchasing decisions before award, recording which artifact was present and whether the gate output matched the decision the institution made. Those designs can also ask how often a § 2157.068 threshold is crossed after award, what amount is encumbered on pay-for-use orders, and whether post-payment audit treats an unused ceiling as over-encumbrance. A separate legal question, flagged here for future analysis, is whether § 2157.068(e-2)'s $10 million limit is applied to the ceiling, the expected value, or cumulative invoices. Until invoice, interview, or case evidence exists, ceilings, quotas, alerts, incremental funding, and human approval remain recommended controls.

The computational analysis complements those three empirical steps. A later study can replace the lognormal demand, the unit prices, and the assumed durations with invoice and time-study records. The multi-file times in Section X-F come from a list schedule, not from a proved optimum, and the three-file budget comparison is an equal split of the cap rather than one coupled mixed-integer program. Solving the Wasserstein layer for a ball that may leave the training atoms, which is formulated here, is a further extension. The tables in Sections X and X-E report solutions of the computational models; testing the legal framework itself requires the empirical designs above.

XVII. Conclusion

This paper contributes a Texas-labeled operating framework for financial clarity when variable-cost AI is priced at a known rate and an unknown quantity, together with an integrated, analytically characterized risk-aware decision model that selects among subscription, consumption, prepaid-credit, and hybrid configurations and reports expected expenditure, CVaR, budget exposure, funding requirements, and—in the extended model—activity timing under funding and human-authorization limits. The

scientific objective is transparent, reproducible analysis of AI procurement decisions under uncertainty. Procurement can place a rate-priced buy only after the file states the amount that may be obligated. Finance can encumber the present legal obligation and treat a higher ceiling, a quota, an alert, and incremental funding as recommended controls rather than as Texas mandates. Information-technology and security staff can keep TX-RAMP and heightened-scrutiny approval separate from the price file. An executive can withhold the next commitment until a named person approves movement past the funded amount. Those actions sit inside duties the law already imposes and add no duties beyond those the reviewed sources state.

Texas-specific stops do not travel: commodity value bands, the DIR statement-of-work rule, TX-RAMP, Chapter 219, Chapter 552's coverage, and Chapter 403 encumbrance and lapse. The variable-cost problem does travel. A known rate and an unknown quantity force a choice among an estimate, a ceiling, and the amount legally owed. GSA's consumption procedure offers a transferable management pattern for that choice, although it is not Texas law.

The evidence base is documentary: the paper draws on statutes, rules, and published guidance rather than Texas invoice data, interviews, or institutional procurement records, so the frequency of the gap remains to be measured. Future work can validate the framework with expert review, historical procurement cases, or application to real AI purchasing files. Until that evidence exists, ceilings, quotas, alerts, incremental funding, and human approval should be labeled as practice.

The decision model supplies the quantitative component of that practice, a calculation the framework may call. In the base case of the first model it selects a credit block and does not open a token stage, yielding a prepaid plan whose expenditure equals its conditional value-at-risk. In the first model's cases with no weight on tail expenditure, or with a token priced below the credit, it selects a configuration that uses tokens and records a higher expenditure tail—demonstrated settings in which risk sensitivity and meter recourse change the planning optimum. The extended experiment, on a different sample, schedules one file: parallel overlap shortens expected cycle time and does not cut staff hours, while the joint model keeps the financial configuration feasible under the same caps. The finite-support Wasserstein layer changes the package when the risk weight is zero and the radius passes the critical value in Corollary 2; on the published profiles that value is 639.6 demand units, which Proposition 1 derives in closed form and an independent transportation linear program confirms. Proposition 2 shows that the separate-suprema risk charge coincides with the single-distribution worst case when the loss is constant or the radius is zero, and exceeds it by at most \$7.2 on the frozen hybrid profile. When the base credit plan is already prepaid, Corollary 1 explains why that optimum is stable: expenditure is constant on the training atoms. Several files on one staff pool wait. One file with two staff does not. The governance-constrained multi-objective extension preserves that validated baseline and, on the same seeds, demonstrates additional decision changes under residual-risk caps, hard security gates, service and CVaR epsilon filters, vendor choice, and ablations. Full COBIT maps and Pareto tables are in the supplement. Those results answer the computational research questions through the reported solves and paired intervals, and together with the formulation and the analytical results they form one body of evidence. Solver outputs are planning estimates rather than contract values, encumbrances, CISO determinations, or purchase recommendations, and an officer uses them only after the legal stops in Section VIII have been applied. The framework labels the Texas questions that the reviewed sources leave open so that agencies and future research can address them directly.

Author Contributions

Both authors contributed to the development and preparation of this manuscript.

Funding

This research received no external funding.

Institutional Review Board Statement

Not applicable.

Informed Consent Statement

Not applicable.

Data Availability Statement

The synthetic data, model assumptions, and numerical results supporting this study are described in the article and its supplementary materials. Reproducibility materials are available from the corresponding author upon reasonable request.

Conflicts of Interest

The authors declare no conflicts of interest.

Acknowledgments

The authors gratefully acknowledge the University of Houston for providing an academic and professional environment supportive of interdisciplinary research in artificial intelligence, engineering management, and technology governance.

References

[1] Texas Legislature, *Texas government code chapter 2157, purchase of automated information systems*, Texas Constitution and Statutes, 2026. Accessed: Sep. 22, 2026. [Online]. Available: https://statutes.capitol.texas.gov/Docs/GV/htm/GV.2157.htm

[2] Texas Legislature, *Senate bill 1964, 89th legislature, regular session, enrolled*, Texas Legislature Online, Effective Sept. 1, 2025. Adds Tex. Gov't Code ch. 2054, subch. S, 2025. Accessed: Sep. 22, 2026. [Online]. Available: https://capitol.texas.gov/tlodocs/89R/billtext/html/SB01964F.HTM

[3] Texas Legislature, *House bill 149, 89th legislature, regular session, enrolled*, Texas Legislature Online, Adds Tex. Bus. & Com. Code ch. 552. Effective Jan. 1, 2026, 2025. Accessed: Sep. 22, 2026. [Online]. Available: https://capitol.texas.gov/tlodocs/89R/billtext/html/HB00149F.htm

[4] Texas Department of Information Resources, *Adoption of 1 Texas Administrative Code chapter 219, artificial intelligence*, 51 Tex. Reg. (Mar. 13, 2026), Filed Feb. 26, 2026. Effective Mar. 18, 2026. TRD-202600994 and TRD-202600995, 2026. Accessed: Sep. 22, 2026. [Online]. Available: https://www.sos.texas.gov/texreg/archive/March132026/Adopted%20Rules/1.ADMINISTRATION.html

[5] A. Rah and D. W. Johnson, "Beyond AI Productivity: Enterprise AI Systems, Human–AI Responsibility, and the Future of Work," 2026.

[6] Texas Legislature, *Texas government code chapter 2054, information resources*, Texas Constitution and Statutes, 2026. Accessed: Sep. 22, 2026. [Online]. Available: https://statutes.capitol.texas.gov/Docs/GV/htm/GV.2054.htm

[7] Texas Legislature, *Texas government code chapter 403, comptroller of public accounts*, Texas Constitution and Statutes, Secs. 403.021 and 403.071, 2026. Accessed: Sep. 22, 2026. [Online]. Available: https://statutes.capitol.texas.gov/Docs/GV/htm/GV.403.htm

[8] Texas Comptroller of Public Accounts, *Encumbrance reporting and lapsing of appropriations (APS 018), FPP A.019*, Fiscal Management Division, 2026. Accessed: Sep. 22, 2026. [Online]. Available: https://fmx.cpa.texas.gov/fm/pubs/aps/18/aps018_all.php

[9] Texas Comptroller of Public Accounts, *Appropriations, eXpendit, fiscal matters*, Texas Comptroller of Public Accounts, Restates Tex. Const. art. VIII, § 6, and Tex. Gov't Code § 403.071(b), 2026. Accessed: Sep. 22, 2026. [Online]. Available: https://fmx.cpa.texas.gov/fm/pubs/purchase/fiscalmatters/index.php?page=appropriations&section=fiscal

[10] Texas Legislature, *Texas education code §§ 51.9335 and 51.9337, acquisition of goods and services and purchasing authority conditional*, Texas Constitution and Statutes, Section 51.9335(d) provides that, subject to Section 51.9337, Subtitle D, Title 10, Government Code does not apply to an acquisition under that section, except historically underutilized-business and disability-procurement provisions. The institution may, but need not, use Subtitle D. Subsection (e) excludes The University of Texas M. D. Anderson Cancer Center and a public junior college from that section's definition., 2026. Accessed: Nov. 6, 2026. [Online]. Available: https://statutes.capitol.texas.gov/Docs/ED/htm/ED.51.htm

[11] U.S. Government Accountability Office, *Artificial intelligence acquisitions: Agencies should collect and apply lessons learned to improve future procurements*, GAO-26-107859, Washington, DC, Apr. 13, 2026. Accessed: Sep. 22, 2026. [Online]. Available: https://www.gao.gov/assets/gao-26-107859.pdf

[12] J. R. Birge and F. Louveaux, *Introduction to Stochastic Programming*, 2nd ed. New York: Springer, 2011.

[13] R. T. Rockafellar and S. Uryasev, "Optimization of conditional value-at-risk," *The Journal of Risk*, vol. 2, no. 3, pp. 21–41, 2000.

[14] Office of Management and Budget, *Memorandum M-25-22, driving efficient acquisition of artificial intelligence in government*, Executive Office of the President, Rescinds and replaces OMB Memorandum M-24-18, Apr. 2025. Accessed: Sep. 22, 2026. [Online]. Available: https://www.whitehouse.gov/wp-content/uploads/2025/02/M-25-22-Driving-Efficient-Acquisition-of-Artificial-Intelligence-in-Government.pdf

[15] General Services Administration, *Acquisition letter MV-21-06, procurement of cloud computing on a consumption basis under the federal supply schedule program*, GSA Office of Government-wide Policy, with Supplement 1, In effect until rescinded or incorporated. Operative clause: GSAR 552.238-199, Dec. 2021. Accessed: Sep. 23, 2026. [Online]. Available: https://www.gsa.gov/system/files/MV-21-06%20with%20sup%201_0.pdf

[16] Texas Comptroller of Public Accounts, *State of Texas procurement and contract management guide, version 4.0*, Statewide Procurement Division, Publication 96-1809, Sept. 2025, 2025. Accessed: Sep. 23, 2026. [Online]. Available: https://comptroller.texas.gov/purchasing/publications/procurement-contract.php

[17] Texas Department of Information Resources, *1 Texas Administrative Code § 212.1, commodity items definitions*, 1 Tex. Admin. Code § 212.1, amended eff. Feb. 10, 2022, 2022. Accessed: Sep. 23, 2026. [Online]. Available: https://www.law.cornell.edu/regulations/texas/1-Tex-Admin-Code-SS-212-1

[18] P. Mohajerin Esfahani and D. Kuhn, "Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations," *Mathematical Programming*, vol. 171, no. 1–2, pp. 115–166, 2018.

[19] Texas Department of Information Resources, *Buying through DIR*, Texas Department of Information Resources, 2026. Accessed: Sep. 22, 2026. [Online]. Available: https://dir.texas.gov/it-solutions-and-services/buying-through-dir

[20] Texas Department of Information Resources, *Statement of work (SOW) process*, Texas Department of Information Resources, 2026. Accessed: Sep. 22, 2026. [Online]. Available: https://dir.texas.gov/it-solutions-and-services/buying-through-dir/statement-work-sow

[21] Texas Secretary of State, *34 Texas Administrative Code § 20.82, delegated purchases*, 34 Tex. Admin. Code § 20.82, 2017. Accessed: Sep. 22, 2026. [Online]. Available: https://www.law.cornell.edu/regulations/texas/34-Tex-Admin-Code-SS-20-82

[22] Texas Secretary of State, *34 Texas Administrative Code § 20.181, state agency compliance and documentation responsibilities*, 34 Tex. Admin. Code § 20.181, 2017. Accessed: Sep. 22, 2026. [Online]. Available: https://www.law.cornell.edu/regulations/texas/34-Tex-Admin-Code-SS-20-181

[23] Texas Department of Information Resources, *Procurement professional's guide to DIR*, Texas Department of Information Resources, DIR states that the Section 2157.068 solicitation counts apply to state agencies and do not apply to higher education, K–12, local governments, or other voluntary customers., 2026. Accessed: Nov. 6, 2026. [Online]. Available: https://dir.texas.gov/it-solutions-and-services/buying-through-dir/procurement-professionals-guide-dir

[24] Texas Comptroller of Public Accounts, *Capital budget rider appropriations (APS 026), FPP A.012*, Fiscal Management Division, 2026. Accessed: Sep. 22, 2026. [Online]. Available: https://fmx.cpa.texas.gov/fm/pubs/aps/26/a012_all.php

[25] Texas Comptroller of Public Accounts, *Manual of accounts, fiscal 2026, general ledger account 9001, encumbrances*, Texas Comptroller of Public Accounts, 2026. Accessed: Sep. 22, 2026. [Online]. Available: https://fmcpa.cpa.state.tx.us/fiscalmoa/glacct.jsp?num=9001

[26] National Institute of Standards and Technology, *The NIST definition of cloud computing*, U.S. Department of Commerce, Tex. Gov't Code § 2157.007 incorporates the definition as it existed on Jan. 1, 2015, Sep. 2011. Accessed: Sep. 22, 2026. [Online]. Available: https://csrc.nist.gov/publications/detail/sp/800-145/final

[27] Texas Legislature, *Senate bill 475, 87th legislature, regular session, enrolled*, Texas Legislature Online, Originally enacted Tex. Gov't Code § 2054.0593 (cloud risk-and-authorization program), 2021. Accessed: Sep. 22, 2026. [Online]. Available: https://capitol.texas.gov/tlodocs/87R/billtext/html/SB00475F.HTM

[28] Texas Legislature, *House bill 150, 89th legislature, regular session, enrolled*, Texas Legislature Online, Section 16 transfers, redesignates, and amends former Gov't Code § 2054.0593 as § 2063.408. Section 50(c) continues the department's pre-Act Chapters 2054 and 2059 duties until the memorandum date. Section 50(d) requires the Chapter 2063 cybersecurity functions to transfer not later than Dec. 31, 2026., 2025. Accessed: Nov. 8, 2026. [Online]. Available: https://capitol.texas.gov/tlodocs/89R/billtext/html/HB00150F.htm

[29] Texas Legislature, *Tex. gov't code § 2063.408, cloud computing state risk and authorization management program*, Texas Constitution and Statutes, Transferred, redesignated, and amended from former § 2054.0593 by Acts 2025, 89th Leg., R.S., Ch. 331 (H.B. 150), Sec. 16, eff. Sept. 1, 2025. Subsection (b) assigns establishment of the program to Texas Cyber Command., 2025. Accessed: Nov. 8, 2026. [Online]. Available: https://statutes.capitol.texas.gov/Docs/GV/htm/GV.2063.htm

[30] Texas Secretary of State, *Texas cyber command rule transfer: 1 TAC chapter 202 to chapter 477*, Texas Register transfer notice, Effective Aug. 1, 2026. Information Security Standards moved from Title 1, Part 10 (DIR) Chapter 202 to Title 1, Part 19 (Texas Cyber Command) Chapter 477, 2026. Accessed: Nov. 7, 2026. [Online]. Available: https://www.sos.texas.gov/texreg/transfers/cyber062626.html

[31] Texas Cyber Command, *1 Texas Administrative Code § 477.27, Texas Risk and Authorization Management Program for state agencies*, 1 Tex. Admin. Code § 477.27, Formerly 1 Tex. Admin. Code § 202.27 (amended eff. Nov. 16, 2023); transferred to § 477.27 eff. Aug. 1, 2026 (TRD-202602469). Cited text is that of former § 202.27, 2026. Accessed: Nov. 7, 2026. [Online]. Available: https://www.sos.texas.gov/texreg/transfers/cyber062626.html

[32] Texas Department of Information Resources, *Texas risk and authorization management program manual, version 4.0*, Texas Department of Information Resources, Cites Tex. Gov't Code § 2063.408 and 1 Tex. Admin. Code ch. 202 (transferred to ch. 477 eff. Aug. 1, 2026), Feb. 2026. Accessed: Nov. 8, 2026. [Online]. Available: https://dir.texas.gov/sites/default/files/2026-02/TX-RAMP%20Program%20Manual%204.0.pdf

[33] Texas Cyber Command, *1 Texas Administrative Code § 477.5, Texas Risk and Authorization Management Program responsibilities and minimum standards*, 1 Tex. Admin. Code § 477.5, Formerly 1 Tex. Admin. Code § 202.5 (adopted eff. Nov. 16, 2023); transferred to § 477.5 eff. Aug. 1, 2026. Cited text is that of former § 202.5, 2026. Accessed: Nov. 7, 2026. [Online]. Available: https://www.sos.texas.gov/texreg/transfers/202602468-1.pdf

[34] Texas Cyber Command, *1 Texas Administrative Code § 477.77, Texas Risk and Authorization Management Program for institutions of higher education*, 1 Tex. Admin. Code § 477.77, Formerly 1 Tex. Admin. Code § 202.77 (amended eff. Nov. 16, 2023); transferred to § 477.77 eff. Aug. 1, 2026. Cited text is that of former § 202.77, 2026. Accessed: Nov. 7, 2026. [Online]. Available: https://www.sos.texas.gov/texreg/transfers/202602468-1.pdf

[35] National Institute of Standards and Technology, *Artificial intelligence risk management framework (AI RMF 1.0)*, U.S. Department of Commerce, Incorporated by reference in Tex. Gov't Code §§ 2054.702–2054.703, Jan. 2023. Accessed: Sep. 22, 2026. [Online]. Available: https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf

[36] Texas Department of Information Resources, *AI templates and resources*, Texas Department of Information Resources, 2026. Accessed: Sep. 22, 2026. [Online]. Available: https://dir.texas.gov/ai-and-innovation/ai-templates-and-resources

[37] United States Congress, *31 U.S.C. § 1341, limitations on expending and obligating amounts*, United States Code, 2026. Accessed: Sep. 22, 2026. [Online]. Available: https://www.law.cornell.edu/uscode/text/31/1341

[38] United States Congress, *31 U.S.C. § 1502, balances available*, United States Code, 2026. Accessed: Sep. 22, 2026. [Online]. Available: https://uscode.house.gov/view.xhtml?req=title:31+section:1502+edition:prelim

[39] General Services Administration, *GSA expands OneGov AI offerings with discounted, consumption-based access to OpenAI's ChatGPT*, GSA Newsroom, Sep. 2026. Accessed: Sep. 22, 2026. [Online]. Available: https://www.gsa.gov/about-gsa/newsroom/news-releases/gsa-expands-onegov-ai-offerings-with-discounted-openais-chatgpt-09102026

[40] Texas Department of Information Resources, *Bulk purchase agreement BP2026-0015, posted contract terms on authorized use of customer data and AI governance*, Texas Department of Information Resources, DIR-posted contract document, 2026. Accessed: Sep. 22, 2026. [Online]. Available: https://dir.texas.gov/sites/default/files/2026-08/Carahsoft%20Bulk%20Purchase%20Agreement%20BP2026-0015..pdf

[41] W.-K. Mak, D. P. Morton, and R. K. Wood, "Monte Carlo bounding techniques for determining solution quality in stochastic programs," *Operations Research Letters*, vol. 24, no. 1–2, pp. 47–56, 1999.

[42] Texas Department of Information Resources, *1 Texas Administrative Code § 212.41, statement of work*, 1 Tex. Admin. Code § 212.41, amended eff. Feb. 10, 2022, 2022. Accessed: Sep. 23, 2026. [Online]. Available: https://www.law.cornell.edu/regulations/texas/1-Tex-Admin-Code-SS-212-41

[43] Texas Comptroller of Public Accounts, *34 Texas Administrative Code § 20.25, definitions*, 34 Tex. Admin. Code § 20.25, amended eff. Mar. 14, 2021, See subsec. (b)(13), 2021. Accessed: Sep. 23, 2026. [Online]. Available: https://www.law.cornell.edu/regulations/texas/34-Tex-Admin-Code-SS-20-25

[44] Texas Department of Information Resources, *Texas public cloud services*, DIR Shared Technology Services, Data Center Services, 2026. Accessed: Sep. 23, 2026. [Online]. Available: https://dir.texas.gov/shared-technology-services/data-center-services/texas-public-cloud-services

[45] Texas Department of Information Resources, *Statement of work application user guide*, Texas Department of Information Resources, Oct. 2022. Accessed: Sep. 23, 2026. [Online]. Available: https://dir.texas.gov/sites/default/files/2022-10/DIR%20SOW%20Application%20User%20Guide_0.pdf

[46] A. Rah, *Agentic and multi-agent systems: A systematic review of tool use, benchmarks, and governance*, Research Square preprint, Posted May 29, 2026, 2026.

[47] A. Rah and D. W. Johnson, *Enterprise AI at the decision point: A U.S.-focused systematic review of productivity, ROI, workforce change, and governance*, Research Square preprint, 2026.

[48] A. Rah and S. Hahues, *Retrieval-augmented generation (RAG), generative AI, and agentic AI governance: An integrated enterprise governance prioritization architecture*, Research Square preprint, Posted June 15, 2026, 2026.

[49] A. Rah, *A digital-twin multi-agent framework for enterprise cognitive networking: ECNetBench and leakage-safe multi-seed evaluation*, SSRN preprint, 2026.

[50] P. Artzner, F. Delbaen, J.-M. Eber, and D. Heath, "Coherent measures of risk," *Mathematical Finance*, vol. 9, no. 3, pp. 203–228, 1999.

[51] P. Brucker, A. Drexl, R. Möhring, K. Neumann, and E. Pesch, "Resource-constrained project scheduling: Notation, classification, models, and methods," *European Journal of Operational Research*, vol. 112, no. 1, pp. 3–41, 1999.

[52] S. Chaisiri, B.-S. Lee, and D. Niyato, "Optimization of resource provisioning cost in cloud computing," *IEEE Transactions on Services Computing*, vol. 5, no. 2, pp. 164–177, 2012.

[53] A. Rah, *Enterprise generative AI: A systematic review of security risks and governance*, Research Square preprint, 2026.

[54] A. Rah, *Cybersecurity risk controls in IT organizations*, ResearchGate preprint, Posted May 2026, 2026. [Online]. Available: https://www.researchgate.net/publication/404348824_Cybersecurity_Risk_Controls_in_IT_Organizations

[55] A. Rah, *Auditing operational data records to improve trust and risk assessments*, ResearchGate preprint, Posted May 2026, 2026. [Online]. Available: https://www.researchgate.net/publication/404347907_Auditing_Operational_Data_Records_to_Improve_Trust_and_Risk_Assessments