Work-Order Authorization in Maintenance Service Networks: Optimal Mechanisms with Supplier Inducements

Chenxi Xu
Mitchell E. Daniels, Jr. School of Business, Purdue University, xu2301@purdue.edu

Ziwei Zhu
Samuel Curtis Johnson Graduate School of Management, Cornell University, zz575@cornell.edu

Abstract. In decentralized maintenance service networks, work-order authorization—the decision of whether to approve a requested repair or replacement task—often depends on information held by different parties. Field evaluators observe diagnostic evidence and operating context that determine the work-order class, whereas suppliers privately know fulfillment costs. This separation creates two adverse-selection problems: suppliers may overstate costs, and field evaluators may classify low-impact or weakly supported work orders as high impact to increase authorization likelihood. A further cross-agent incentive problem arises because suppliers may benefit from high-impact classifications and may therefore offer report-contingent side payments or other inducements that make classification inflation attractive. We develop a mechanism-design model in which the administrator jointly chooses authorization probabilities, documentation requirements, and supplier payments. The optimal mechanisms balance the value of screening low-impact work orders against the cost of sustaining truthful classification. When low-impact work orders have low value, class-based screening is valuable and may be supported by documentation or class-contingent supplier compensation. As low-impact value increases, the optimal mechanism softens the class distinction through probabilistic authorization; when low-impact work orders are sufficiently valuable, pooling becomes optimal. We show that class-contingent expected net transfers can internalize supplier-financed inducements and implement the optimal mechanism. The results imply that documentation is an incentive instrument whose value depends on authorization and payment design, and that administrators should design authorization, documentation, and supplier compensation jointly.

Key words: mechanism design, service operations, maintenance contracts, supplier inducements, work-order authorization

History: The current version was updated on June 25, 2026.

1. Introduction

Repair and maintenance services are often organized through decentralized service-contract networks. In these networks, a work order refers to a service request, claim, or maintenance task submitted for authorization. Administrators decide whether requested repairs or replacements should be authorized, field evaluators document the operational basis for authorization, and suppliers complete or support approved tasks. We study this task-level decision, which we call *work-order authorization*: the joint design of authorization rules, documentation requirements, and supplier compensation.

A leading application is maintenance contracting for building systems and facility equipment in schools, hospitals, universities, airports, transit authorities, public agencies, and large enterprise facilities. These organizations often rely on outside suppliers or contractors to perform standardized repair or replacement tasks, while local facility managers, field certifiers, technicians, or service representatives observe asset conditions and classify work orders according to operational impact and evidentiary support. Consider, for example, a standardized task such as replacing a specified HVAC component. The physical task may be similar across locations, but its authorization value can differ sharply across operating contexts. Authorizing the replacement may have relatively low value when the problem affects a noncritical area, but much higher value when it affects a hospital room, a school's heating system, a critical equipment room, or another high-value operating environment. Thus, the work-order class captures the operational importance and evidentiary support for authorization, rather than the physical complexity of the task itself.

This separation between diagnosis, authorization, and fulfillment creates two adverse-selection problems and one cross-agent incentive problem. The field evaluator privately observes diagnostic and contextual information about the work-order class, while the supplier privately observes the cost of completing the task. The supplier may overstate its cost to increase compensation, and the evaluator may inflate a low-impact or weakly supported work order to increase the probability of authorization. Moreover, because a high-impact classification can also benefit the supplier, the supplier may reimburse substantiation effort, provide documentation support, or offer other report-contingent inducements that make classification inflation more attractive to the evaluator.

Documentation is a natural governance tool in this environment, but its ability to deter inflation depends on the incentives created by authorization and supplier compensation. Requiring diagnostic reports, photographs, maintenance records, teardown notes, or other supporting evidence can make inflated classifications more costly for the field evaluator. Supplier-financed payment, however, limits documentation as a stand-alone governance tool: if a high-impact classification raises the probability of authorization or the supplier's expected net compensation, the supplier may offset the evaluator's incremental documentation burden. A documentation policy that appears effective when the evaluator is considered in isolation may therefore fail once supplier-financed inducements are taken into account. The administrator must design documentation jointly with authorization and compensation.

This raises the central question of the paper: *How should an administrator jointly design work-order authorization, documentation requirements, and supplier payments when field evaluators privately observe* *diagnostic information that determines work-order classes, suppliers privately observe fulfillment costs, and suppliers may influence evaluator reports?*

We develop a mechanism-design model in which the administrator acts as the principal and interacts with two agents: a field evaluator and a supplier. The evaluator privately observes whether the proposed standardized task has low or high authorization value, while the supplier privately observes fulfillment cost. The administrator commits to a mechanism specifying three instruments: the probability of authorizing the work order, the expected net transfer to the supplier, and the probability of requiring supporting documentation. After observing the mechanism, the supplier reports its cost and may offer report-contingent inducements to influence the evaluator's classification report. The evaluator then reports the work-order class, after which authorization, transfer, and documentation rules are implemented.

Supplier inducements make the design problem nonstandard. In a standard direct mechanism, the administrator would use transfers to elicit supplier costs and documentation to discipline evaluator reports. Here, however, supplier compensation also affects the supplier's incentive to finance inducements that alter another agent's report. The administrator must therefore account for both direct reporting incentives and cross-agent influence.

We proceed in two steps. First, we analyze an implementable no-side-payment benchmark that induces truthful classification without supplier-financed inducements in equilibrium. Second, we study an auxiliary direct problem that treats intended supplier-financed transfers to the evaluator as decision variables. We show that, under class-contingent expected net transfers, the auxiliary mechanism can be implemented in the original contracting environment: by changing the supplier's relative payoff from inducing different reported classes, the administrator can make it weakly optimal for the supplier to offer the intended inducements.

The optimal mechanisms can be organized into three broad authorization patterns. When low-impact work orders generate sufficiently low value, screening the work-order class is valuable: both classes are authorized at low fulfillment costs, only high-impact work orders are authorized at intermediate costs, and neither class is authorized at high costs. This separation may be supported either by documentation or by class-contingent supplier compensation, depending on the relative costs of documentation and inducements. When low-impact work orders become more valuable, deterministic class-contingent authorization becomes more costly to sustain; the optimal mechanism then relies on probabilistic authorization, combined with documentation requirements and, in some regimes, supplier-financed payments. When low-impact work orders are sufficiently valuable, screening the work-order class is no longer worthwhile, and the administrator adopts a pooling mechanism that authorizes both classes below a fulfillment-cost threshold and denies both above it.

This paper makes three contributions. First, it contributes to the operations management literature on service contracts and maintenance operations by formulating work-order authorization as a decentralized service-network governance problem. Prior work on warranty, service-contract, and maintenance operations has emphasized coverage design, pricing, repair-or-replacement decisions, service quality, outsourcing, and customer adoption. We complement this literature by studying task-level authorization decisions in decentralized maintenance networks, where local evaluators classify work orders and separate suppliers privately observe fulfillment costs.

Second, the paper contributes to research on costly documentation and hassle-cost screening by showing that documentation is an incentive instrument rather than merely an administrative requirement. In our model, documentation creates a differential burden between truthful and inflated classifications. Its value depends not only on documentation and review costs, but also on authorization value, supplier cost, and the supplier-payment system. This perspective explains why stricter documentation is not always optimal and why documentation policy must be designed jointly with authorization and compensation.

Third, the paper identifies a side-contracting friction in decentralized service networks: suppliers may benefit from high-impact classifications and may therefore use side payments or other inducements to influence evaluator reports. This turns work-order authorization into a multi-agent governance problem. We show that class-contingent expected net transfers can internalize supplier-financed inducements and implement the optimal mechanism, thereby linking side-contracting concerns to the operational design of authorization, documentation, and supplier compensation.

The remainder of the paper is organized as follows. Section 2 reviews the related literature. Section 3 introduces the three-party work-order authorization model and formulates the administrator's optimization problem. Section 4 studies an implementable no-side-payment benchmark. Section 5 characterizes the optimal mechanism with supplier-financed inducements and shows how the auxiliary direct mechanism can be implemented through class-contingent expected net transfers. Section 6 examines how documentation and review costs affect the optimal policy and participants' payoffs. Section 7 discusses managerial implications, and Section 8 concludes.

2. Literature Review

This paper relates to three streams of research. First, it connects to the operations management literature on service contracts, maintenance operations, and after-sales service supply chains. Second, it relates to research on auditing, verification, and documentation in contract design and operational governance. Third, it connects to the literature on multi-agent mechanism design, collusion, and side contracting.

2.1. Maintenance Service Contracts and Work-Order Authorization

This paper relates to the operations management literature on maintenance service contracts, performance-based contracting, and after-sales service supply chains. Within this broad area, performance-based service-contracting research studies performance-based incentives in dynamic operational systems, compensation tied to after-sales service outcomes, and the effect of performance-based contracts on product reliability (Plambeck and Zenios 2000, Kim et al. 2007, Guajardo et al. 2012). Related maintenance-service models analyze contracts for mission-critical recovery, outsourced repair and restoration, machine repair and maintenance, and medical-equipment maintenance services (Kim et al. 2010, Jain et al. 2013, Tian et al. 2021, Chan et al. 2019). A related stream on warranties and service plans examines warranty policy and extended service contracts, interactions between retailer service plans and manufacturer warranties, and reliability signaling through after-sales service contracts (Padmanabhan and Rao 1993, Jiang and Zhang 2011, Bakshi et al. 2015). More broadly, service-operations research has examined how information and inference shape decisions in service systems under uncertainty; for example, Yang et al. (2023) study customer joining behavior when service capacity is uncertain and queue length is observable.

We share with this stream of literature an interest in how service-contract design shapes maintenance and repair outcomes, but our focus is different: we study work-order authorization before fulfillment. Rather than tying compensation to realized service performance, choosing warranty coverage, or designing relationship-level service contracts, we examine how an administrator should authorize individual repair or replacement tasks when field diagnosis and task fulfillment are separated across parties. In our model, a field evaluator privately observes the operational value and evidentiary support for the work order, while a supplier privately observes the cost of completing the authorized task. This separation creates a classification-and-approval problem that is largely absent from existing service-contract models: the reported work-order class affects the authorization decision and may also affect the supplier's expected payoff, giving the supplier an incentive to induce classification inflation. We therefore complement the maintenance-contract and warranty/service-plan literatures by studying how authorization probabilities, documentation requirements, and supplier payments interact within decentralized service networks.

2.2. Costly Documentation and Hassle-Cost Screening

Our paper relates to research on costly controls for private-information problems. A useful economic benchmark is the costly state verification literature, beginning with Townsend (1979) and applied in the debt-contract setting by Gale and Hellwig (1985), in which privately observed outcomes can be verified only by incurring a cost. Our model departs from this framework in a key respect: documentation is not modeled as a state-revelation or detection technology, but as a differential documentation burden. Supporting a truthful report and supporting an inflated high-impact classification impose different costs on the field evaluator, so documentation disciplines classification through self-selection rather than through information acquisition. This distinction also separates our mechanism from models of optimal auditing, which study how audit probabilities and transfers discipline reports when verification is costly (Border and Sobel 1987, Mookherjee and Png 1989), and from work on partially verifiable information, which studies how type-dependent message spaces or evidentiary restrictions alter implementability (Green and Laffont 1986).

At the mechanism level, a closer analogue is the use of costly procedural frictions, often referred to as ordeal mechanisms, to induce self-selection. In public-economics models of targeting, restrictions and work requirements can improve targeting by making claimants bear nonpecuniary participation costs (Nichols and Zeckhauser 1982, Besley and Coate 1992); related empirical work shows that modest application, travel, and waiting costs can improve self-targeting, although making ordeals more onerous need not further improve targeting (Alatas et al. 2016). Similar logic appears in operations and marketing studies of return-policy design, where partial refunds, restocking fees, return hassle, and hassle-free return policies affect purchase and return decisions without independently verifying consumers' underlying valuations or product fit (Su 2009, Shulman et al. 2011, Hsiao and Chen 2014). In service settings, Dukes and Zhu (2019) show that tiered customer-service structures can impose escalation hassles that screen less severe or illegitimate claims. Across these settings, costly procedural frictions screen behavior by changing the cost of taking an action, rather than by independently verifying the underlying state.

An adjacent operations management stream studies inspection, appraisal, certification, deferred payment, and contingency payments as operational governance tools in supplier-quality, product-integrity, and responsible-sourcing settings (Baiman et al. 2000, Hwang et al. 2006, Babich and Tang 2012, Rui and Lai 2015, Chen and Lee 2017). We build on this stream's insight that costly operational controls and payment design interact, but our mechanism differs in both the object of control and the source of offsetting incentives. Rather than assessing supplier quality, product adulteration, or responsibility risk, documentation in our model is attached to an informed field evaluator's work-order classification. This distinction matters because the classification is valuable to another contracting party: the supplier may benefit when a work order is classified as high impact and may therefore reimburse documentation effort, assist with substantiation, or provide other inducements. This third-party offsetting channel is where our hassle-cost screening problem connects to side contracting: the party that benefits from a report can endogenously finance the removal of the screening friction. As a result, documentation design cannot be separated from supplier-payment design, because supplier compensation is endogenous and can neutralize the screening friction created by documentation.

2.3. Side Contracting and Supplier Inducements

Our paper also relates to the literature on multi-agent mechanism design, collusion, and side contracting. A benchmark mechanism-design formulation imposes individual incentive-compatibility constraints to rule out unilateral misreports within the principal's mechanism. In multi-agent environments, however, agents may communicate, coordinate reports, or form side agreements outside the principal's mechanism. Classic work on collusion in organizations and hierarchical agency shows how such side contracting can distort information transmission and constrain incentive design (Tirole 1986, Kofman and Lawarrée 1993). More general models of collusion under asymmetric information and correlated information show that collusive communication and side contracting can alter implementability and optimal incentive schemes (Laffont and Martimort 1997, 2000). This literature highlights that a principal may need mechanisms that are robust not only to individual deviations, but also to cross-agent influence.

Our setting shares this concern, but the direction and object of influence are different. The supplier is not colluding with a monitor or auditor to conceal supplier noncompliance or manipulate a compliance signal. Instead, it offers report-contingent inducements to affect another agent's work-order classification because that classification affects authorization likelihood and the supplier's expected net compensation. The design problem is therefore not simply to rule out side contracting, but to design supplier transfers so that supplier-financed inducements support truthful classification rather than classification inflation.

A related stream of operations management research uses mechanism-design and game-theoretic approaches to examine how hidden information shapes operational governance. This includes work on environmental disclosure and audit evasion (Wang et al. 2016, 2024), as well as procurement under supplier cost private information (Hu and Qi 2018, Davis et al. 2022). More closely related to the inducement channel, Fan et al. (2021) study a procurement service provider that may receive hidden rebates and inflate quotes, and Chen et al. (2020) study responsible sourcing under supplier-auditor collusion. These papers show how hidden information, hidden transfers, intermediary incentives, and audit-related collusion can undermine operational governance.

We complement these studies by analyzing a different cross-agent inducement problem in work-order authorization. Unlike settings in which the distortion concerns the supplier's own compliance report, a procurement intermediary's price quote, or an auditor's compliance signal, the supplier in our model can finance the field evaluator's classification because that classification affects authorization likelihood and the supplier's expected payoff. The administrator must therefore design authorization, documentation, and supplier compensation jointly: supplier payments must both elicit fulfillment-cost information and govern the supplier's incentive to influence the field evaluator's classification report through side payments or other inducements. This structure motivates our no-side-payment benchmark and auxiliary supplier-financed-transfer problem, and leads to the implementation result that class-contingent expected net transfers can internalize the inducement channel.

3. Model

We study a sequential game in which an administrator, acting as the principal, interacts with two agents: a supplier and a field evaluator. We present the model in the context of work-order authorization in decentralized service-contract networks, with maintenance networks as the leading application. We use the term "work order" broadly to include service requests or claims submitted for authorization in warranty, protection-plan, and maintenance-contract settings. In these settings, the administrator authorizes standardized repair or replacement tasks, the field evaluator observes diagnostic evidence and operating context that determine the work-order class, and the supplier privately knows the cost of completing the task. Examples include maintenance contracts for building systems in schools, hospitals, universities, airports, transit authorities, public agencies, and large enterprise facilities.[1] The model therefore applies to programs in which diagnosis and classification are at least partially separated from fulfillment and the core operational decision is whether—and under what documentation and payment rules—to authorize a standardized repair or replacement task. The work-order classes in the model capture the authorization value, operational impact, and evidentiary support for a proposed standardized task, rather than different physical tasks. Because evaluators handling different work orders do not strategically interact, we study a representative work order and the field evaluator assigned to it. All three parties are risk-neutral.

The supplier privately knows a unit fulfillment cost c for a standardized physical task within a given service category, such as replacing a specified HVAC fan module, appliance component, pump controller, sensor, or facility-equipment part. Conditional on authorization, the supplier carries out the task and incurs cost c. The field evaluator, in contrast, privately observes diagnostic and contextual information about the customer's problem or the operating condition of the asset. We summarize this information by a binary work-order class $i \in \{L, H\}$, which captures

[1] Similar incentive problems may arise in home-warranty plans, stand-alone HVAC or appliance protection plans, vehicle service contracts, electronics protection plans, public-sector maintenance procurement, and other extended-service or facilities-management programs. In these applications, the supplier may be an HVAC or appliance parts vendor, a replacement-unit supplier, a specialized maintenance contractor, or a centralized fulfillment network that supplies the components required for fulfillment or completes authorized fulfillment tasks. The field evaluator may be a facility manager, field certifier, authorized technician, or local service representative who observes the operational condition of the asset, documents the evidentiary basis for authorization, and submits the work-order classification on behalf of customers or operating units.

the authorization value, operational impact, and evidentiary support associated with the task. A class-L work order is a low-impact or weakly justified case; for example, the component may be noisy, intermittently problematic, or only weakly supported for replacement. A class-H work order is a high-impact or strongly justified case for the same task; for example, the component may affect a hospital room, a school's heating system, a critical equipment room, or another high-value operating environment. Let θ_i denote the expected value from fulfilling a class-i work order, with $\theta_H > \theta_L$. Because the field evaluator acts as the local representative of the customer or operating unit in the authorization process, we model it as deriving value θ_i from authorization of a class-i work order. All parties share the belief that (1) the fulfillment cost c is drawn from a distribution with support $[c_L, c_H]$, cumulative distribution function F, and density f; and (2) independently of c, the work-order class is i with probability ρ_i for $i \in \{L, H\}$, where $\rho_L + \rho_H = 1$.[2] Throughout the paper, a "class-i field evaluator" refers to a field evaluator handling a work order in class i.

Two distinct adverse-selection problems arise in our setting. First, the supplier has an incentive to overstate its fulfillment cost in order to extract a larger transfer from the administrator. Second, the field evaluator has an incentive to overstate the work-order class to increase the likelihood of authorization. We model this classification distortion as one-sided: a class-L field evaluator may report the work order as class H, whereas class-H work orders are assumed not to be strategically downgraded to class L. This assumption fits service networks in which high-impact status is supported by customer-visible or operationally salient evidence—such as critical-location designations, service-level requirements, shutdown records, safety concerns, or operating-unit escalation. In such settings, downgrading a class-H work order would forgo clear customer or operating-unit value and could be challenged when the supporting evidence is visible to the administrator or other stakeholders. This assumption defines the intended scope of the model: if class-H work orders could also be strategically downgraded, additional incentive constraints would be needed. To discipline class-L-to-H inflation, the administrator may require the field evaluator to submit supporting documentation, such as diagnostic reports, photographs, maintenance records, teardown notes, evidence of operational impact, or technician assessments. If the field evaluator reports the work-order class truthfully, the documentation cost is c_d; if it inflates a class-L work order as class H, the cost is $c'_d \geq c_d$, capturing the additional effort, scrutiny, and reputational or legal risk involved in substantiating an inflated classification. We assume that documentation requests are effectively mandatory under the administrator's network rules: refusal

[2] The independence assumption reflects our focus on a standardized fulfillment task for which physical fulfillment cost is governed by supplier-side cost heterogeneity, while the work-order class captures authorization value, operational impact, and evidentiary support. If high-impact environments systematically changed the physical cost of fulfillment, the model could be extended to allow class-dependent cost distributions.

can delay or prevent reimbursement, damage the field evaluator's standing, or jeopardize continued participation in the approved network. The administrator incurs review cost c_p for submitted documentation.

Beyond the adverse-selection issues described above, a cross-agent side-contracting problem also arises: the supplier may offer report-contingent inducements to field evaluators to encourage class-L work orders to be reported as class H, thereby increasing the probability of authorization or the supplier's expected net compensation under the contract. Such inducements may take the form of reimbursing documentation or other substantiation effort, providing informal technical support, expediting future service, absorbing customer or operating-unit costs, or offering relationship benefits. A field evaluator may accept such inducements because class-L work orders still generate positive expected value from fulfillment, even when fulfillment is not socially efficient $(\theta_L < c)$, and the inducement helps offset the cost of supporting the inflated classification.

The game begins with the administrator publicly committing to a mechanism

$$\mathcal{M} \equiv \{q_i(\cdot), m_i(\cdot), \tau_i(\cdot)\}_{i\in\{L,H\}},$$

where i denotes the reported work-order class. For each reported fulfillment cost $\hat{c} \in [c_L, c_H]$ and reported work-order class i, the mechanism specifies the probability of authorizing the work order, $q_i(\hat{c})$; the expected net transfer to the supplier, $m_i(\hat{c})$; and the probability of requiring the field evaluator to submit supporting documentation, $\tau_i(\hat{c})$. The transfer $m_i(\hat{c})$ represents the supplier's expected net compensation under the administrator–supplier contract.[3] Because the supplier accepts or rejects the contract before the work-order class is reported, its participation constraint is imposed on the expected payoff conditional on its fulfillment cost, rather than separately for each reported work-order class. If the supplier declines the offer, the game ends immediately, and all parties receive zero payoff. If it accepts, the supplier joins the network, reports a fulfillment cost $\hat{c}$, and may offer report-contingent inducements to influence the field evaluator's class report. The field evaluator then observes the work-order class and submits a report, after which documentation requests, authorization decisions, and transfers are implemented according to the mechanism. Figure 1 summarizes the sequence of moves and information flows.

We analyze the game by backward induction. Fix a supplier cost report $\hat{c}$ and consider a class-L field evaluator. The evaluator compares the payoff from truthful reporting with that from inflating the work order to H. We first define these reporting payoffs before any supplier-financed inducement is added. If the evaluator reports truthfully, the work order is authorized with probability

[3] The net transfer may include fulfillment reimbursement, work-order-handling compensation, service-level payments, chargebacks, rebates, retainage, award or incentive payments, and end-of-period reconciliation. Thus, $m_i(\hat{c})$ is interpreted as an expected net payment conditional on the submitted reports, rather than as a stand-alone physical unit price.

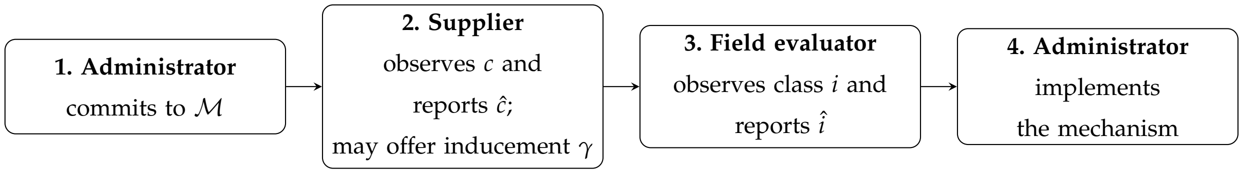

Figure 1 Timeline and information flows in the administrator–supplier–field-evaluator game. Participation decisions are omitted from the timeline and are captured by the individual rationality constraints.

$q_L(\hat{c})$, yielding expected value $\theta_L q_L(\hat{c})$. Documentation is required with probability $\tau_L(\hat{c})$, so the expected documentation cost is $c_d \tau_L(\hat{c})$. Thus, the truthful-report payoff is

$$\theta_L q_L(\hat{c}) - c_d \tau_L(\hat{c}).$$

If instead the evaluator inflates the work order and reports class H, its expected payoff is

$$\theta_L q_H(\hat{c}) - c_d' \tau_H(\hat{c}),$$

where $c_d' \geq c_d$ captures the higher cost of supporting an inflated classification. The authorization value remains θ_L under an inflated report because the underlying work order is still a low-impact or weakly justified class-L case; only the reported class changes, and hence the applicable authorization and documentation probabilities may differ.

From the supplier's perspective, suppose its true fulfillment cost is c and it reports $\hat{c}$ to the administrator. Consider first the case in which the supplier induces a class-L field evaluator to report truthfully. The least costly way to do so is to offer a report-contingent inducement only after an L report, in an amount equal to the positive part of the evaluator's payoff gain from inflating the work order rather than reporting truthfully. Thus, the required inducement is

$$\tilde{\gamma}_L(\hat{c}) \equiv \max\left\{\left[\theta_L\, q_H(\hat{c}) - c_d'\, \tau_H(\hat{c})\right] - \left[\theta_L\, q_L(\hat{c}) - c_d\, \tau_L(\hat{c})\right], 0\right\}. \tag{1}$$

Given this inducement, the supplier's expected payoff is

$$\tilde{u}_L(c, \hat{c}) \equiv \rho_L\, \left[m_L(\hat{c}) - c\, q_L(\hat{c}) - \tilde{\gamma}_L(\hat{c})\right] + \rho_H\, \left[m_H(\hat{c}) - c\, q_H(\hat{c})\right]. \tag{2}$$

Analogously, if the supplier instead seeks to induce an H report from a class-L field evaluator, the minimum H-contingent inducement is the positive part of the opposite payoff difference:

$$\tilde{\gamma}_H(\hat{c}) \equiv \max\left\{\left[\theta_L\, q_L(\hat{c}) - c_d\, \tau_L(\hat{c})\right] - \left[\theta_L\, q_H(\hat{c}) - c_d'\, \tau_H(\hat{c})\right], 0\right\}. \tag{3}$$

Because the supplier cannot observe the field evaluator's private diagnostic and contextual information, any positive H-contingent inducement must be paid whenever an H report is submitted, including when the work order is truly class H. Given this inducement, the supplier's expected payoff is

$$\tilde{u}_H(c,\hat{c}) \equiv m_H(\hat{c}) - c\, q_H(\hat{c}) - \tilde{\gamma}_H(\hat{c}). \tag{4}$$

The supplier compares the expected payoff from inducing truthful L-reporting with the expected payoff from inducing an H report, and chooses the report-contingent inducement that yields the larger expected payoff. We denote this optimized supplier payoff by $\hat{u}_s(c,\hat{c})$, i.e.,

$$\hat{u}_s(c,\hat{c}) \equiv \max\left\{\tilde{u}_L(c,\hat{c}), \tilde{u}_H(c,\hat{c})\right\}. \tag{5}$$

The supplier's cost-reporting strategy is a map

$$\phi_s : [c_L, c_H] \to [c_L, c_H], \tag{6}$$

so that a supplier with true fulfillment cost c reports $\phi_s(c)$. In equilibrium, this strategy maximizes the optimized supplier payoff $\hat{u}_s$ for every true cost c:

$$u_s(c) \equiv \hat{u}_s(c, \phi_s(c)) \geq \hat{u}_s(c,\hat{c}), \quad \forall (c,\hat{c}) \in [c_L, c_H]^2. \tag{7}$$

We normalize the supplier's outside option to zero and impose the following individual rationality constraint to ensure its participation:

$$u_s(c) \geq 0, \quad \forall c \in [c_L, c_H]. \tag{8}$$

The analysis above characterizes the supplier's optimal cost-reporting and report-contingent inducement choices for any mechanism committed by the administrator. We adopt the following tie-breaking convention. If, after accounting for the relevant minimal inducement, the field evaluator is indifferent between reporting L and reporting H, the tie is broken in favor of the supplier-preferred report. If the supplier is indifferent between inducing truthful L-reporting and inducing an H report, the tie is broken in favor of truthful L-reporting. Because a supplier with true fulfillment cost c reports $\phi_s(c)$ in equilibrium, the resulting inducement schedules can be written as functions of c. Given the committed mechanism, define

$$\gamma_L(c) \equiv \tilde{\gamma}_L(\phi_s(c)) \cdot \mathbb{1}_{\{\tilde{u}_L(c,\phi_s(c)) \geq \tilde{u}_H(c,\phi_s(c))\}}, \text{ and} \tag{9}$$

$$\gamma_H(c) \equiv \tilde{\gamma}_H(\phi_s(c)) \cdot \mathbb{1}_{\{\tilde{u}_L(c,\phi_s(c)) < \tilde{u}_H(c,\phi_s(c))\}}. \tag{10}$$

Equations (9)–(10) characterize the equilibrium inducement schedules: only the inducement corresponding to the supplier-preferred reporting outcome can be positive, and it equals the minimal amount required to implement that report. Therefore, the induced report of a class-L field evaluator is

$$\phi_e(c) = \begin{cases} L, & \text{if } \tilde{u}_L(c, \phi_s(c)) \geq \tilde{u}_H(c, \phi_s(c)), \\ H, & \text{if } \tilde{u}_L(c, \phi_s(c)) < \tilde{u}_H(c, \phi_s(c)). \end{cases} \tag{11}$$

The weak inequality in the first branch reflects the tie-breaking convention in favor of truthful reporting when the supplier is indifferent.

We next impose the field-evaluator participation constraints. For any supplier type $c \in [c_L, c_H]$, after the supplier reports $\phi_s(c)$, these constraints are

$$\begin{aligned} u_e(L;c) &\equiv \max\left\{\theta_L\, q_L(\phi_s(c)) - c_d\, \tau_L(\phi_s(c)) + \gamma_L(c), \theta_L\, q_H(\phi_s(c)) - c_d'\, \tau_H(\phi_s(c)) + \gamma_H(c)\right\} \\ &= \max\left\{\theta_L\, q_L(\phi_s(c)) - c_d\, \tau_L(\phi_s(c)), \theta_L\, q_H(\phi_s(c)) - c_d'\, \tau_H(\phi_s(c))\right\} \geq 0, \end{aligned} \tag{12}$$

$$u_e(H;c) \equiv \theta_H\, q_H(\phi_s(c)) - c_d\, \tau_H(\phi_s(c)) + \gamma_H(c) \geq 0. \tag{13}$$

Here $u_e(L;c)$ and $u_e(H;c)$ denote the expected payoffs of class-L and class-H field evaluators, respectively. The equality in (12) follows from the minimal-inducement construction: whenever the supplier induces the report with the lower pre-inducement payoff, it compensates the class-L evaluator exactly for the payoff difference. For a class-H field evaluator, the one-sided reporting assumption implies truthful reporting. Because an H-contingent inducement is tied to the reported class rather than the evaluator's true class, it enters the class-H participation payoff in (13).

Finally, the mechanism is subject to feasibility constraints. For every reported fulfillment cost $\hat{c}$ and reported work-order class i, the documentation and authorization probabilities must lie in $[0, 1]$, and the expected net transfer must be nonnegative:

$$0 \leq \tau_i(\hat{c}) \leq 1, \qquad \forall (\hat{c}, i) \in [c_L, c_H] \times \{L, H\}, \tag{14}$$

$$0 \leq q_i(\hat{c}) \leq 1, \qquad \forall (\hat{c}, i) \in [c_L, c_H] \times \{L, H\}, \text{ and} \tag{15}$$

$$m_i(\hat{c}) \geq 0, \qquad \forall (\hat{c}, i) \in [c_L, c_H] \times \{L, H\}. \tag{16}$$

The administrator maximizes the expected customer or operating-unit value generated by work-order fulfillment, less expected net transfers to the supplier and documentation-review costs, plus a fraction $\alpha \in [0, 1]$ of the supplier's profit (Baron and Myerson 1982).[4] Given the equilibrium behavior described above, a supplier with fulfillment cost c reports $\phi_s(c)$. A class-L work order is reported as $\phi_e(c)$, whereas a class-H work order is reported as H under the one-sided

[4] An administrator may internalize part of the supplier's surplus in order to sustain service-network capacity and ensure the supplier's continued participation.

reporting assumption. Accounting for the corresponding inducements $\gamma_i(c)$, the administrator's expected payoff conditional on c, denoted by $U(c)$, is

$$\begin{aligned}
U(c) \equiv \rho_L \, \Big[& \theta_L \, q_{\phi_e(c)}(\phi_s(c)) - m_{\phi_e(c)}(\phi_s(c)) - c_p \, \tau_{\phi_e(c)}(\phi_s(c)) \\
& + \alpha \, \left[m_{\phi_e(c)}(\phi_s(c)) - c \, q_{\phi_e(c)}(\phi_s(c)) - \gamma_{\phi_e(c)}(c) \right] \Big] \\
& + \rho_H \, \Big[\theta_H \, q_H(\phi_s(c)) - m_H(\phi_s(c)) - c_p \, \tau_H(\phi_s(c)) + \alpha \, \left[m_H(\phi_s(c)) - c \, q_H(\phi_s(c)) - \gamma_H(c) \right] \Big] .
\end{aligned} \tag{17}$$

The administrator's problem can be written as

$$U^* \equiv \max_{\{\mathcal{M}, \phi_s, \phi_e, \gamma_L, \gamma_H\} \in \Omega} \int_{c_L}^{c_H} U(c) \, \mathrm{d}F(c), \tag{18}$$

where the feasible set Ω is defined by the supplier cost-reporting optimality and participation constraints (7)–(8); the induced field-evaluator reporting rule and participation constraints (11)–(13); the equilibrium inducement schedules (9)–(10); and the feasibility constraints (14)–(16).

The standard revelation principle cannot be invoked directly in the original game, because the supplier may use report-contingent inducements outside the administrator's mechanism to influence the field evaluator's reporting incentives. We therefore proceed in two steps. First, we study an implementable no-side-payment benchmark in which truthful work-order classification is sustained without supplier-financed inducements in equilibrium. Second, we formulate an auxiliary problem that treats intended supplier-financed inducements as decision variables. This auxiliary problem relaxes the original game, because the administrator cannot directly choose such inducements in the contracting environment. We then show that the relaxation is tight under class-contingent expected net transfers: by adjusting the supplier's expected net transfer across reported work-order classes, the administrator can make the intended inducement policy a best response for the supplier. Hence, the auxiliary optimum is implementable in the original game and characterizes the optimal solution to the administrator's problem.

4. No-Side-Payment Benchmark

In this section, we study an implementable no-side-payment benchmark. We restrict attention to direct mechanisms that induce truthful reporting by both the supplier and the class-L field evaluator along the equilibrium path, i.e.,

$$\phi_s(c) = c, \quad \phi_e(c) = L, \quad \forall c \in [c_L, c_H].$$

Because the model allows only class-L-to-H classification inflation, a class-H field evaluator has no strategic downgrade deviation; hence, no separate class-H incentive-compatibility constraint is required. We further require the equilibrium supplier-financed inducement schedules to be identically zero:

$$\gamma_L(c) = \gamma_H(c) = 0, \quad \forall c \in [c_L, c_H].$$

For the no-side-payment mechanism described above to induce truthful reporting in equilibrium, the class-L field evaluator's incentive-compatibility constraint is given by

$$\theta_L\, q_L(c) - c_d\, \tau_L(c) \geq \theta_L\, q_H(c) - c_d'\, \tau_H(c), \quad \forall c \in [c_L, c_H]. \tag{19}$$

This constraint says that, absent supplier-financed inducements, a class-L field evaluator weakly prefers truthful reporting to inflating the work order as class H. Under these conditions, the field-evaluator participation constraints defined in (12) and (13) reduce to

$$\theta_i\, q_i(c) - c_d\, \tau_i(c) \geq 0, \quad \forall (c, i) \in [c_L, c_H] \times \{L, H\}. \tag{20}$$

The supplier's payoff functions $\hat{u}_s(c, \hat{c})$ and $u_s(c)$, introduced in (5) and (7), then reduce to

$$\hat{u}_s(c, \hat{c}) = \sum_{i \in \{L,H\}} \rho_i\, [m_i(\hat{c}) - c\, q_i(\hat{c})], \quad \forall (c, \hat{c}) \in [c_L, c_H]^2, \tag{21}$$

and

$$u_s(c) = \sum_{i \in \{L,H\}} \rho_i\, [m_i(c) - c\, q_i(c)], \quad \forall c \in [c_L, c_H], \tag{22}$$

respectively. Thus, truthful reporting by the supplier requires

$$u_s(c) \geq \hat{u}_s(c, \hat{c}), \quad \forall (c, \hat{c}) \in [c_L, c_H]^2. \tag{23}$$

To ensure that the no-side-payment mechanism is implementable, we also require that the supplier not gain from a joint deviation in which it reports a cost $\hat{c}$ and offers a report-contingent inducement that leads class-L work orders to be reported as class H:

$$\sum_{i \in \{L,H\}} \rho_i\, [m_i(c) - c q_i(c)] \geq m_H(\hat{c}) - c q_H(\hat{c}) - \tilde{\gamma}_H(\hat{c}), \quad \forall (c, \hat{c}) \in [c_L, c_H]^2, \tag{24}$$

where $\tilde{\gamma}_H(\hat{c})$ is defined in (3).

Along the equilibrium path, conditional on the supplier's fulfillment cost c, the administrator's payoff function defined in (17) becomes

$$U(c) = \sum_{i \in \{L,H\}} \rho_i \left[\theta_i\, q_i(c) - m_i(c) - c_p\, \tau_i(c) + \alpha\, [m_i(c) - c\, q_i(c)]\right]. \tag{25}$$

The no-side-payment benchmark can be formulated as follows:

$$\underline{U} \equiv \max_{\{q_L, q_H, m_L, m_H, \tau_L, \tau_H\} \in \underline{\Omega}} \int_{c_L}^{c_H} U(c)\, \mathrm{d}F(c), \tag{26}$$

where the feasible set $\underline{\Omega}$ is defined by (8), (14), (15), (16), (19), (20), (23), and (24). The payoff functions $\hat{u}_s(c, \hat{c})$, $u_s(c)$, and $U(c)$ are defined in (21), (22), and (25), respectively.

The following proposition establishes that the benchmark value $\underline{U}$ is a lower bound on U^*.

PROPOSITION 1. *$\underline{U} \le U^*$.*

The following lemma shows that constraint (24) can be omitted from the benchmark problem (26) without loss of optimality.

LEMMA 1. *The optimal value of the benchmark problem* (26) *is unchanged if constraint* (24) *is omitted.*

By Lemma 1, we can aggregate the expected net transfers from the administrator to the supplier into a single decision variable:

$$m(c) \equiv \sum_{i \in \{L,H\}} \rho_i \, m_i(c), \quad \forall c \in [c_L, c_H]. \tag{27}$$

The aggregate transfer $m(c)$ is the administrator's expected per-work-order net transfer to the supplier, conditional on the supplier's fulfillment cost c. By Lemma 1, any feasible aggregate transfer $m(c)$ can be implemented through class-contingent net transfers that preserve $m(c)$ and satisfy (24).

Using the aggregate transfer $m(c)$, the supplier's payoff functions defined in (21) and (22) can be rewritten as

$$\hat{u}_s(c, \hat{c}) = m(\hat{c}) - c \sum_{i \in \{L,H\}} \rho_i \, q_i(\hat{c}), \qquad \forall (c, \hat{c}) \in [c_L, c_H]^2, \tag{28}$$

$$u_s(c) = m(c) - c \sum_{i \in \{L,H\}} \rho_i \, q_i(c), \qquad \forall c \in [c_L, c_H]. \tag{29}$$

Similarly, the administrator's conditional expected payoff defined in (25) can be reformulated as

$$U(c) = -(1-\alpha)\, m(c) + \sum_{i \in \{L,H\}} \rho_i \, \left[\theta_i \, q_i(c) - c_p \, \tau_i(c) - \alpha \, c \, q_i(c)\right], \quad \forall c \in [c_L, c_H]. \tag{30}$$

The closed-form characterizations below rely on the following regularity assumption.

ASSUMPTION 1. *The virtual valuation function*

$$w(c; \theta) \equiv \theta - c - (1-\alpha)\, \frac{F(c)}{f(c)}, \quad \forall c \in [c_L, c_H],$$

is non-increasing in c.

Assumption 1 is equivalent to requiring that the cost-dependent component

$$c + (1-\alpha)\frac{F(c)}{f(c)}$$

be nondecreasing in c. This condition holds under several commonly used distributions, including the uniform distribution and truncated exponential and Pareto distributions. Appendix EC.2 relaxes Assumption 1: for general fulfillment-cost distributions, the same work-order screening problem applies after replacing the virtual valuation function $w(c;\theta)$ with its ironed counterpart and pooling supplier cost types over ironing intervals.

In the aggregate-transfer formulation, the following proposition characterizes an optimal solution to the benchmark problem (26) when the class-L fulfillment value θ_L is sufficiently low.

PROPOSITION 2. *Under Assumption 1, if*

$$\theta_L \le \min\left\{\theta_H - \frac{c_p\,\theta_L}{c_d'\rho_L},\, c_d'\right\},$$

then an optimal solution to the benchmark problem (26) *is given as follows:*

$$\tau_L(c) = 0, \quad \tau_H(c) = \frac{\theta_L}{c_d'}\,\mathbb{1}_{c\in[\acute{c},\check{c})}, \quad q_L(c) = \mathbb{1}_{c\in[c_L,\acute{c})}, \quad q_H(c) = \mathbb{1}_{c\in[c_L,\check{c})}, \quad \forall c \in [c_L, c_H], \tag{31}$$

and

$$m(c) = \begin{cases} \rho_L\,\acute{c} + \rho_H\,\check{c}, & c \in [c_L, \acute{c}), \\ \rho_H\,\check{c}, & c \in [\acute{c}, \check{c}), \\ 0, & c \in [\check{c}, c_H], \end{cases} \tag{32}$$

where the thresholds $\acute{c}$ and $\check{c}$ are defined by

$$\begin{aligned} \acute{c} &\equiv \sup\left\{c \in [c_L, c_H] : c_d'\rho_L\, w(c;\theta_L) + c_p\rho_H\,\theta_L \ge 0\right\}, \\ \check{c} &\equiv \sup\left\{c \in [c_L, c_H] : c_d'\, w(c;\theta_H) - c_p\,\theta_L \ge 0\right\}. \end{aligned}$$

The mechanism in Proposition 2 has the following interpretation. When the supplier's fulfillment cost is low $(c \in [c_L, \acute{c}))$, both class-L and class-H work orders are authorized, and no documentation is required. When the fulfillment cost is intermediate $(c \in [\acute{c}, \check{c}))$, only class-$H$ work orders are authorized. In this region, documentation is required for reported class-H work orders with probability θ_L/c_d', which makes a class-L field evaluator indifferent between reporting truthfully and inflating the work-order class. This region may involve inefficiency relative to the first best: the threshold $\acute{c}$ may lie below the first-best cutoff θ_L, so some class-L work orders that would be efficient to fulfill are denied authorization. Finally, when the fulfillment cost is sufficiently high $(c \in [\check{c}, c_H])$, neither class is authorized.

As the class-L fulfillment value θ_L increases, the documentation probability θ_L/c_d' prescribed in Proposition 2 may exceed its upper bound of one. The following proposition characterizes the corresponding regime in which the documentation-probability constraint binds.

PROPOSITION 3. *Under Assumption 1, if*

$$c_d' < \theta_L \le \theta_H - \frac{c_p\theta_L}{c_d'\rho_L},$$

then an optimal solution to the benchmark problem (26) *is given as follows:*

$$\tau_L(c) = 0, \quad \tau_H(c) = \mathbb{1}_{c\in[\acute{c},\grave{c})}, \quad \forall c \in [c_L, c_H],$$
$$q_L(c) = \begin{cases} 1, & c \in [c_L, \acute{c}), \\ 1 - c_d'/\theta_L, & c \in [\acute{c}, \check{c}), \\ 0, & c \in [\check{c}, c_H], \end{cases} \qquad q_H(c) = \begin{cases} 1, & c \in [c_L, \check{c}), \\ c_d'/\theta_L, & c \in [\check{c}, \grave{c}), \\ 0, & c \in [\grave{c}, c_H]. \end{cases} \tag{33}$$

and

$$m(c) = \begin{cases} \dfrac{c_d'}{\theta_L}\left(\rho_L \acute{c} + \rho_H \grave{c}\right) + \left(1 - \dfrac{c_d'}{\theta_L}\right)\check{c}, & c \in [c_L, \acute{c}), \\ \left(1 - \dfrac{c_d'}{\theta_L}\right)\check{c} + \dfrac{c_d'\rho_H}{\theta_L}\grave{c}, & c \in [\acute{c}, \check{c}), \\ \dfrac{c_d'\rho_H}{\theta_L}\grave{c}, & c \in [\check{c}, \grave{c}), \\ 0, & c \in [\grave{c}, c_H], \end{cases} \tag{34}$$

where the thresholds $\acute{c}$, $\check{c}$, and $\grave{c}$ are defined by

$$\acute{c} \equiv \sup\left\{c \in [c_L, c_H] : c_d'\rho_L\, w(c;\theta_L) + c_p\rho_H\theta_L \geq 0\right\},$$
$$\check{c} \equiv \sup\left\{c \in [c_L, c_H] : \rho_L\, w(c;\theta_L) + \rho_H\, w(c;\theta_H) \geq 0\right\},$$
$$\grave{c} \equiv \sup\left\{c \in [c_L, c_H] : c_d'\, w(c;\theta_H) - c_p\theta_L \geq 0\right\}.$$

The optimal mechanism characterized in Proposition 3 has a similar structure. When the fulfillment cost is low $(c \in [c_L, \acute{c}))$, both work-order classes are authorized with probability one. For intermediate costs $(c \in [\acute{c}, \check{c}))$, class-$H$ work orders remain fully authorized. Because the class-L fulfillment value is high relative to the cost of supporting an inflated report, even requiring documentation with probability one would not deter inflation if class-L work orders were denied entirely. The administrator therefore authorizes class-L work orders with probability $1 - c_d'/\theta_L$, making the field evaluator indifferent between truthful reporting and inflation. As the fulfillment cost increases further $(c \in [\check{c}, \grave{c}))$, the administrator denies class-L work orders, requires documentation for reported class-H work orders, and authorizes reported class-H work orders with probability c_d'/θ_L. This again leaves a class-L field evaluator indifferent between truthful reporting and inflation. Finally, when the fulfillment cost is sufficiently high $(c \in [\grave{c}, c_H])$, neither class is authorized.

As θ_L increases further, the class-L fulfillment value becomes high enough that screening across work-order classes is no longer optimal. The following proposition characterizes the resulting pooling mechanism.

PROPOSITION 4. *Under Assumption 1, if*

$$\theta_L > \theta_H - \frac{c_p\theta_L}{c_d'\rho_L},$$

then an optimal solution to the benchmark problem (26) *is given as follows:*

$$\tau_L(c) = \tau_H(c) = 0, \quad q_L(c) = q_H(c) = \mathbb{1}_{c \in [c_L, \check{c})}, \quad m(c) = \check{c}\, \mathbb{1}_{c \in [c_L, \check{c})}, \quad \forall c \in [c_L, c_H], \tag{35}$$

where the threshold $\check{c}$ is defined by

$$\check{c} \equiv \sup \{ c \in [c_L, c_H] : \rho_L\, w(c; \theta_L) + \rho_H\, w(c; \theta_H) \geq 0 \} .$$

Because $q_L(c) = q_H(c)$ for all c, classification inflation has no effect on the authorization probability; hence, no documentation is required.

Together, Propositions 2–4 characterize the optimal no-side-payment benchmark across the parameter regimes considered above. By Proposition 1, the benchmark value $\underline{U}$ provides a lower bound on the administrator's optimal value U^* in (18).

5. Optimal Mechanism with Supplier-Financed Inducements

This section characterizes the administrator's optimal mechanism when supplier-financed inducements are available. Because the administrator cannot directly dictate the supplier's report-contingent payments to the field evaluator, we first formulate an auxiliary problem that treats the intended payments as decision variables. This relaxation yields an upper bound on the administrator's expected payoff. We then show that the relaxation is tight: in the contracting framework described in Section 3, the administrator can choose work-order-class-contingent expected net transfers to make the intended payments optimal for the supplier. Hence, the auxiliary solution characterizes the optimal mechanism in the original game.

In the auxiliary problem, with a slight abuse of notation, for each $(\hat{c}, i) \in [c_L, c_H] \times \{L, H\}$, $\gamma_i(\hat{c})$ denotes the intended payment from the supplier to the field evaluator when the supplier reports $\hat{c}$ and the field evaluator reports the work order as class i.

For a given supplier report $\hat{c}$, a class-L field evaluator obtains payoff

$$\theta_L q_L(\hat{c}) - c_d \tau_L(\hat{c}) + \gamma_L(\hat{c})$$

from truthful reporting and payoff

$$\theta_L q_H(\hat{c}) - c_d' \tau_H(\hat{c}) + \gamma_H(\hat{c})$$

from inflating the work order to class H. Therefore, when the work order is class L, the field evaluator's reporting rule in the auxiliary problem is

$$\phi_e(\hat{c}) = \begin{cases} L, & \text{if } \theta_L q_L(\hat{c}) - c_d \tau_L(\hat{c}) + \gamma_L(\hat{c}) \geq \theta_L q_H(\hat{c}) - c_d' \tau_H(\hat{c}) + \gamma_H(\hat{c}), \\ H, & \text{otherwise.} \end{cases} \tag{36}$$

The corresponding participation constraint for a class-L field evaluator is

$$\max\left\{\theta_L q_L(\hat{c}) - c_d\tau_L(\hat{c}) + \gamma_L(\hat{c}), \theta_L q_H(\hat{c}) - c_d'\tau_H(\hat{c}) + \gamma_H(\hat{c})\right\} \geq 0, \quad \forall \hat{c} \in [c_L, c_H]. \tag{37}$$

Under the one-sided reporting structure, a class-H field evaluator cannot misreport. Its participation constraint is

$$\theta_H q_H(\hat{c}) - c_d\tau_H(\hat{c}) + \gamma_H(\hat{c}) \geq 0, \quad \forall \hat{c} \in [c_L, c_H]. \tag{38}$$

Suppose a supplier with fulfillment cost c reports $\hat{c}$. Given the field evaluator's reporting rule ϕ_e for class-L work orders, the supplier's interim payoff in the auxiliary problem is

$$\hat{u}_s(c, \hat{c}) = \rho_L\left[m_{\phi_e(\hat{c})}(\hat{c}) - cq_{\phi_e(\hat{c})}(\hat{c}) - \gamma_{\phi_e(\hat{c})}(\hat{c})\right] + \rho_H\left[m_H(\hat{c}) - cq_H(\hat{c}) - \gamma_H(\hat{c})\right]. \tag{39}$$

The supplier's reporting-optimality and participation constraints take the same form as (7) and (8) in Section 3, evaluated with the auxiliary payoff in (39).

Under the prescribed reporting rules, the administrator's payoff conditional on a supplier fulfillment cost c is

$$\begin{aligned} U(c) = \rho_L \big[&\theta_L q_{\phi_e(\phi_s(c))}(\phi_s(c)) - m_{\phi_e(\phi_s(c))}(\phi_s(c)) - c_p\tau_{\phi_e(\phi_s(c))}(\phi_s(c)) \\ &+ \alpha\left(m_{\phi_e(\phi_s(c))}(\phi_s(c)) - cq_{\phi_e(\phi_s(c))}(\phi_s(c)) - \gamma_{\phi_e(\phi_s(c))}(\phi_s(c))\right)\big] \\ &+ \rho_H\left[\theta_H q_H(\phi_s(c)) - m_H(\phi_s(c)) - c_p\tau_H(\phi_s(c)) + \alpha\left(m_H(\phi_s(c)) - cq_H(\phi_s(c)) - \gamma_H(\phi_s(c))\right)\right]. \end{aligned} \tag{40}$$

The auxiliary problem is

$$\bar{U} \equiv \max_{\{q_i, m_i, \tau_i, \phi_s, \phi_e, \gamma_i\} \in \bar{\Omega}_0} \int_{c_L}^{c_H} U(c)\,\mathrm{d}F(c), \tag{41}$$

where $\bar{\Omega}_0$ is defined by the field evaluator's reporting rule for class-L work orders (36) and participation constraints (37)–(38), the supplier's reporting-optimality and participation constraints, which take the form of (7) and (8) with the supplier payoff $\hat{u}_s$ defined in (39), the feasibility constraints (14)–(16), and the nonnegativity constraint on intended supplier-financed payments:

$$\gamma_i(\hat{c}) \geq 0, \quad \forall (\hat{c}, i) \in [c_L, c_H] \times \{L, H\}. \tag{42}$$

The following proposition formalizes the upper-bound property of the auxiliary relaxation.

Proposition 5.

$$\bar{U} \geq U^*.$$

The next lemma justifies solving the auxiliary problem by restricting attention to direct incentive-compatible mechanisms.

LEMMA 2. *For any mechanism feasible for the auxiliary problem* (41)*, there exists a direct incentive-compatible mechanism that preserves the administrator's expected payoff. Moreover, for every* $c \in [c_L, c_H]$*, the direct mechanism preserves the supplier's interim payoff and the payoff of a class-H field evaluator, and weakly increases the payoff of a class-L field evaluator. Hence, restricting attention to direct incentive-compatible mechanisms is without loss of optimality in the auxiliary problem.*

By Lemma 2, it is without loss to solve the auxiliary problem over direct incentive-compatible mechanisms. Using the aggregate expected net transfer $m(c)$ defined in (27), the payoff functions can be rewritten as

$$u_e(i,i;c) = \theta_i \, q_i(c) - c_d \, \tau_i(c) + \gamma_i(c), \qquad \forall (c,i) \in [c_L, c_H] \times \{L, H\}, \tag{43}$$

$$u_e(L,H;c) = \theta_L \, q_H(c) - c_d' \, \tau_H(c) + \gamma_H(c), \qquad \forall c \in [c_L, c_H], \tag{44}$$

$$\hat{u}_s(c,\hat{c}) = m(\hat{c}) - \sum_{i \in \{L,H\}} \rho_i \left[c \, q_i(\hat{c}) + \gamma_i(\hat{c}) \right], \qquad \forall (c,\hat{c}) \in [c_L, c_H]^2, \tag{45}$$

$$u_s(c) = m(c) - \sum_{i \in \{L,H\}} \rho_i \left[c \, q_i(c) + \gamma_i(c) \right], \qquad \forall c \in [c_L, c_H], \tag{46}$$

$$U(c) = -(1-\alpha)\, m(c) + \sum_{i \in \{L,H\}} \rho_i \left[(\theta_i - \alpha \, c)\, q_i(c) - c_p \, \tau_i(c) - \alpha \, \gamma_i(c) \right], \quad \forall c \in [c_L, c_H]. \tag{47}$$

Under the one-sided reporting structure, only class-L work orders can be reported as class H. Thus, the field-evaluator incentive-compatibility constraint is

$$u_e(L,L;c) \geq u_e(L,H;c), \quad \forall c \in [c_L, c_H]. \tag{48}$$

The field-evaluator participation constraints are

$$u_e(i,i;c) \geq 0, \quad \forall (c,i) \in [c_L, c_H] \times \{L, H\}. \tag{49}$$

The supplier's direct incentive-compatibility and participation constraints are

$$u_s(c) \geq \hat{u}_s(c,\hat{c}), \qquad \forall (c,\hat{c}) \in [c_L, c_H]^2, \tag{50}$$

$$u_s(c) \geq 0, \qquad \forall c \in [c_L, c_H]. \tag{51}$$

The auxiliary problem can therefore be written in direct form as

$$\bar{U} = \max_{\{q_L, q_H, m, \tau_L, \tau_H, \gamma_L, \gamma_H\} \in \bar{\Omega}} \int_{c_L}^{c_H} U(c) \, \mathrm{d}F(c), \tag{52}$$

where the payoff functions u_e, $\hat{u}_s$, u_s, and U are defined in (43)–(47). The feasible set $\bar{\Omega}$ is defined by the field-evaluator incentive-compatibility and participation constraints (48)–(49), the supplier incentive-compatibility and participation constraints (50)–(51), the feasibility constraints on documentation and authorization probabilities (14)–(15), and the nonnegativity constraints on intended payments (42).

We now characterize the optimal solutions to the direct auxiliary problem (52) under Assumption 1. We begin with parameter regimes in which the no-side-payment benchmark mechanisms from Section 4 remain optimal even when supplier-financed inducements are available.

PROPOSITION 6. *Under Assumption 1, the no-side-payment benchmark mechanisms characterized in Propositions 2–4 remain optimal for* (52) *in the following parameter regimes:*

1. *If $\theta_H \geq \theta_L + c_p\theta_L/(c_d'\rho_L)$, $\theta_L \leq c_d'$, and $c_d'\rho_L - c_p\rho_H \geq 0$, an optimal mechanism is given by* (31), (32), *and $\gamma_L(c) = \gamma_H(c) = 0$ for all $c \in [c_L, c_H]$.*
2. *If $\theta_L + c_p\theta_L/(c_d'\rho_L) \leq \theta_H \leq \theta_L + \theta_L/\rho_H$, $\theta_L \geq c_d'$, and $c_d'\rho_L - c_p\rho_H \geq 0$, an optimal mechanism is given by* (33), (34), *and $\gamma_L(c) = \gamma_H(c) = 0$ for all $c \in [c_L, c_H]$.*
3. *If one of the following conditions holds: (i) $\theta_H \leq \theta_L + c_p\theta_L/(c_d'\rho_L)$ and $c_d'\rho_L - c_p\rho_H \geq 0$, or (ii) $\theta_H \leq \theta_L + \theta_L/\rho_H$ and $c_d'\rho_L - c_p\rho_H \leq 0$, an optimal mechanism is given by* (35) *and $\gamma_L(c) = \gamma_H(c) = 0$ for all $c \in [c_L, c_H]$.*

Proposition 6 identifies regimes in which the direct auxiliary problem (52) attains the same objective value as the no-side-payment benchmark (26). In these regimes, there exists an optimal auxiliary mechanism with $\gamma_L(c) = \gamma_H(c) = 0$ for all $c \in [c_L, c_H]$.

Outside these regimes, supplier-financed payments can increase the objective value relative to the no-side-payment benchmark. The next proposition characterizes the first case with positive supplier-financed payments. The resulting mechanism mirrors the two-threshold structure of the separating benchmark in Proposition 2: class-H work orders are authorized over a larger range of supplier costs than class-L work orders. The key difference is that, in the intermediate interval, truthful class-L reporting is sustained by a supplier-financed payment rather than by documentation.

PROPOSITION 7. *Under Assumption 1, if $\theta_H \geq \theta_L + \theta_L/\rho_H$ and $c_d'\rho_L - c_p\rho_H \leq 0$, then an optimal solution to* (52) *is given by*

$$\tau_L(c) = \tau_H(c) = \gamma_H(c) = 0, \quad \gamma_L(c) = \theta_L \cdot \mathbb{1}_{c \in [\bar{c}, \acute{c})}, \quad q_L(c) = \mathbb{1}_{c \in [c_L, \bar{c})}, \quad q_H(c) = \mathbb{1}_{c \in [c_L, \acute{c})}, \quad \forall c \in [c_L, c_H],$$

and

$$m(c) = \begin{cases} \rho_L \bar{c} + \rho_H \acute{c}, & c \in [c_L, \bar{c}), \\ \rho_L \theta_L + \rho_H \acute{c}, & c \in [\bar{c}, \acute{c}), \\ 0, & c \in [\acute{c}, c_H], \end{cases}$$

where the thresholds $\bar{c}$ and $\acute{c}$ are defined by

$$\bar{c} \equiv \sup\left\{c \in [c_L, c_H] : w(c;\theta_L) + \theta_L \geq 0\right\},$$
$$\acute{c} \equiv \sup\left\{c \in [c_L, c_H] : -\rho_L\theta_L + \rho_H w(c;\theta_H) \geq 0\right\}.$$

The mechanism in Proposition 7 has the following interpretation. When the supplier's fulfillment cost is low, $c \in [c_L, \bar{c})$, both work-order classes are authorized with probability one, and the mechanism uses neither documentation nor supplier-financed payments to the field evaluator. When the fulfillment cost is intermediate, $c \in [\bar{c}, \acute{c})$, only class-$H$ work orders are authorized. Since no documentation is required in this interval, a class-L field evaluator would prefer to inflate the work order to class H in the absence of a supplier-financed payment. The auxiliary mechanism instead sets $\gamma_L(c) = \theta_L$, making the field evaluator indifferent between truthful reporting and inflation. This payment serves as a substitute for documentation as an incentive instrument and avoids the administrator's review cost. Finally, when the fulfillment cost is high, $c \in [\acute{c}, c_H]$, neither work-order class is authorized, and no documentation or supplier-financed payment is used.

The improvement in the administrator's objective value need not come at the expense of the supplier or the field evaluator. In some parameterizations, supplier-financed payments generate a weak Pareto improvement relative to the no-side-payment benchmark: the supplier's interim payoff and the field evaluator's conditional payoff for each work-order class are weakly higher as well. Figure EC.1 in Appendix EC.1 provides a numerical illustration.

The second positive-payment case arises when θ_L is high enough relative to c_d', while class-H work orders remain sufficiently more valuable than class-L work orders. The resulting mechanism mirrors the multi-threshold structure in Proposition 3, but now documentation and supplier-financed payments coexist.

PROPOSITION 8. *Under Assumption 1, if $\theta_H \geq \theta_L + \theta_L/\rho_H$, $c_d'\rho_L - c_p\rho_H \geq 0$, and $c_d' \leq \theta_L$, then an optimal solution to* (52) *is given by*

$$\tau_L(c) = \gamma_H(c) = 0, \quad \tau_H(c) = \mathbb{1}_{c \in [\acute{c}, \dot{c})}, \quad \gamma_L(c) = (\theta_L - c_d') \cdot \mathbb{1}_{c \in [\check{c}, \grave{c})}, \quad \forall c \in [c_L, c_H],$$

$$q_L(c) = \begin{cases} 1, & c \in [c_L, \acute{c}), \\ 1 - c_d'/\theta_L, & c \in [\acute{c}, \check{c}), \\ 0, & c \in [\check{c}, c_H], \end{cases} \qquad q_H(c) = \begin{cases} 1, & c \in [c_L, \grave{c}), \\ c_d'/\theta_L, & c \in [\grave{c}, \dot{c}), \\ 0, & c \in [\dot{c}, c_H], \end{cases}$$

and

$$m(c) = \begin{cases} \dfrac{c_d'}{\theta_L}\rho_L \acute{c} + \left(1 - \dfrac{c_d'}{\theta_L}\right)\rho_L \check{c} + \left(1 - \dfrac{c_d'}{\theta_L}\right)\rho_H \grave{c} + \dfrac{c_d'}{\theta_L}\rho_H \dot{c}, & c \in [c_L, \acute{c}), \\ \left(1 - \dfrac{c_d'}{\theta_L}\right)\rho_L \check{c} + \left(1 - \dfrac{c_d'}{\theta_L}\right)\rho_H \grave{c} + \dfrac{c_d'}{\theta_L}\rho_H \dot{c}, & c \in [\acute{c}, \check{c}), \\ \rho_L(\theta_L - c_d') + \left(1 - \dfrac{c_d'}{\theta_L}\right)\rho_H \grave{c} + \dfrac{c_d'}{\theta_L}\rho_H \dot{c}, & c \in [\check{c}, \grave{c}), \\ \dfrac{c_d'}{\theta_L}\rho_H \dot{c}, & c \in [\grave{c}, \dot{c}), \\ 0, & c \in [\dot{c}, c_H], \end{cases}$$

where the thresholds $\acute{c}$, $\check{c}$, $\grave{c}$, and $\dot{c}$ are defined by

$$\acute{c} \equiv \sup\left\{c \in [c_L, c_H] : c_d' \rho_L w(c;\theta_L) + c_p \rho_H \theta_L \geq 0\right\},$$
$$\check{c} \equiv \sup\left\{c \in [c_L, c_H] : w(c;\theta_L) + \theta_L \geq 0\right\},$$
$$\grave{c} \equiv \sup\left\{c \in [c_L, c_H] : -\rho_L \theta_L + \rho_H w(c;\theta_H) \geq 0\right\},$$
$$\dot{c} \equiv \sup\left\{c \in [c_L, c_H] : c_d' w(c;\theta_H) - c_p \theta_L \geq 0\right\}.$$

The mechanism in Proposition 8 can be interpreted as follows. For $c \in [c_L, \acute{c})$, both work-order classes are authorized with probability one, and the mechanism uses neither documentation nor supplier-financed payments. For $c \in [\acute{c}, \check{c})$, the authorization and documentation rules coincide with those in Proposition 3: class-H work orders are fully authorized, class-L work orders are authorized with probability $1 - c_d'/\theta_L$, and reported class-H work orders are documented with probability one. These rules already make a class-L field evaluator indifferent between truthful reporting and inflation, so no supplier-financed payment is needed. For $c \in [\check{c}, \grave{c})$, only class-$H$ work orders are authorized, and reported class-H work orders are documented with probability one. Absent a supplier-financed payment, a class-L field evaluator would obtain zero from truthful reporting but $\theta_L - c_d'$ from inflation. The auxiliary mechanism therefore sets $\gamma_L(c) = \theta_L - c_d'$, restoring indifference between truthful reporting and inflation. For $c \in [\grave{c}, \dot{c})$, class-$H$ work orders are authorized only with probability c_d'/θ_L, while reported class-H work orders continue to be documented with probability one. This reduced authorization probability drives the class-L evaluator's inflation payoff to zero, so no supplier-financed payment is needed. Finally, for $c \in [\dot{c}, c_H]$, neither work-order class is authorized, and the mechanism again uses neither documentation nor supplier-financed payments.

Figure EC.2 in Appendix EC.1 provides the analogous payoff comparison for Proposition 8, where documentation and supplier-financed inducements coexist. In this example, the supplier-financed mechanism again generates a weak Pareto improvement relative to the corresponding no-side-payment benchmark.

Propositions 6–8 characterize the optimal solutions to (52). It remains to show that the relaxation is tight: for each solution $(q_i, \tau_i, \gamma_i, m)$ characterized above, there exists a mechanism feasible for the original problem (18) that implements the same authorization and documentation rules (q_i, τ_i), makes it optimal for the supplier to offer the intended supplier-financed payments γ_i, and preserves the aggregate expected net transfer $m(c)$ defined in (27).

The key implementation step is to choose, for each supplier cost report $\hat{c}$, the expected net transfers $m_L(\hat{c})$ and $m_H(\hat{c})$ from the administrator to the supplier, contingent on the reported work-order class. The split must first satisfy

$$\rho_L m_L(\hat{c}) + \rho_H m_H(\hat{c}) = m(\hat{c}), \quad \forall \hat{c} \in [c_L, c_H],$$

so that the aggregate expected net transfer is preserved. These class-contingent transfers must also satisfy the supplier's implementation incentives. In all mechanisms characterized above, $\gamma_H(\hat{c}) = 0$, and $\gamma_L(\hat{c})$ is the minimal payment that makes a class-L field evaluator weakly prefer truthful reporting to inflation. The transfer split is chosen so that, conditional on a class-L work order, the supplier weakly prefers to offer $\gamma_L(\hat{c})$ and induce truthful reporting rather than induce a class-H report. The proof of Proposition 9 verifies that such a split exists for each characterized mechanism.

PROPOSITION 9. *The optimal auxiliary mechanisms characterized in Propositions 6–8 are implementable in the original problem* (18)*. Consequently,*

$$U^* = \bar{U}.$$

Proposition 9 establishes that the auxiliary relaxation is tight. The administrator need not directly control supplier-to-field-evaluator payments: by reallocating the aggregate expected net transfer $m(c)$ across reported work-order classes, it can make the intended supplier-financed payments an equilibrium choice for the supplier. Thus, the no-side-payment problem in Section 4 serves as a benchmark, while the mechanisms characterized in this section solve the original problem with supplier-financed inducements.

6. Comparative Statics: Documentation Technology

This section studies how documentation and review technologies affect the optimal mechanisms characterized in Section 5. We interpret these technologies broadly: they may reduce the field evaluator's cost of preparing documentation for a truthful report, c_d, the cost of supporting an inflated class-H report, c_d', or the administrator's review cost, c_p. For example, electronic work-order platforms, standardized checklists, mobile photo capture, and AI-assisted report drafting may reduce the routine burden of preparing documentation, while automated triage, duplicate detection, and AI-assisted review may reduce the administrator's review cost. At the same time, some of these tools may also make weakly supported high-impact classifications easier to justify, thereby reducing the cost of substantiating an inflated report. The distinction between these costs is important. A reduction in review or routine filing costs lowers the cost of using documentation as an incentive instrument, whereas a reduction in the cost of substantiating inflated reports may weaken the screening role of documentation.

The results below are local comparative statics. That is, they apply to changes in costs that leave the model parameters in the same policy regime. Changes that cross a regime boundary fall outside these local results; after such a change, the optimal mechanism must be obtained by applying the relevant characterization in Propositions 6–8 to the new parameter region.

For each work-order class $i \in \{L, H\}$, let

$$\mathcal{V}_i \equiv \int_{c_L}^{c_H} [\theta_i q_i(c) - c_d \tau_i(c) + \gamma_i(c)] \,\mathrm{d}F(c)$$

denote the field evaluator's expected payoff in equilibrium, conditional on handling a class-i work order, under the optimal mechanism in the relevant regime.

We consider two local cost changes. First, all documentation and review costs fall proportionally:

$$\hat{c}_d = \beta c_d, \quad \hat{c}'_d = \beta c'_d, \quad \hat{c}_p = \beta c_p, \quad \beta \in (0, 1].$$

Second, only the field evaluator's cost of supporting an inflated class-H report falls:

$$\hat{c}_d = c_d, \quad \hat{c}_p = c_p, \quad \hat{c}'_d < c'_d.$$

We begin with the regimes covered by Proposition 6, where the auxiliary problem is solved by one of the no-side-payment benchmark mechanisms from Section 4. Because these mechanisms have $\gamma_L(c) = \gamma_H(c) = 0$ for all $c \in [c_L, c_H]$, the comparative statics follow from how documentation and review costs affect the benchmark authorization and documentation rules.

PROPOSITION 10. *Suppose the parameters lie in one of the regimes covered by Proposition 6, so that an optimal mechanism coincides with one of the no-side-payment benchmark mechanisms. Then the following local comparative statics hold.*

1. *In the low-θ_L separating regime of Proposition 6(1), a proportional reduction in c_d, c'_d, and c_p leaves U^*, $\mathcal{V}_L$, $\mathcal{V}_H$, and the supplier's interim payoff $u_s(c)$ unchanged for all $c \in [c_L, c_H]$. A reduction in c'_d alone, holding c_d and c_p fixed, weakly decreases U^*, weakly increases $\mathcal{V}_L$, and has an ambiguous effect on $\mathcal{V}_H$.*
2. *In the intermediate-θ_L separating regime of Proposition 6(2), both a proportional reduction in c_d, c'_d, and c_p and a reduction in c'_d alone, holding c_d and c_p fixed, weakly decrease U^*, weakly increase $\mathcal{V}_L$, and have an ambiguous effect on $\mathcal{V}_H$.*
3. *In the pooling regime of Proposition 6(3), both cost changes leave U^*, $\mathcal{V}_L$, $\mathcal{V}_H$, and the supplier's interim payoff $u_s(c)$ unchanged for all $c \in [c_L, c_H]$.*

Proposition 10 highlights two opposing effects of technologies that affect documentation and review costs. Reductions in review costs or routine filing burdens for truthful reports make documentation less costly as an incentive instrument. By contrast, a reduction in c'_d, the cost of supporting an inflated class-H report, makes inflation easier and weakens documentation's screening power. In the separating regimes, this weakly increases the expected payoff of a class-L field evaluator, but it can reduce the administrator's objective value and leaves the effect on the class-H field evaluator ambiguous.

We next consider the regimes in which supplier-financed payments are positive.

PROPOSITION 11. *Suppose the parameters lie in one of the regimes covered by Propositions 7 or 8, and consider local cost changes that keep the parameters in the same regime. The following comparative statics hold.*

1. *If the relevant mechanism is given by Proposition 7, then neither a proportional reduction in c_d, c'_d, and c_p, nor a reduction in c'_d alone, holding c_d and c_p fixed, has any local effect on the optimal mechanism, the administrator's objective value U^*, the supplier's interim payoff $u_s(c)$ for any $c \in [c_L, c_H]$, or the field evaluator's expected payoffs $\mathcal{V}_L$ and $\mathcal{V}_H$.*

2. *If the relevant mechanism is given by Proposition 8, then each of the two cost changes considered above—a proportional reduction in c_d, c'_d, and c_p, and a reduction in c'_d alone holding c_d and c_p fixed—weakly decreases the administrator's objective value U^*, weakly increases the class-L field evaluator's expected payoff $\mathcal{V}_L$, and has an ambiguous effect on the class-H field evaluator's expected payoff $\mathcal{V}_H$.*

Proposition 11 distinguishes the two positive-payment regimes. When the mechanism in Proposition 7 applies, documentation is not used: the thresholds $\bar{c}$ and $\acute{c}$ are independent of c_d, c'_d, and c_p, so the local cost changes considered above do not affect the mechanism or equilibrium payoffs. By contrast, when the mechanism in Proposition 8 applies, documentation and supplier-financed payments coexist. A lower c'_d makes class inflation easier for a class-L field evaluator, and the optimal response weakly increases $\mathcal{V}_L$, weakly decreases U^*, and leaves the effect on $\mathcal{V}_H$ ambiguous. Together with Proposition 10, this shows that documentation and review technologies are most valuable to the administrator when they reduce routine filing or review burdens without also lowering the cost of supporting inflated reports.

7. Managerial Implications

The analysis offers several implications for administrators of decentralized maintenance service networks. A central message is that authorization rules, documentation requirements, and supplier payments should be designed as an integrated system rather than as separate governance tools.

First, documentation should be treated as a targeted incentive instrument, not as a default response to classification risk. Documentation can deter classification inflation by making inflated high-impact reports more costly for field evaluators. However, its value depends on whether class-based screening is worth preserving. When low-impact work orders generate low value, distinguishing low- from high-impact work orders can prevent inefficient authorization, and this screening may be supported by documentation or by class-contingent supplier compensation. As low-impact work orders become more valuable, maintaining sharp class distinctions becomes more costly, and the optimal policy may soften the distinction through probabilistic authorization.

When low-impact work orders are sufficiently valuable, pooling becomes optimal. Thus, administrators should first evaluate whether the benefit of separating work-order classes justifies the cost of sustaining truthful classification.

Second, supplier compensation should be evaluated not only for cost elicitation, but also for its effect on classification incentives. A supplier may benefit from high-impact classifications because they increase authorization likelihood or expected net compensation. Payment rules that make such classifications more profitable can therefore unintentionally encourage suppliers to influence field evaluators. In practice, class-contingent reimbursement, service-level payments, award fees, chargebacks, retainage, and end-of-period reconciliation should be calibrated together with authorization and documentation rules. When class-contingent expected net transfers are feasible, compensation can be designed so that offering the intended inducements is weakly optimal for the supplier.

Third, supplier inducements should be viewed as endogenous to contract design rather than merely as behavior to prohibit after the fact. In service networks, inducements may take noncash forms such as reimbursement for diagnostic effort, documentation assistance, expedited future service, informal technical support, or absorption of customer or operating-unit costs. Some of these practices may improve operations, but they can also lower the field evaluator's effective cost of supporting an inflated high-impact classification. The model therefore suggests that administrators should examine whether the contract gives suppliers a strong reason to encourage high-impact reports. In particular, administrator-to-supplier transfers, authorization rules, and documentation or review requirements should be designed jointly so that supplier support does not make inflated classifications attractive.

Fourth, documentation technology can have mixed effects. Technologies that reduce routine filing or review costs can improve processing efficiency. However, if they also reduce the incremental burden of substantiating inflated high-impact classifications, they may make classification inflation easier. For example, standardized templates, automated evidence generation, or AI-assisted documentation may reduce administrative effort while also making weakly supported high-impact claims easier to justify. Such technologies are most valuable when they reduce routine processing costs while preserving the evidentiary burden, audit exposure, or reputational risk associated with inflated classifications.

Overall, effective governance requires aligning all three instruments: documentation affects the evaluator's cost of substantiating reports, supplier payments affect the supplier's incentive to encourage particular classifications, and authorization probabilities determine the value of classification inflation.

8. Conclusion

This paper studies work-order authorization in decentralized maintenance service networks. In these networks, diagnosis, authorization, and fulfillment are separated across parties: a field evaluator observes diagnostic and contextual information about the work-order class, while a supplier privately observes the fulfillment cost. This separation creates two adverse-selection problems—cost overstatement by the supplier and classification inflation by the field evaluator—as well as a cross-agent incentive problem: the supplier may benefit from high-impact classifications and may therefore offer report-contingent inducements that make classification inflation attractive. We develop a mechanism-design model to examine how an administrator should jointly design authorization probabilities, documentation requirements, and supplier payments in this environment.

The analysis highlights a simple tradeoff: class-based screening can prevent low-value authorization, but truthful reporting is costly to sustain when supplier inducements may offset the evaluator's documentation burden. As the value of low-impact work orders rises, the optimal policy moves from class-based screening, supported by documentation or class-contingent supplier compensation, to probabilistic authorization that softens the class distinction, and eventually to pooling. This progression shows why stricter documentation is not always better: documentation is valuable only when the benefit of separating low- and high-impact work orders exceeds the cost of sustaining truthful classification.

Supplier inducements make the design problem nonstandard because the compensation scheme affects not only cost-reporting incentives but also the supplier's incentive to influence the evaluator's reported classification. We address this issue by comparing a no-side-payment benchmark with an auxiliary problem that treats supplier-financed transfers as decision variables. We show that the resulting upper bound is tight in the expected-net-transfer contracting environment: by reallocating aggregate expected net transfers across reported work-order classes, the administrator can make the intended supplier-financed payments optimal for the supplier. Thus, class-contingent expected net transfers can internalize supplier inducements without requiring the administrator to observe or directly control off-contract transfers.

Our analysis suggests several directions for future research. One direction is to study richer service networks with multiple suppliers, multiple field evaluators, or endogenous matching between evaluators and suppliers. Another is to analyze dynamic authorization policies in which documentation histories, prior classification behavior, and suppliers' historical cost reports affect future authorization and payment rules. A third direction is to enrich the documentation technology by allowing documentation to generate verifiable signals, vary in precision or manipulability, or differ across tools such as standardized templates, automated evidence generation, and AI-assisted documentation. Empirical work could examine how authorization thresholds, documentation requirements, and supplier compensation rules affect classification behavior in maintenance, warranty, and facilities-management settings.

References

Alatas V, Purnamasari R, Wai-Poi M, Banerjee A, Olken BA, Hanna R (2016) Self-targeting: Evidence from a field experiment in indonesia. *Journal of Political Economy* 124(2):371–427.

Babich V, Tang CS (2012) Managing opportunistic supplier product adulteration: Deferred payments, inspection, and combined mechanisms. *Manufacturing & Service Operations Management* 14(2):301–314.

Baiman S, Fischer PE, Rajan MV (2000) Information, contracting, and quality costs. *Management Science* 46(6):776–789.

Bakshi N, Kim SH, Savva N (2015) Signaling new product reliability with after-sales service contracts. *Management Science* 61(8):1812–1829.

Baron DP, Myerson RB (1982) Regulating a monopolist with unknown costs. *Econometrica* 911–930.

Besley T, Coate S (1992) Workfare versus welfare: Incentive arguments for work requirements in poverty-alleviation programs. *American Economic Review* 82(1):249–261.

Border KC, Sobel J (1987) Samurai accountant: A theory of auditing and plunder. *Review of Economic Studies* 54(4):525–540.

Chan TH, de Véricourt F, Besbes O (2019) Contracting in medical equipment maintenance services: An empirical investigation. *Management Science* 65(3):1136–1150.

Chen L, Lee HL (2017) Sourcing under supplier responsibility risk: The effects of certification, audit, and contingency payment. *Management Science* 63(9):2795–2812.

Chen L, Yao S, Zhu K (2020) Responsible sourcing under supplier-auditor collusion. *Manufacturing & Service Operations Management* 22(6):1234–1250.

Davis AM, Hu B, Hyndman K, Qi A (2022) Procurement for assembly under asymmetric information: Theory and evidence. *Management Science* 68(4):2694–2713.

Dukes A, Zhu Y (2019) Why customer service frustrates consumers: Using a tiered organizational structure to exploit hassle costs. *Marketing Science* 38(3):500–515.

Fan X, Chen YJ, Tang CS (2021) Mechanism design for managing hidden rebates and inflated quotes of a procurement service provider. *Manufacturing & Service Operations Management* 23(5):1275–1296.

Gale D, Hellwig M (1985) Incentive-compatible debt contracts: The one-period problem. *Review of Economic Studies* 52(4):647–663.

Green JR, Laffont JJ (1986) Partially verifiable information and mechanism design. *Review of Economic Studies* 53(3):447–456.

Guajardo JA, Cohen MA, Kim SH, Netessine S (2012) Impact of performance-based contracting on product reliability: An empirical analysis. *Management Science* 58(5):961–979.

Hsiao L, Chen YJ (2014) Return policy: Hassle-free or your money-back guarantee? *Naval Research Logistics* 61(5):403–417.

Hu B, Qi A (2018) Optimal procurement mechanisms for assembly. *Manufacturing & Service Operations Management* 20(4):655–666.

Hwang I, Radhakrishnan S, Su L (2006) Vendor certification and appraisal: Implications for supplier quality. *Management Science* 52(10):1472–1482.

Jain N, Hasija S, Popescu DG (2013) Optimal contracts for outsourcing of repair and restoration services. *Operations Research* 61(6):1295–1311.

Jiang B, Zhang X (2011) How does a retailer's service plan affect a manufacturer's warranty? *Management Science* 57(4):727–740.

Kim SH, Cohen MA, Netessine S (2007) Performance contracting in after-sales service supply chains. *Management Science* 53(12):1843–1858.

Kim SH, Cohen MA, Netessine S, Veeraraghavan S (2010) Contracting for infrequent restoration and recovery of mission-critical systems. *Management Science* 56(9):1551–1567.

Kofman F, Lawarrée J (1993) Collusion in hierarchical agency. *Econometrica* 629–656.

Krishna V (2009) *Auction theory* (Academic Press).

Laffont JJ, Martimort D (1997) Collusion under asymmetric information. *Econometrica* 875–911.

Laffont JJ, Martimort D (2000) Mechanism design with collusion and correlation. *Econometrica* 68(2):309–342.

Mookherjee D, Png I (1989) Optimal auditing, insurance, and redistribution. *Quarterly Journal of Economics* 104(2):399–415.

Nichols AL, Zeckhauser RJ (1982) Targeting transfers through restrictions on recipients. *American Economic Review* 72(2):372–377.

Padmanabhan V, Rao RC (1993) Warranty policy and extended service contracts: Theory and an application to automobiles. *Marketing Science* 12(3):230–247.

Plambeck EL, Zenios SA (2000) Performance-based incentives in a dynamic principal-agent model. *Manufacturing & Service Operations Management* 2(3):240–263.

Rui H, Lai G (2015) Sourcing with deferred payment and inspection under supplier product adulteration risk. *Production and Operations Management* 24(6):934–946.

Shulman JD, Coughlan AT, Savaskan RC (2011) Managing consumer returns in a competitive environment. *Management Science* 57(2):347–362.

Su X (2009) Consumer returns policies and supply chain performance. *Manufacturing & Service Operations Management* 11(4):595–612.

Tian F, Sun P, Duenyas I (2021) Optimal contract for machine repair and maintenance. *Operations Research* 69(3):916–949.

Tirole J (1986) Hierarchies and bureaucracies: On the role of collusion in organizations. *Journal of Law, Economics, and Organization* 2(2):181–214.

Townsend RM (1979) Optimal contracts and competitive markets with costly state verification. *Journal of Economic Theory* 21(2):265–293.

Wang S, de Véricourt F, Sun P (2024) Audit and remediation strategies in the presence of evasion capabilities. *Operations Research* 72(5):1843–1860.

Wang S, Sun P, de Véricourt F (2016) Inducing environmental disclosures: A dynamic mechanism design approach. *Operations Research* 64(2):371–389.

Yang L, Shang W, Liu L (2023) Follow the crowd with uncertain service capacity. *Manufacturing & Service Operations Management* 25(1):341–352.

Online Appendix

EC.1. Additional Figures

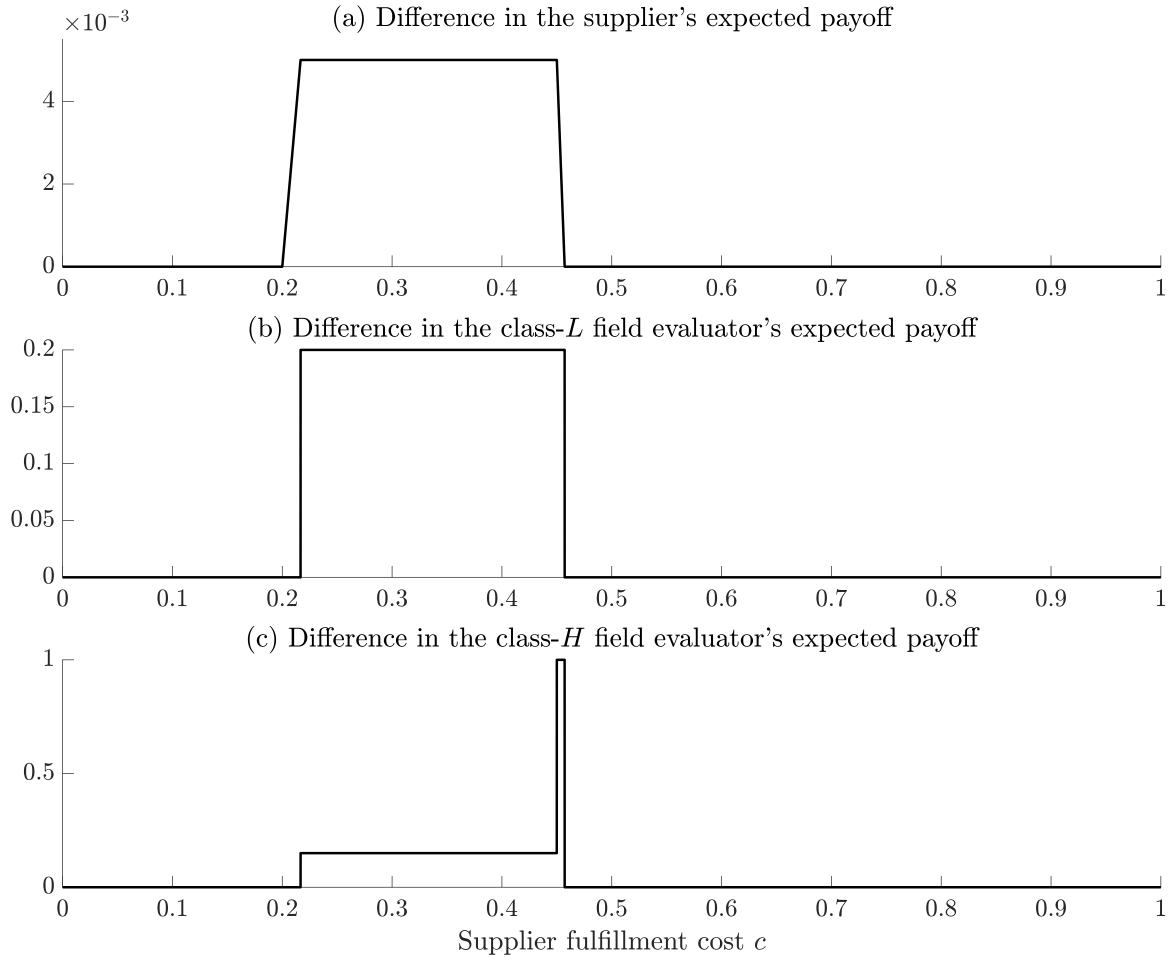

Figure EC.1 Payoff differences between the supplier-financed mechanism and the no-side-payment benchmark as a function of the supplier's fulfillment cost c. Panel (a) reports the supplier's interim payoff difference. Panels (b) and (c) report the corresponding payoff differences for field evaluators handling class-L and class-H work orders, respectively. The model parameters are $c_L = 0$, $c_H = 1$, $\theta_L = 0.2$, $\theta_H = 1$, $\alpha = 0$, $c_d = 0.15$, $c_d' = 0.2$, $c_p = 0.1$, $\rho_L = 0.3$, $\rho_H = 0.7$, and $f(c) = 1$ for all $c \in [0, 1]$.

EC.2. General Fulfillment-Cost Distributions

The closed-form characterizations in Sections 4 and 5 rely on Assumption 1, which ensures that the supplier's virtual fulfillment cost

$$\eta(c) \equiv c + (1-\alpha)\frac{F(c)}{f(c)}$$

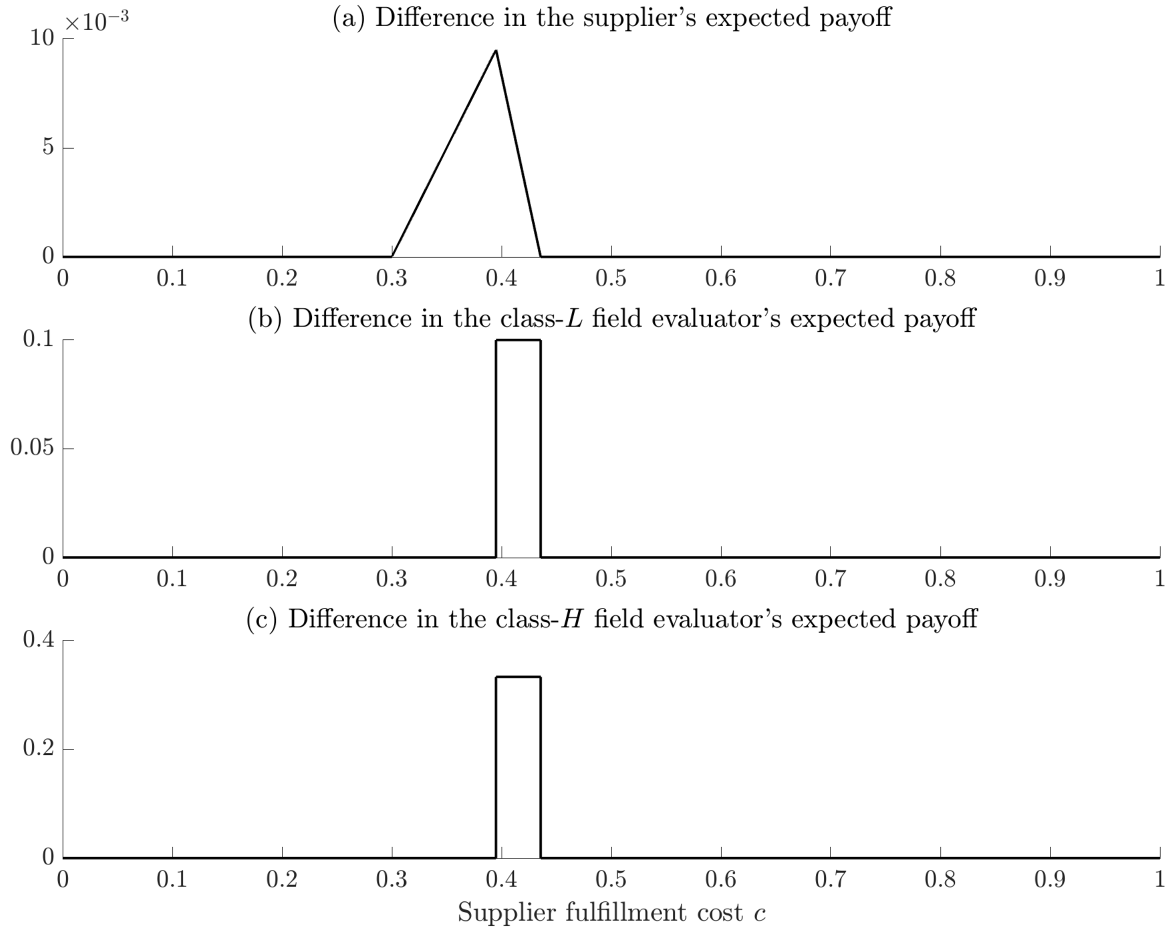

Figure EC.2 Payoff differences between the supplier-financed mechanism and the no-side-payment benchmark as a function of the supplier's fulfillment cost c. Panel (a) reports the supplier's interim payoff difference. Panels (b) and (c) report the corresponding payoff differences for field evaluators handling class-L and class-H work orders, respectively. The model parameters are $c_L = 0$, $c_H = 1$, $\theta_L = 0.3$, $\theta_H = 1$, $\alpha = 0$, $c_d = 0.15$, $c'_d = 0.2$, $c_p = 0.08$, $\rho_L = 0.3$, $\rho_H = 0.7$, and $f(c) = 1$ for all $c \in [0,1]$.

is nondecreasing in c. Under this regularity condition, the pointwise optimal authorization rule satisfies the supplier monotonicity constraint. We now show that this assumption is not essential for the no-side-payment benchmark and the auxiliary supplier-financed problem. When the fulfillment-cost distribution is irregular, the field-evaluator screening subproblem conditional on a fixed aggregate authorization probability Q remains the same; only the supplier-cost term in the reduced objective is replaced by its ironed counterpart.

Throughout this appendix, assume that F is absolutely continuous on $[c_L, c_H]$, with density $f(c) > 0$ on (c_L, c_H), and that η is integrable.

We use $k \in \{\mathrm{LB}, \mathrm{UB}\}$ to index two problems: LB denotes the no-side-payment benchmark, and UB denotes the auxiliary supplier-financed problem. Define

$$Q(c) \equiv \rho_L q_L(c) + \rho_H q_H(c),$$

the aggregate authorization probability conditional on supplier cost c. For each fixed Q, the administrator chooses the class-specific authorization probabilities, documentation probabilities, and, in the auxiliary problem, intended supplier-to-evaluator payments. We define the corresponding conditional screening values as follows.

For the no-side-payment benchmark, define

$$\begin{aligned} B^{\mathrm{LB}}(Q) \equiv \max_{q_L, q_H, \tau_L, \tau_H} \quad & \sum_{i \in \{L,H\}} \rho_i \left(\theta_i q_i - c_p \tau_i\right) \\ \text{s.t.} \quad & \rho_L q_L + \rho_H q_H = Q, \\ & \theta_L q_L - c_d \tau_L \geq \theta_L q_H - c_d' \tau_H, \\ & \theta_i q_i - c_d \tau_i \geq 0, \quad i \in \{L,H\}, \\ & 0 \leq q_i \leq 1, \quad 0 \leq \tau_i \leq 1, \quad i \in \{L,H\}. \end{aligned} \tag{EC.1}$$

For the auxiliary supplier-financed problem, define

$$\begin{aligned} B^{\mathrm{UB}}(Q) \equiv \max_{q_L, q_H, \tau_L, \tau_H, \gamma_L, \gamma_H} \quad & \sum_{i \in \{L,H\}} \rho_i \left(\theta_i q_i - c_p \tau_i - \gamma_i\right) \\ \text{s.t.} \quad & \rho_L q_L + \rho_H q_H = Q, \\ & \theta_L q_L - c_d \tau_L + \gamma_L \geq \theta_L q_H - c_d' \tau_H + \gamma_H, \\ & \theta_i q_i - c_d \tau_i + \gamma_i \geq 0, \quad i \in \{L,H\}, \\ & 0 \leq q_i \leq 1, \quad 0 \leq \tau_i \leq 1, \quad \gamma_i \geq 0, \quad i \in \{L,H\}. \end{aligned} \tag{EC.2}$$

Thus, $B^{\mathrm{LB}}(Q)$ and $B^{\mathrm{UB}}(Q)$ are the value functions of the conditional field-evaluator screening problems at aggregate authorization probability Q. They measure the maximal authorization value, net of administrator review costs and, in the auxiliary supplier-financed problem, net of intended supplier-to-evaluator payments. The supplier's virtual fulfillment cost is subtracted separately below through the term $\eta(c)Q(c)$. Because the objective and constraints are linear, convex combinations of policies feasible at Q_1 and Q_2 are feasible at $\lambda Q_1 + (1-\lambda)Q_2$ and deliver the corresponding convex combination of objective values. Hence B^{LB} and B^{UB} are concave on $[0,1]$.

By definition of η, we have $w(c;\theta) = \theta - \eta(c)$. After normalizing the highest-cost supplier's information rent to zero and substituting the envelope condition, both problems can be written in the generic form

$$V^k \equiv \max_{Q(\cdot) \text{ non-increasing}} \int_{c_L}^{c_H} \left[B^k(Q(c)) - \eta(c)Q(c)\right] \mathrm{d}F(c), \quad k \in \{\mathrm{LB}, \mathrm{UB}\}, \tag{EC.3}$$

where $V^{\mathrm{LB}} = \underline{U}$ and $V^{\mathrm{UB}} = \bar{U}$. After this normalization, the only remaining constraint that links supplier cost types is the monotonicity of $Q(c)$; all work-order classification, documentation, and field-evaluator incentive constraints are embedded in the frontier B^k.

To state the ironing result, let $z \equiv F(c)$ and let $C(z) \equiv F^{-1}(z)$ denote the cost quantile function. Define the virtual fulfillment cost in quantile space by

$$\tilde{\eta}(z) \equiv C(z) + (1-\alpha)\frac{z}{f(C(z))}, \quad z \in (0,1), \tag{EC.4}$$

with endpoint values defined by one-sided limits when needed. Let

$$R(z) \equiv \int_0^z \tilde{\eta}(s)\,\mathrm{d}s, \tag{EC.5}$$

and let R^I denote the lower convex envelope of R on $[0,1]$, i.e., the greatest convex minorant of R satisfying

$$R^I(0) = R(0), \qquad R^I(1) = R(1).$$

The ironed virtual fulfillment cost is any nondecreasing measurable selection from the subdifferential of R^I:

$$\eta^I(z) \in \partial R^I(z), \quad z \in (0,1). \tag{EC.6}$$

Equivalently, $\eta^I(z) = \frac{\mathrm{d}}{\mathrm{d}z}R^I(z)$ wherever R^I is differentiable, with arbitrary subgradient values at nondifferentiability points and endpoint values chosen by one-sided limits when needed. These choices at nondifferentiability points do not affect the integrals below. A maximal open interval (a,b) is an ironing interval if $R(z) > R^I(z)$ for all $z \in (a,b)$. On every ironing interval, R^I is affine and $\eta^I(z)$ is constant.

The next proposition characterizes optimal aggregate authorization schedules and the associated mechanism components for the no-side-payment benchmark and the auxiliary supplier-financed problem, without imposing Assumption 1.

PROPOSITION EC.1. *Under the absolute-continuity conditions maintained above, for each $k \in \{\mathrm{LB}, \mathrm{UB}\}$, there exists an optimal aggregate authorization schedule Q^k obtained as follows. For each cost quantile $z \in [0,1]$, choose*

$$Q^k(z) \in \operatorname*{arg\,max}_{Q \in [0,1]} \left\{ B^k(Q) - \eta^I(z) Q \right\}, \tag{EC.7}$$

with a measurable selection that is non-increasing in z and constant on every ironing interval. Given Q^k, choose a solution to the corresponding frontier problem (EC.1) *or* (EC.2) *with $Q = Q^k(F(c))$ to obtain the class-specific authorization probabilities, documentation probabilities, and, in the auxiliary supplier-financed problem, intended supplier-to-evaluator payments. The aggregate expected administrator-to-supplier transfer is recovered from the envelope condition after normalizing the highest-cost supplier's net payoff to zero:*

$$m^k(c) = G^k(c) + \int_c^{c_H} Q^k(F(y))\,\mathrm{d}y + cQ^k(F(c)), \tag{EC.8}$$

where $G^{\mathrm{LB}}(c) = 0$ *and*

$$G^{\mathrm{UB}}(c) = \rho_L \gamma_L(c) + \rho_H \gamma_H(c).$$

Proof of Proposition EC.1. Fix $k \in \{\mathrm{LB}, \mathrm{UB}\}$. We first establish the properties of the frontier B^k. For every $Q \in [0,1]$, the feasible sets of the frontier programs (EC.1) and (EC.2) are nonempty: setting $q_L = q_H = Q$, $\tau_L = \tau_H = 0$, and, for B^{UB}, $\gamma_L = \gamma_H = 0$, satisfies all constraints. The feasible set for B^{LB} is compact. The feasible set for B^{UB} is unbounded as written because the γ_i's have no explicit upper bound. However, for any fixed $(q_L, q_H, \tau_L, \tau_H)$, the objective is weakly improved by choosing the smallest nonnegative (γ_L, γ_H) satisfying the field-evaluator incentive and participation constraints. Since $(q_L, q_H, \tau_L, \tau_H) \in [0,1]^4$, these minimum required payments are uniformly bounded. Hence, without changing the value of $B^{\mathrm{UB}}(Q)$, we may impose a sufficiently large upper bound on the γ_i's. With this restriction, the feasible set is compact and the optimum defining $B^{\mathrm{UB}}(Q)$ is attained.

We next show that each frontier is concave. Fix $Q_1, Q_2 \in [0,1]$ and $\lambda \in [0,1]$, and let $Q_\lambda = \lambda Q_1 + (1-\lambda) Q_2$. For $j = 1,2$, let x^j be an optimal solution to the program defining $B^k(Q_j)$, where $x^j = (q_L^j, q_H^j, \tau_L^j, \tau_H^j)$ for $k = \mathrm{LB}$ and $x^j = (q_L^j, q_H^j, \tau_L^j, \tau_H^j, \gamma_L^j, \gamma_H^j)$ for $k = \mathrm{UB}$. Because all constraints in (EC.1) and (EC.2) are affine and the bound constraints define convex sets, the convex combination $x^\lambda = \lambda x^1 + (1-\lambda) x^2$ is feasible for the same frontier program at aggregate authorization probability Q_λ. Since the objective is linear in the decision variables, the objective value of x^λ equals $\lambda B^k(Q_1) + (1-\lambda) B^k(Q_2)$. Therefore,

$$B^k(Q_\lambda) \geq \lambda B^k(Q_1) + (1-\lambda) B^k(Q_2),$$

which proves that B^k is concave on $[0,1]$.

Next, consider the reduced formulations. For the no-side-payment benchmark, after normalizing the highest-cost supplier's information rent to zero, the supplier envelope characterization implies that incentive compatibility is equivalent to the monotonicity of

$$Q(c) = \rho_L q_L(c) + \rho_H q_H(c),$$

with the aggregate transfer recovered from the supplier envelope formula

$$m(c) = \int_c^{c_H} Q(y)\,\mathrm{d}y + cQ(c).$$

Substituting the envelope condition into the administrator's objective gives

$$\int_{c_L}^{c_H} \sum_{i \in \{L,H\}} \rho_i \left[w(c;\theta_i) q_i(c) - c_p \tau_i(c) \right] \mathrm{d}F(c).$$

For any fixed c, the contribution of $(q_L, q_H, \tau_L, \tau_H)$ is bounded above by $B^{\mathrm{LB}}(Q(c)) - \eta(c)Q(c)$, and this bound is attained by choosing a solution to (EC.1). Hence, the no-side-payment benchmark is equivalent to (EC.3) with $k = \mathrm{LB}$.

The same argument applies to the auxiliary supplier-financed problem after applying the envelope condition to the supplier's net transfer $m(c) - G(c)$, where

$$G(c) = \rho_L \gamma_L(c) + \rho_H \gamma_H(c).$$

After normalizing the highest-cost supplier's net payoff to zero, the corresponding administrator-to-supplier transfer is

$$m(c) = G(c) + \int_c^{c_H} Q(y)\,\mathrm{d}y + cQ(c).$$

Substituting out m gives

$$\int_{c_L}^{c_H} \sum_{i \in \{L,H\}} \rho_i \left[w(c;\theta_i) q_i(c) - c_p \tau_i(c) - \gamma_i(c) \right] \mathrm{d}F(c),$$

so the auxiliary supplier-financed problem is equivalent to (EC.3) with $k = \mathrm{UB}$. Standard measurable-selection arguments allow the frontier maximizers to be selected measurably as functions of $Q(c)$.

We now solve the generic problem. Change variables from cost c to quantile $z = F(c)$, and write $Q(z)$ for $Q(C(z))$. The problem becomes

$$J^k(Q) \equiv \int_0^1 \left[B^k(Q(z)) - \tilde{\eta}(z) Q(z) \right] \mathrm{d}z, \tag{EC.9}$$

subject to $Q(z)$ being non-increasing. Let

$$D(z) \equiv R(z) - R^I(z).$$

Then $D(z) \geq 0$, $D(0) = D(1) = 0$, and the ironing intervals are the connected components of $\{z \in (0,1) : D(z) > 0\}$. Moreover, since $R'(z) = \tilde{\eta}(z)$ and $(R^I)'(z) = \eta^I(z)$ almost everywhere,

$$D'(z) = \tilde{\eta}(z) - \eta^I(z)$$

almost everywhere. Hence, for any non-increasing Q, viewed as a function of bounded variation,

$$\int_0^1 \tilde{\eta}(z) Q(z)\,\mathrm{d}z = \int_0^1 \eta^I(z) Q(z)\,\mathrm{d}z + \int_0^1 Q(z)\,\mathrm{d}D(z).$$

By Stieltjes integration by parts,

$$\int_0^1 Q(z)\,\mathrm{d}D(z) = [D(z)Q(z)]_0^1 - \int_0^1 D(z)\,\mathrm{d}Q(z).$$

The boundary term is zero because $D(0) = D(1) = 0$. Therefore,

$$\int_0^1 \tilde{\eta}(z)Q(z)\,\mathrm{d}z = \int_0^1 \eta^I(z)Q(z)\,\mathrm{d}z - \int_0^1 D(z)\,\mathrm{d}Q(z). \tag{EC.10}$$

Since $\mathrm{d}Q$ is a non-positive Stieltjes measure and $D(z) \geq 0$, the last term on the right-hand side of (EC.10) is nonnegative. Therefore, for every feasible Q,

$$J^k(Q) \leq \int_0^1 \left[B^k(Q(z)) - \eta^I(z)Q(z)\right] \mathrm{d}z. \tag{EC.11}$$

Moreover, the inequality is tight whenever Q is constant on every ironing interval, because then all variation in Q occurs only where $D(z) = 0$, so $\int_0^1 D(z)\,\mathrm{d}Q(z) = 0$.

For each z, let $Q^k(z)$ solve the pointwise ironed problem (EC.7). Since R^I is convex, $\eta^I(z)$ is nondecreasing in z. Because B^k is finite and concave on the compact interval $[0,1]$, the pointwise argmax correspondence is nonempty and compact-valued. A measurable monotone selection can be obtained, for example, by selecting the largest maximizer at each z. To see the monotonicity, take $z_1 < z_2$, write $\eta_j \equiv \eta^I(z_j)$, and choose $Q_j \in \arg\max_Q\{B^k(Q) - \eta_j Q\}$ for $j = 1,2$, with $\eta_1 \leq \eta_2$. Pointwise optimality implies

$$\begin{aligned} B^k(Q_1) - \eta_1 Q_1 &\geq B^k(Q_2) - \eta_1 Q_2, \\ B^k(Q_2) - \eta_2 Q_2 &\geq B^k(Q_1) - \eta_2 Q_1. \end{aligned}$$

Adding these inequalities yields

$$(\eta_2 - \eta_1)(Q_1 - Q_2) \geq 0.$$

Thus any maximizer at the lower virtual cost is weakly larger than any maximizer at the higher virtual cost whenever $\eta_2 > \eta_1$. If $\eta_1 = \eta_2$, the pointwise objective is the same at both quantiles, so the selection can be chosen constant. In particular, Q^k can be chosen to be constant on every ironing interval.

For this choice of Q^k, inequality (EC.11) is tight. Pointwise optimality gives, for every non-increasing feasible Q,

$$J^k(Q) \leq \int_0^1 \max_{\tilde{Q}\in[0,1]} \left\{B^k(\tilde{Q}) - \eta^I(z)\tilde{Q}\right\} \mathrm{d}z = J^k(Q^k).$$

Thus Q^k is optimal for the generic problem. The class-specific authorization probabilities, documentation probabilities, and intended supplier-to-evaluator payments are obtained from the corresponding frontier problem at $Q = Q^k(F(c))$. The transfer formula (EC.8) follows from the same envelope condition used in the no-side-payment benchmark and the auxiliary supplier-financed problem, with $G^{\mathrm{LB}}(c) = 0$ and $G^{\mathrm{UB}}(c) = \rho_L\gamma_L(c) + \rho_H\gamma_H(c)$.

Finally, if Assumption 1 holds, then $\tilde{\eta}(z)$ is nondecreasing, so R is convex. Hence, $R^I = R$, $\eta^I = \tilde{\eta}$, and there are no ironing intervals. The characterization therefore collapses to the pointwise characterization used in the regular case. $\square$

Proposition EC.1 shows that Assumption 1 is a regularity condition for solving the no-side-payment benchmark and the auxiliary supplier-financed problem by pointwise maximization. When the fulfillment-cost distribution is regular, $\tilde{\eta}(z)$ is nondecreasing, R is convex, and hence $R^I = R$ and $\eta^I = \tilde{\eta}$. The ironed characterization then reduces to the regular-case pointwise mechanisms characterized earlier. When the distribution is irregular, the administrator pools supplier cost types over ironing intervals by assigning them the same aggregate authorization probability. Conditional on that aggregate authorization probability, the work-order screening instruments are chosen from the same frontier B^k as in the regular case. The proposition characterizes V^{LB} and V^{UB}; applying an ironed auxiliary solution to the original side-payment game requires the corresponding class-contingent transfer-splitting implementation step, as in the regular case.

EC.3. Proofs of Section 4

Proof of Proposition 1. Let

$$\underline{\mathcal{M}} \equiv \{q_L(c), q_H(c), m_L(c), m_H(c), \tau_L(c), \tau_H(c)\}$$

denote a feasible solution to the optimization problem (26). Construct a mechanism for the original problem by setting

$$\mathcal{M} \equiv \{q_L(c), q_H(c), m_L(c), m_H(c), \tau_L(c), \tau_H(c), \gamma_L(c), \gamma_H(c), \phi_s(c), \phi_e(c)\},$$

where

$$\gamma_L(c) = \gamma_H(c) = 0, \quad \phi_s(c) = c, \quad \phi_e(c) = L, \quad \forall c \in [c_L, c_H].$$

By the constraints defining $\underline{\Omega}$, the constructed mechanism $\mathcal{M}$ is feasible for the original optimization problem (18). Along the equilibrium path, $\mathcal{M}$ coincides with $\underline{\mathcal{M}}$ and involves no side payments, so the two mechanisms generate the same objective value. Hence, every feasible no-side-payment mechanism is feasible for the original problem, implying $\underline{U} \leq U^*$. $\square$

Proof of Lemma 1. Consider any mechanism

$$\tilde{\mathcal{M}} \equiv \{\tilde{q}_i(c), \tilde{m}_i(c), \tilde{\tau}_i(c)\}_{i \in \{L,H\}, c \in [c_L, c_H]}$$

that satisfies all constraints of the no-side-payment benchmark except (24). Define

$$\hat{\mathcal{M}} \equiv \{\hat{q}_i(c), \hat{m}_i(c), \hat{\tau}_i(c)\}_{i \in \{L,H\}, c \in [c_L, c_H]},$$

where

$$\begin{aligned}
&\hat{q}_i(c) = \tilde{q}_i(c), && \forall (c,i) \in [c_L, c_H] \times \{L, H\}, \\
&\hat{\tau}_i(c) = \tilde{\tau}_i(c), && \forall (c,i) \in [c_L, c_H] \times \{L, H\}, \\
&\hat{m}_L(c) = \tilde{m}_L(c) + \frac{\rho_H}{\rho_L} \tilde{m}_H(c), && \forall c \in [c_L, c_H], \\
&\hat{m}_H(c) = 0, && \forall c \in [c_L, c_H].
\end{aligned}$$

The supplier's payoff function defined in (22) is the same under the two aforementioned mechanisms, i.e.,

$$\sum_{i\in\{L,H\}} \rho_i \left[\tilde{m}_i(\hat{c}) - c\,\tilde{q}_i(\hat{c})\right] = \sum_{i\in\{L,H\}} \rho_i \left[\hat{m}_i(\hat{c}) - c\,\hat{q}_i(\hat{c})\right], \quad \forall (c,\hat{c}) \in [c_L, c_H]^2,$$

satisfying its incentive constraints automatically. Since the administrator's monetary transfers to the supplier do not affect the field evaluator's incentives, it remains weakly optimal for the field evaluator to report its type truthfully by construction. It remains to verify constraint (24), which follows from the fact that, under $\hat{\mathcal{M}}$, the right-hand side of (24) is negative, while the nonnegativity of the left-hand side implies the desired inequality.

It is straightforward to verify that all remaining constraints are satisfied under $\hat{\mathcal{M}}$, thereby ensuring its feasibility. Since the administrator's expected payoff remains the same, this completes the proof. □

The following lemma is standard in static mechanism design, and we state it here for completeness.

LEMMA EC.1. *Let*

$$Q(c) \equiv \sum_{i\in\{L,H\}} \rho_i q_i(c).$$

When the supplier's payoff functions are given by (28) *and* (29), *the incentive compatibility constraint* (7) *holds if and only if*

$$Q(c) \text{ non-increasing, and} \tag{EC.12}$$

$$m(c) = u_s(c_H) + \int_c^{c_H} Q(y)\,\mathrm{d}y + c\,Q(c), \quad \forall c \in [c_L, c_H]. \tag{EC.13}$$

Moreover, without loss of optimality, we can focus on mechanisms satisfying $u_s(c_H) = 0$.

Proof of Lemma EC.1. See Proposition 5.2 on page 66 of Krishna (2009). □

PROPOSITION EC.2. *The optimization problem* (26) *is equivalent to the following:*

$$\max_{\{q_L,q_H,\tau_L,\tau_H\}\in\underline{\Omega}_{\mathrm{EC}}} \int_{c_L}^{c_H} \sum_{i\in\{L,H\}} \rho_i \left[\left(\theta_i - c - (1-\alpha)\frac{F(c)}{f(c)}\right) q_i(c) - c_p\tau_i(c)\right] \mathrm{d}F(c), \tag{EC.14}$$

where $\underline{\Omega}_{\mathrm{EC}}$ is determined by (14), (15), (19), (20), *and* (EC.12).

Proof of Proposition EC.2. The objective function follows from substituting m using (EC.13), i.e.,

$$\begin{aligned}
&\int_{c_L}^{c_H} \left\{-(1-\alpha)m(c) + \sum_{i\in\{L,H\}} \rho_i \left[\theta_i q_i(c) - c_p\tau_i(c) - \alpha c q_i(c)\right]\right\} \mathrm{d}F(c) \\
&= \int_{c_L}^{c_H} \left\{-(1-\alpha)\int_c^{c_H} Q(y)\,\mathrm{d}y - cQ(c) + \sum_{i\in\{L,H\}} \rho_i \left[\theta_i q_i(c) - c_p\tau_i(c)\right]\right\} \mathrm{d}F(c) \\
&= \int_{c_L}^{c_H} \sum_{i\in\{L,H\}} \rho_i \left[\left(\theta_i - c - (1-\alpha)\frac{F(c)}{f(c)}\right) q_i(c) - c_p\tau_i(c)\right] \mathrm{d}F(c).
\end{aligned}$$

Invoking Lemma EC.1 completes the proof. □

We first simplify the pointwise problem underlying the characterization results. Fix c and write $w_i(c) \equiv w(c;\theta_i)$ for $i \in \{L, H\}$. In any pointwise optimum, $\tau_L(c) = 0$: lowering $\tau_L(c)$ improves the objective and relaxes both the field-evaluator incentive constraint (19) and participation constraint (20). We can also restrict attention to $q_H(c) \geq q_L(c)$. If $q_H(c) < q_L(c)$ and $w_H(c) \geq 0$, raising $q_H(c)$ to $q_L(c)$ weakly improves the objective and preserves feasibility. If $q_H(c) < q_L(c)$ and $w_H(c) < 0$, then $w_L(c) < 0$, so lowering $q_L(c)$ to $q_H(c)$ weakly improves the objective and preserves feasibility. Hence, with $q_H(c) \geq q_L(c)$, the field-evaluator incentive constraint (19) requires

$$\tau_H(c) \geq \frac{\theta_L}{c_d'}[q_H(c) - q_L(c)].$$

Because documentation is costly, this constraint binds whenever $q_H(c) > q_L(c)$, so

$$\tau_H(c) = \frac{\theta_L}{c_d'}[q_H(c) - q_L(c)].$$

The feasibility constraints (14)–(15) then imply

$$q_H(c) - q_L(c) \leq d, \quad \text{where } d \equiv \min\left\{1, \frac{c_d'}{\theta_L}\right\}.$$

The participation constraints (20) are automatically satisfied because $c_d' \geq c_d$ and $\theta_H \geq \theta_L$. Suppressing the argument c, the pointwise problem is therefore equivalent to

$$\max_{0 \leq q_L \leq q_H \leq 1,\ q_H - q_L \leq d} a(c)q_L + b(c)q_H, \tag{EC.15}$$

where

$$a(c) \equiv \rho_L w(c;\theta_L) + \frac{c_p \rho_H \theta_L}{c_d'},$$
$$b(c) \equiv \rho_H \left[w(c;\theta_H) - \frac{c_p \theta_L}{c_d'}\right].$$

Let

$$S(c) \equiv a(c) + b(c) = \rho_L w(c;\theta_L) + \rho_H w(c;\theta_H).$$

Under Assumption 1, the functions $a(c)$, $b(c)$, and $S(c)$ are non-increasing in c. The proofs below compare the extreme points of the pointwise problem (EC.15). In each case, the induced aggregate authorization probability

$$Q(c) = \rho_L q_L(c) + \rho_H q_H(c)$$

is non-increasing, so the supplier monotonicity constraint (EC.12) is satisfied. We use the convention that, if a cutoff-defining set is empty, the corresponding cutoff is c_L, and if the set is the entire support, the corresponding cutoff is c_H; thus boundary cases are handled by allowing intervals to be empty.

Proof of Proposition 2. In this regime, $\theta_L \le c_d'$, so $d = 1$. Since the pointwise feasible set in (EC.15) already imposes $0 \le q_L \le q_H \le 1$, the constraint $q_H - q_L \le d$ imposes no additional restriction. The feasible set in the (q_L, q_H)-plane is therefore the triangle with vertices

$$(0,0), \quad (0,1), \quad (1,1).$$

The corresponding objective values in (EC.15) are 0, $b(c)$, and $S(c) \equiv a(c) + b(c)$, respectively.

By the definitions of $\acute{c}$ and $\check{c}$, $a(c) \ge 0$ on $[c_L, \acute{c})$, while $b(c) \ge 0$ on $[c_L, \check{c})$. The condition

$$\theta_L \le \theta_H - \frac{c_p \theta_L}{c_d' \rho_L}$$

implies $\acute{c} \le \check{c}$. Hence, for $c < \acute{c}$, both $a(c)$ and $b(c)$ are nonnegative, so $S(c) \ge b(c)$ and $S(c) \ge 0$; the vertex $(1,1)$ is optimal. For $\acute{c} \le c < \check{c}$, we have $a(c) \le 0 \le b(c)$, so $b(c) \ge S(c)$ and $b(c) \ge 0$; the vertex $(0,1)$ is optimal. For $c \ge \check{c}$, both $a(c)$ and $b(c)$ are nonpositive, so $(0,0)$ is optimal. Thus the optimal authorization rule is

$$q_L(c) = \mathbb{1}_{c<\acute{c}}, \qquad q_H(c) = \mathbb{1}_{c<\check{c}}.$$

Using the binding form of the field-evaluator incentive constraint (19), the documentation rule is

$$\tau_L(c) = 0, \qquad \tau_H(c) = \frac{\theta_L}{c_d'} \mathbb{1}_{\acute{c} \le c < \check{c}},$$

as presented in (31).

The induced aggregate authorization probability is

$$Q(c) = \begin{cases} 1, & c \in [c_L, \acute{c}), \\ \rho_H, & c \in [\acute{c}, \check{c}), \\ 0, & c \in [\check{c}, c_H], \end{cases}$$

which is non-increasing in c. Applying the envelope formula (EC.13) gives

$$\begin{aligned} m(c) &= \rho_L \acute{c} + \rho_H \check{c}, & c &\in [c_L, \acute{c}), \\ m(c) &= \rho_H \check{c}, & c &\in [\acute{c}, \check{c}), \\ m(c) &= 0, & c &\in [\check{c}, c_H]. \end{aligned}$$

The candidate therefore satisfies the monotonicity constraint (EC.12) and solves the pointwise problem (EC.15) for every c. Hence, it is globally optimal for (26). □

Proof of Proposition 3. In this regime, $\theta_L \ge c_d'$, so $d = c_d'/\theta_L \le 1$. The feasible set of (EC.15) has vertices

$$E_0 = (0,0), \quad E_1 = (1,1), \quad E_2 = (1-d, 1), \quad E_3 = (0, d),$$

with objective values

$$V_0(c) = 0, \quad V_1(c) = S(c), \quad V_2(c) = (1-d)a(c) + b(c), \quad V_3(c) = db(c),$$

respectively. The condition

$$\theta_L \le \theta_H - \frac{c_p \theta_L}{c_d' \rho_L}$$

implies the ordering $\acute{c} \le \check{c} \le \grave{c}$, where $\acute{c}$, $\check{c}$, and $\grave{c}$ are the thresholds at which $a(c)$, $S(c)$, and $b(c)$ cross zero, respectively.

For $c < \acute{c}$, we have $a(c) \ge 0$, $S(c) \ge 0$, and $b(c) \ge 0$. Thus

$$V_1(c) - V_2(c) = da(c) \ge 0, \qquad V_1(c) - V_3(c) = a(c) + (1-d)b(c) \ge 0,$$

and $V_1(c) \ge V_0(c)$. Hence, E_1 is optimal.

For $\acute{c} \le c < \check{c}$, we have $a(c) \le 0$, $S(c) \ge 0$, and $b(c) \ge 0$. Since

$$V_2(c) - V_1(c) = -da(c) \ge 0, \qquad V_2(c) - V_3(c) = (1-d)S(c) \ge 0,$$

and $V_2(c) = S(c) - da(c) \ge 0 = V_0(c)$, the vertex E_2 is optimal.

For $\check{c} \le c < \grave{c}$, we have $a(c) \le 0$, $S(c) \le 0$, and $b(c) \ge 0$. Then

$$V_3(c) - V_2(c) = -(1-d)S(c) \ge 0, \qquad V_3(c) - V_1(c) = db(c) - S(c) \ge 0,$$

and $V_3(c) \ge V_0(c)$. Hence, E_3 is optimal.

Finally, for $c \ge \grave{c}$, both $a(c)$ and $b(c)$ are nonpositive, and hence $S(c) \le 0$. Therefore, $V_1(c) \le 0$, $V_2(c) \le 0$, and $V_3(c) \le 0$, so E_0 is optimal. This gives the authorization and documentation rules in (33).

The aggregate allocation is

$$Q(c) = \begin{cases} 1, & c \in [c_L, \acute{c}), \\ 1 - \rho_L d, & c \in [\acute{c}, \check{c}), \\ \rho_H d, & c \in [\check{c}, \grave{c}), \\ 0, & c \in [\grave{c}, c_H], \end{cases}$$

which is non-increasing because $d \le 1$. Applying (EC.13) yields

$$\begin{aligned} m(c) &= d\left[\rho_L \acute{c} + \left(\frac{1}{d} - 1\right)\check{c} + \rho_H \grave{c}\right], && c \in [c_L, \acute{c}), \\ m(c) &= (1-d)\check{c} + d\rho_H \grave{c}, && c \in [\acute{c}, \check{c}), \\ m(c) &= d\rho_H \grave{c}, && c \in [\check{c}, \grave{c}), \\ m(c) &= 0, && c \in [\grave{c}, c_H]. \end{aligned}$$

Substituting $d = c_d'/\theta_L$ gives (34). Since the candidate solves the pointwise relaxation and satisfies monotonicity, it is globally optimal. □

Proof of Proposition 4. Let $\Delta \equiv \theta_H - \theta_L$. The condition in the proposition is equivalent to

$$\rho_L \Delta \leq \frac{c_p \theta_L}{c_d'}.$$

Consider the pointwise problem (EC.15). Since the problem is linear, it suffices to compare the vertices of its feasible set. Using the value of d defined above, if $d < 1$, the feasible set has vertices

$$E_0 = (0,0), \quad E_1 = (1,1), \quad E_2 = (1-d,1), \quad E_3 = (0,d).$$

If $d = 1$, then $E_2 = E_3 = (0,1)$, so the same notation covers the triangular boundary case by allowing duplicate vertices. The corresponding objective values are

$$V_0(c) = 0, \quad V_1(c) = S(c), \quad V_2(c) = (1-d)a(c) + b(c), \quad V_3(c) = db(c).$$

Using $w(c;\theta_H) - w(c;\theta_L) = \Delta$, we can write

$$w(c;\theta_H) = S(c) + \rho_L \Delta, \qquad w(c;\theta_L) = S(c) - \rho_H \Delta.$$

Therefore,

$$a(c) = \rho_L S(c) + \rho_H \left[\frac{c_p \theta_L}{c_d'} - \rho_L \Delta \right] \geq \rho_L S(c),$$

and

$$b(c) = \rho_H \left[S(c) + \rho_L \Delta - \frac{c_p \theta_L}{c_d'} \right] \leq \rho_H S(c),$$

where both inequalities follow from $\rho_L \Delta \leq c_p \theta_L / c_d'$.

By the definition of $\check{c}$ and Assumption 1, we have

$$S(c) \geq 0 \quad \text{for } c \in [c_L, \check{c}), \qquad S(c) \leq 0 \quad \text{for } c \in [\check{c}, c_H],$$

up to payoff-irrelevant tie cases at the cutoff.

First consider $c \in [c_L, \check{c})$. Then $S(c) \geq 0$, and the inequality $a(c) \geq \rho_L S(c)$ implies $a(c) \geq 0$. Hence

$$V_1(c) - V_2(c) = S(c) - [(1-d)a(c) + b(c)] = da(c) \geq 0.$$

Also, $V_1(c) = S(c) \geq 0 = V_0(c)$. Finally, since $b(c) \leq \rho_H S(c)$,

$$V_1(c) - V_3(c) = S(c) - db(c) \geq S(c)(1 - d\rho_H) \geq 0,$$

because $d \leq 1$ and $\rho_H \leq 1$. Therefore, $E_1 = (1,1)$ is optimal for $c \in [c_L, \check{c})$.

Now consider $c \in [\check{c}, c_H]$. Then $S(c) \leq 0$. The inequality $b(c) \leq \rho_H S(c)$ implies $b(c) \leq 0$. Hence

$$V_1(c) = S(c) \leq 0, \qquad V_3(c) = db(c) \leq 0.$$

Moreover,

$$V_2(c) = (1-d)a(c) + b(c) = (1-d)S(c) + db(c) \leq 0.$$

Thus $E_0 = (0,0)$ is optimal for $c \in [\check{c}, c_H]$.

The pointwise optimum is therefore pooling:

$$q_L(c) = q_H(c) = \mathbb{1}_{c_L \leq c < \check{c}}.$$

Since both work-order classes are authorized with the same probability, the field-evaluator incentive constraint (19) holds without documentation. Because documentation is costly, the optimal documentation rule is

$$\tau_L(c) = \tau_H(c) = 0.$$

The induced aggregate authorization probability is

$$Q(c) = \rho_L q_L(c) + \rho_H q_H(c) = \mathbb{1}_{c_L \leq c < \check{c}},$$

which is non-increasing. Applying the envelope formula (EC.13) gives

$$m(c) = \check{c} \cdot \mathbb{1}_{c_L \leq c < \check{c}}.$$

Thus the candidate satisfies the supplier monotonicity constraint (EC.12), solves the pointwise problem (EC.15) for every c, and is globally optimal for the benchmark problem (26). □

EC.4. Proofs of Section 5

Proof of Proposition 5. The auxiliary problem (41) relaxes the original side-payment problem (18): it allows the administrator to specify the intended supplier-financed payments γ_i directly, subject to (42). In the original problem, by contrast, the administrator specifies only the mechanism (q_i, m_i, τ_i), and any side payments to field evaluators must be chosen by the supplier in equilibrium. Therefore, any equilibrium outcome attainable in the original problem, together with its equilibrium side-payment schedule, is feasible for the auxiliary problem and yields the same objective value. Maximizing over this larger feasible set gives $\bar{U} \geq U^*$. □

Proof of Lemma 2. Let

$$\tilde{\mathcal{M}} \equiv \{\tilde{q}_L, \tilde{q}_H, \tilde{m}_L, \tilde{m}_H, \tilde{\tau}_L, \tilde{\tau}_H, \tilde{\gamma}_L, \tilde{\gamma}_H, \tilde{\phi}_s, \tilde{\phi}_e\}$$

denote any feasible mechanism for the auxiliary problem (41). We construct a direct mechanism that asks the supplier to report its fulfillment cost and asks the field evaluator to report the work-order class. For each supplier type c, let

$$a(c) \equiv \tilde{\phi}_e(\tilde{\phi}_s(c))$$

denote the on-path report of a class-L field evaluator under $\tilde{\mathcal{M}}$. Define the direct mechanism by

$$
\begin{aligned}
q_L(c) &= \tilde{q}_{a(c)}(\tilde{\phi}_s(c)), & q_H(c) &= \tilde{q}_H(\tilde{\phi}_s(c)),\\
m_L(c) &= \tilde{m}_{a(c)}(\tilde{\phi}_s(c)), & m_H(c) &= \tilde{m}_H(\tilde{\phi}_s(c)),\\
\tau_L(c) &= \tilde{\tau}_{a(c)}(\tilde{\phi}_s(c)), & \tau_H(c) &= \tilde{\tau}_H(\tilde{\phi}_s(c)),\\
\gamma_L(c) &= \tilde{\gamma}_{a(c)}(\tilde{\phi}_s(c)), & \gamma_H(c) &= \tilde{\gamma}_H(\tilde{\phi}_s(c)),\\
\phi_s(c) &= c, & \phi_e(c) &= L.
\end{aligned}
$$

The direct mechanism preserves the class-H outcome induced by $\tilde{\mathcal{M}}$ after the supplier's equilibrium report. For a class-L work order, truthful report L is assigned the outcome that $\tilde{\mathcal{M}}$ assigned to the on-path report $a(c)$. If $a(c) = L$, the class-L field evaluator's incentive constraint follows from the on-path optimality of $a(c)$ in $\tilde{\mathcal{M}}$:

$$\theta_L q_L(c) - c_d \tau_L(c) + \gamma_L(c) \geq \theta_L q_H(c) - c_d' \tau_H(c) + \gamma_H(c).$$

If $a(c) = H$, then by construction $q_L(c) = q_H(c)$, $\tau_L(c) = \tau_H(c)$, and $\gamma_L(c) = \gamma_H(c)$. The payoff from truthful report L then exceeds the payoff from report H by

$$(c_d' - c_d)\tau_H(c) \geq 0.$$

Thus the class-L field evaluator's incentive constraint holds. Moreover, the class-L payoff is unchanged when $a(c) = L$ and weakly higher when $a(c) = H$, while the class-H payoff is preserved by construction.

Now consider the supplier. For any true cost c and reported cost $\hat{c}$, the supplier's payoff from reporting c in the direct mechanism equals its payoff in $\tilde{\mathcal{M}}$ when following the reporting strategy $\tilde{\phi}_s(c)$. Since $\tilde{\phi}_s(c)$ is optimal in $\tilde{\mathcal{M}}$,

$$
\begin{aligned}
&\sum_{i\in\{L,H\}} \rho_i \left[m_i(c) - c q_i(c) - \gamma_i(c)\right]\\
&= \rho_L \left[\tilde{m}_{a(c)}(\tilde{\phi}_s(c)) - c\tilde{q}_{a(c)}(\tilde{\phi}_s(c)) - \tilde{\gamma}_{a(c)}(\tilde{\phi}_s(c))\right] + \rho_H \left[\tilde{m}_H(\tilde{\phi}_s(c)) - c\tilde{q}_H(\tilde{\phi}_s(c)) - \tilde{\gamma}_H(\tilde{\phi}_s(c))\right]\\
&\geq \rho_L \left[\tilde{m}_{a(\hat{c})}(\tilde{\phi}_s(\hat{c})) - c\tilde{q}_{a(\hat{c})}(\tilde{\phi}_s(\hat{c})) - \tilde{\gamma}_{a(\hat{c})}(\tilde{\phi}_s(\hat{c}))\right] + \rho_H \left[\tilde{m}_H(\tilde{\phi}_s(\hat{c})) - c\tilde{q}_H(\tilde{\phi}_s(\hat{c})) - \tilde{\gamma}_H(\tilde{\phi}_s(\hat{c}))\right]\\
&= \sum_{i\in\{L,H\}} \rho_i \left[m_i(\hat{c}) - c q_i(\hat{c}) - \gamma_i(\hat{c})\right].
\end{aligned}
$$

Thus supplier incentive compatibility holds. Supplier participation is preserved because the supplier's interim payoff is preserved.

By construction, the direct mechanism reproduces the on-path authorization rule, documentation rule, expected net supplier transfer, and intended supplier-financed payment induced by $\tilde{\mathcal{M}}$.

Hence, the administrator's expected payoff and the supplier's interim payoff are preserved. The class-H field evaluator's payoff is preserved, while the class-L field evaluator's payoff is weakly higher, so the participation constraints remain satisfied. Therefore, the constructed direct mechanism is feasible and establishes the direct-representation result. □

PROPOSITION EC.3. *The optimization problem* (52) *is equivalent to the following:*

$$\max_{\{q_L,q_H,\tau_L,\tau_H\,\gamma_L,\gamma_H\}\in\bar{\Omega}_{EC}} \int_{c_L}^{c_H} \sum_{i\in\{L,H\}} \rho_i \, [w(c;\theta_i)\, q_i(c) - c_p\, \tau_i(c) - \gamma_i(c)]\, \mathrm{d}F(c), \tag{EC.16}$$

where $\bar{\Omega}_{\mathrm{EC}}$ *is determined by the feasibility constraints* (14)–(15)*, the nonnegativity constraints* (42)*, the field-evaluator incentive constraints* (48)–(49)*, and the monotonicity constraint* (EC.12)*.*

Proof of Proposition EC.3. The proof proceeds along the same lines as those of Lemma EC.1 and Proposition EC.2, and is omitted for brevity.

We use the same pointwise notation as in the no-side-payment benchmark: fix c, suppress the argument c, and write $w_i \equiv w(c;\theta_i)$ for $i \in \{L,H\}$. Since $w_H \geq w_L$, we may restrict attention to $q_H \geq q_L$. If $q_H < q_L$, then either raising q_H to q_L when $w_H \geq 0$, or lowering q_L to q_H when $w_H < 0$, weakly improves the pointwise objective. Also, $\tau_L = 0$ in any optimum, because lowering τ_L improves the objective and relaxes the field-evaluator incentive and participation constraints (48)–(49).

Let

$$x \equiv \theta_L(q_H - q_L) \geq 0$$

denote the gross gain to a class-L field evaluator from inflating the work order as class H. With $\tau_L = 0$, the class-L incentive constraint (48) becomes

$$c_d'\tau_H + \gamma_L - \gamma_H \geq x.$$

We first show that $\gamma_H = 0$ without loss of optimality. Let

$$x' \equiv c_d'\tau_H$$

denote the portion of the incentive gap closed through documentation. Since $\tau_H \leq 1$, we have $x' \leq c_d'$. Moreover, it is never optimal to use documentation to close more than the gap x, because documentation is costly and any excess documentation can be reduced without weakening the class-L incentive constraint. Hence, we can restrict attention to

$$0 \leq x' \leq \min\{x, c_d'\}.$$

For any such x', class-H participation is satisfied even with $\gamma_H = 0$, which follows from

$$x' \leq x = \theta_L(q_H - q_L) \leq \theta_L q_H$$

and therefore

$$\theta_H q_H - c_d \tau_H = \theta_H q_H - c_d \frac{x'}{c_d'} \geq \theta_H q_H - c_d \frac{\theta_L q_H}{c_d'} \geq 0,$$

where the last inequality follows from $\theta_H \geq \theta_L$ and $c_d \leq c_d'$. Thus γ_H is not needed for class-H participation. Since γ_H is costly and tightens the class-L incentive constraint, we set $\gamma_H = 0$.

Given an incentive gap x, define $\Psi(x)$ as the least expected cost of deterring class-L inflation:

$$\Psi(x) \equiv \min_{\tau_H, \gamma_L} \left\{ c_p \rho_H \tau_H + \rho_L \gamma_L : c_d' \tau_H + \gamma_L \geq x,\ 0 \leq \tau_H \leq 1,\ \gamma_L \geq 0 \right\}. \tag{EC.17}$$

Here, documentation for reported class-H work orders closes $c_d' \tau_H$ units of the class-L incentive gap and costs $c_p \rho_H \tau_H$ in expectation, while the supplier-financed payment γ_L closes the remaining gap and costs $\rho_L \gamma_L$ in expectation. Thus,

$$\Psi(x) = \min_{0 \leq x' \leq \min\{x, c_d'\}} \left\{ \frac{c_p \rho_H}{c_d'} x' + \rho_L (x - x') \right\}.$$

Solving this one-dimensional minimization yields

$$\Psi(x) = \begin{cases} \rho_L x, & \text{if } c_d' \rho_L - c_p \rho_H \leq 0, \\ \dfrac{c_p \rho_H}{c_d'} x, & \text{if } c_d' \rho_L - c_p \rho_H > 0 \text{ and } x \leq c_d', \\ c_p \rho_H + \rho_L (x - c_d'), & \text{if } c_d' \rho_L - c_p \rho_H > 0 \text{ and } x > c_d'. \end{cases} \tag{EC.18}$$

Having shown that $\tau_L = 0$ and $\gamma_H = 0$ are without loss of optimality, fix an authorization pair (q_L, q_H) with $0 \leq q_L \leq q_H \leq 1$. The remaining choice is how to close the class-L incentive gap

$$x = \theta_L (q_H - q_L)$$

at minimum expected cost. By definition of Ψ, this minimum cost is $\Psi(x)$. Therefore, after replacing the documentation and incentive-payment subproblem by its value Ψ, the pointwise problem reduces to

$$\max_{0 \leq q_L \leq q_H \leq 1} \rho_L w_L q_L + \rho_H w_H q_H - \Psi\left(\theta_L (q_H - q_L)\right). \tag{EC.19}$$

The candidate mechanisms below solve (EC.19) for every c. Their induced aggregate authorization probability $Q(c) = \rho_L q_L(c) + \rho_H q_H(c)$ is non-increasing, so the monotonicity constraint (EC.12) is satisfied.

Proof of Proposition 6. We use the pointwise formulation (EC.19). We prove the three parts by comparing the objective values of the vertices.

First consider Part 1. Since $c_d' \rho_L - c_p \rho_H \geq 0$ and $\theta_L \leq c_d'$, we have $x = \theta_L (q_H - q_L) \leq c_d'$ for all feasible (q_L, q_H). Hence, the cost of closing the incentive gap $\Psi(x)$ reduces to the documentation cost $(c_p \rho_H / c_d') x$. The feasible set has vertices

$$E_0 = (0, 0), \qquad E_1 = (1, 1), \qquad E_2 = (0, 1),$$

with values

$$V_0(c) = 0, \qquad V_1(c) = S(c), \qquad V_2(c) = \rho_H \left[w(c;\theta_H) - \frac{c_p \theta_L}{c_d'} \right],$$

where $S(c) = \rho_L w(c;\theta_L) + \rho_H w(c;\theta_H)$. The adjacent differences are

$$V_1(c) - V_2(c) = \rho_L w(c;\theta_L) + \frac{c_p \rho_H \theta_L}{c_d'}, \qquad V_2(c) - V_0(c) = \rho_H \left[w(c;\theta_H) - \frac{c_p \theta_L}{c_d'} \right].$$

These expressions cross zero at the thresholds $\acute{c}$ and $\check{c}$ in Proposition 2. The condition $\theta_H \geq \theta_L + c_p\theta_L/(c_d'\rho_L)$ implies $\acute{c} \leq \check{c}$. Hence, for $c \in [c_L, \acute{c})$, both adjacent differences are nonnegative and E_1 is optimal; for $c \in [\acute{c}, \check{c})$, $V_2(c) \geq V_1(c)$ and $V_2(c) \geq V_0(c)$, so E_2 is optimal; and for $c \in [\check{c}, c_H]$, both $V_1(c)$ and $V_2(c)$ are nonpositive, so E_0 is optimal. Thus the no-side-payment mechanism in Proposition 2, with $\gamma_L = \gamma_H = 0$, is optimal for the auxiliary problem. This proves Part 1.

Next consider Part 2. Suppose $c_d'\rho_L - c_p\rho_H \geq 0$, $\theta_L \geq c_d'$, and

$$\theta_L + \frac{c_p \theta_L}{c_d' \rho_L} \leq \theta_H \leq \theta_L + \frac{\theta_L}{\rho_H}.$$

Let $d \equiv c_d'/\theta_L \leq 1$. Notice that $q_H - q_L = d$ is the kink at which the cost function Ψ changes slope. Hence, the objective is linear on each of the two regions

$$\mathcal{P}^- = \{0 \leq q_L \leq q_H \leq 1,\ q_H - q_L \leq d\}$$

and

$$\mathcal{P}^+ = \{0 \leq q_L \leq q_H \leq 1,\ q_H - q_L \geq d\}.$$

The vertices of $\mathcal{P}^-$ are

$$E_0 = (0,0), \quad E_1 = (1,1), \quad E_2 = (1-d, 1), \quad E_4 = (0, d),$$

and the vertices of $\mathcal{P}^+$ are

$$E_2 = (1-d, 1), \quad E_3 = (0,1), \quad E_4 = (0, d).$$

Thus it is enough to compare the five candidate points

$$E_0 = (0,0), \quad E_1 = (1,1), \quad E_2 = (1-d,1), \quad E_3 = (0,1), \quad E_4 = (0,d),$$

with duplicate vertices allowed when $d = 1$.

The points E_0, E_1, E_2, and E_4 are exactly the candidate points from the no-side-payment pointwise problem in Proposition 3. The only additional candidate is the full-separation point $E_3 = (0,1)$, which is implemented with $\tau_H = 1$ and $\gamma_L = \theta_L - c_d'$. Its value is

$$V_3(c) = \rho_H w(c;\theta_H) - c_p \rho_H - \rho_L(\theta_L - c_d').$$

The adjacent no-side-payment candidates E_2 and E_4 have values

$$V_2(c) = \rho_L(1-d)w(c;\theta_L) + \rho_H w(c;\theta_H) - c_p\rho_H,$$

and

$$V_4(c) = d\rho_H w(c;\theta_H) - c_p\rho_H.$$

A direct comparison gives

$$V_3(c) - V_2(c) = -\rho_L(1-d)\left[w(c;\theta_L) + \theta_L\right],$$

and

$$V_3(c) - V_4(c) = (1-d)\left[\rho_H w(c;\theta_H) - \rho_L\theta_L\right].$$

If $w(c;\theta_L) + \theta_L \geq 0$, then $V_3(c) \leq V_2(c)$. If $w(c;\theta_L) + \theta_L < 0$, then

$$w(c;\theta_H) = w(c;\theta_L) + \theta_H - \theta_L < \theta_H - 2\theta_L,$$

and therefore

$$\rho_H w(c;\theta_H) - \rho_L\theta_L < \rho_H(\theta_H - 2\theta_L) - \rho_L\theta_L = \rho_H(\theta_H - \theta_L) - \theta_L \leq 0,$$

where the last inequality follows from

$$\theta_H \leq \theta_L + \frac{\theta_L}{\rho_H}.$$

Hence, $V_3(c) \leq V_4(c)$. Therefore, the full-separation candidate E_3 is dominated by one of the no-side-payment candidates for each c.

It follows that the auxiliary pointwise problem reduces to the no-side-payment pointwise problem in this parameter region. The interval-by-interval comparison in Proposition 3 therefore applies: the optimal point is E_1 on $[c_L, \acute{c})$, E_2 on $[\acute{c}, \check{c})$, E_4 on $[\check{c}, \grave{c})$, and E_0 on $[\grave{c}, c_H]$. Hence, the mechanism in Proposition 3, with $\gamma_L = \gamma_H = 0$, remains optimal for the auxiliary problem. This proves Part 2.

It remains to prove Part 3. First consider case (i), where $c_d'\rho_L - c_p\rho_H \geq 0$ and $\theta_H \leq \theta_L + \frac{c_p\theta_L}{c_d'\rho_L}$. The no-side-payment proof of Proposition 4 shows that, among the no-side-payment candidate vertices, the pooling vertex $(1,1)$ is optimal whenever $S(c) \geq 0$, and the no-authorization vertex $(0,0)$ is optimal whenever $S(c) < 0$. If $\theta_L \leq c_d'$, then $x = \theta_L(q_H - q_L) \leq c_d'$ for every feasible authorization pair (q_L, q_H). Hence, the auxiliary pointwise problem coincides with the no-side-payment pointwise problem, and the pooling mechanism remains optimal.

Now suppose $\theta_L > c_d'$. Then $d = c_d'/\theta_L < 1$. Since $c_d'\rho_L - c_p\rho_H \geq 0$, the cost function Ψ is piecewise linear, with a possible kink at $q_H - q_L = d$. The kink points $(1-d, 1)$ and $(0, d)$ are already among the no-side-payment candidates considered in Proposition 4; by the no-side-payment pooling proof, they are dominated by either the pooling point $(1,1)$ when $S(c) \geq 0$ or the no-authorization point $(0,0)$ when $S(c) < 0$.

Thus, the only candidate in the pointwise problem not already covered by the no-side-payment comparison is the full-separation point $(0,1)$, which is implemented with $\tau_H = 1$ and $\gamma_L = \theta_L - c_d'$. Since $c_d'\rho_L - c_p\rho_H \geq 0$ implies

$$\theta_L + \frac{c_p\theta_L}{c_d'\rho_L} \leq \theta_L + \frac{\theta_L}{\rho_H},$$

the maintained condition in case (i) implies

$$\theta_H \leq \theta_L + \frac{\theta_L}{\rho_H}.$$

The domination argument in Part 2 then rules out the full-separation point $(0,1)$. Hence the pooling mechanism in Proposition 4, with $\gamma_L = \gamma_H = 0$, remains optimal in case (i).

Now consider case (ii), where $c_d'\rho_L - c_p\rho_H \leq 0$ and $\theta_H \leq \theta_L + \frac{\theta_L}{\rho_H}$. In this case, the cost function is given by $\Psi(x) = \rho_L x$. Thus, the only possible competitor to pooling or no authorization is the full-separation point $(0,1)$, whose value is

$$V_3(c) = \rho_H w(c;\theta_H) - \rho_L\theta_L.$$

When $S(c) \geq 0$, the pooling vertex $(1,1)$ dominates $(0,1)$, because

$$S(c) - V_3(c) = \rho_L\left[w(c;\theta_L) + \theta_L\right] = \rho_L\left[S(c) + \theta_L - \rho_H(\theta_H - \theta_L)\right] \geq 0,$$

where the inequality follows from $\theta_H \leq \theta_L + \theta_L/\rho_H$. Hence $(1,1)$ is optimal whenever $S(c) \geq 0$.

When $S(c) < 0$, the full-separation point has nonpositive value:

$$V_3(c) = \rho_H S(c) + \rho_L\left[\rho_H(\theta_H - \theta_L) - \theta_L\right] \leq 0.$$

Thus the no-authorization vertex $(0,0)$ is optimal whenever $S(c) < 0$. Therefore the pooling mechanism in Proposition 4, with $\gamma_L = \gamma_H = 0$, is optimal in case (ii). This proves Part 3 and completes the proof. □

Proof of Proposition 7. Since $c_d'\rho_L - c_p\rho_H \leq 0$, the objective function is linear, and the feasible set has vertices

$$E_0 = (0,0), \qquad E_1 = (1,1), \qquad E_3 = (0,1),$$

with values

$$V_0(c) = 0, \quad V_1(c) = \rho_L w(c;\theta_L) + \rho_H w(c;\theta_H), \quad V_3(c) = \rho_H w(c;\theta_H) - \rho_L\theta_L.$$

The threshold $\bar{c}$ is the point at which

$$V_1(c) - V_3(c) = \rho_L[w(c;\theta_L) + \theta_L]$$

crosses zero, while $\acute{c}$ is the point at which $V_3(c) - V_0(c)$ crosses zero. The condition $\theta_H \geq \theta_L + \theta_L/\rho_H$ implies $\bar{c} \leq \acute{c}$.

For $c \in [c_L, \bar{c})$, we have $V_1(c) - V_3(c) \geq 0$. Since $c < \acute{c}$, we also have $V_3(c) \geq 0 = V_0(c)$. Hence, $E_1 = (1,1)$ is optimal on $[c_L, \bar{c})$. For $c \in [\bar{c}, \acute{c})$, we have $V_1(c) - V_3(c) \leq 0$ and $V_3(c) \geq 0 = V_0(c)$, so $E_3 = (0,1)$ is optimal. For $c \in [\acute{c}, c_H]$, we have $V_3(c) \leq 0$, and because $c \geq \bar{c}$, also $V_1(c) \leq V_3(c) \leq 0$. Thus $E_0 = (0,0)$ is optimal.

The implementation has $\tau_L = \tau_H = \gamma_H = 0$. On $[\bar{c}, \acute{c})$, the gap is $x = \theta_L$, so $\gamma_L = \theta_L$; outside this interval, the gap is zero and $\gamma_L = 0$. The induced aggregate authorization probability is

$$Q(c) = \begin{cases} 1, & c \in [c_L, \bar{c}), \\ \rho_H, & c \in [\bar{c}, \acute{c}), \\ 0, & c \in [\acute{c}, c_H], \end{cases}$$

which is non-increasing. The aggregate expected payment m can be recovered as follows:

$$\begin{aligned} m(c) &= \rho_L \bar{c} + \rho_H \acute{c}, && c \in [c_L, \bar{c}), \\ m(c) &= \rho_L \theta_L + \rho_H \acute{c}, && c \in [\bar{c}, \acute{c}), \\ m(c) &= 0, && c \in [\acute{c}, c_H]. \end{aligned}$$

The candidate mechanism solves (EC.19) for every c and satisfies the monotonicity constraint (EC.12), it is globally optimal for the auxiliary problem. □

Proof of Proposition 8. Here $c_d' \rho_L - c_p \rho_H \geq 0$ and $c_d' \leq \theta_L$. Let $d = c_d'/\theta_L \leq 1$. The kink in Ψ occurs at $q_H - q_L = d$. The pointwise objective is piecewise linear, so it suffices to compare the vertices of the induced polyhedral subdivision:

$$E_0 = (0,0), \quad E_1 = (1,1), \quad E_2 = (1-d, 1), \quad E_3 = (0,1), \quad E_4 = (0,d),$$

allowing duplicate vertices when $d = 1$. Their objective values are

$$\begin{aligned} V_0(c) &= 0, \\ V_1(c) &= \rho_L w(c;\theta_L) + \rho_H w(c;\theta_H), \\ V_2(c) &= \rho_L (1-d) w(c;\theta_L) + \rho_H w(c;\theta_H) - c_p \rho_H, \\ V_3(c) &= \rho_H w(c;\theta_H) - c_p \rho_H - \rho_L(\theta_L - c_d'), \\ V_4(c) &= d \rho_H w(c;\theta_H) - c_p \rho_H. \end{aligned}$$

The adjacent differences are

$$\begin{aligned}
V_1(c) - V_2(c) &= \frac{1}{\theta_L}\left[c_d'\rho_L w(c;\theta_L) + c_p\rho_H\theta_L\right],\\
V_2(c) - V_3(c) &= \rho_L(1-d)[w(c;\theta_L) + \theta_L],\\
V_3(c) - V_4(c) &= (1-d)[-\rho_L\theta_L + \rho_H w(c;\theta_H)],\\
V_4(c) - V_0(c) &= \frac{\rho_H}{\theta_L}\left[c_d' w(c;\theta_H) - c_p\theta_L\right].
\end{aligned}$$

The four thresholds $\acute{c}$, $\check{c}$, $\grave{c}$, and $\dot{c}$ are the zero-crossings of these adjacent differences, respectively. Assumption 1 gives single crossing, and the conditions in the proposition imply

$$\acute{c} \le \check{c} \le \grave{c} \le \dot{c}.$$

Indeed, $\acute{c} \le \check{c}$ and $\grave{c} \le \dot{c}$ are both equivalent to $c_d'\rho_L - c_p\rho_H \ge 0$, while $\check{c} \le \grave{c}$ is equivalent to $\theta_H \ge \theta_L + \theta_L/\rho_H$.

For $c \in [c_L, \acute{c})$, all four adjacent differences above are nonnegative, so

$$V_1(c) \ge V_2(c) \ge V_3(c) \ge V_4(c) \ge V_0(c).$$

Hence, $E_1 = (1,1)$ is optimal. For $c \in [\acute{c}, \check{c})$, the first adjacent difference is nonpositive and the remaining three are nonnegative, so

$$V_2(c) \ge V_1(c), \qquad V_2(c) \ge V_3(c) \ge V_4(c) \ge V_0(c),$$

and $E_2 = (1-d, 1)$ is optimal. For $c \in [\check{c}, \grave{c})$, the first two adjacent differences are nonpositive and the last two are nonnegative, so

$$V_3(c) \ge V_2(c) \ge V_1(c), \qquad V_3(c) \ge V_4(c) \ge V_0(c),$$

and $E_3 = (0,1)$ is optimal. For $c \in [\grave{c}, \dot{c})$, the first three adjacent differences are nonpositive and the last one is nonnegative, so

$$V_4(c) \ge V_3(c) \ge V_2(c) \ge V_1(c), \qquad V_4(c) \ge V_0(c),$$

and $E_4 = (0,d)$ is optimal. Finally, for $c \in [\dot{c}, c_H]$, all four adjacent differences are nonpositive, so

$$V_0(c) \ge V_4(c) \ge V_3(c) \ge V_2(c) \ge V_1(c),$$

and $E_0 = (0,0)$ is optimal. These interval-by-interval comparisons give the authorization probabilities stated in the proposition.

The values of τ_H and γ_L follow from (EC.18). For E_1, no deterrence is needed. For E_2 and E_4, the gap is $x = c'_d$, so the mechanism sets $\tau_H = 1$ and $\gamma_L = 0$. For E_3, the gap is $x = \theta_L > c'_d$, so the mechanism sets $\tau_H = 1$ and $\gamma_L = \theta_L - c'_d$. Therefore

$$\tau_L(c) = \gamma_H(c) = 0, \quad \tau_H(c) = \mathbb{1}_{c \in [\acute{c}, \dot{c})}, \quad \gamma_L(c) = (\theta_L - c'_d)\mathbb{1}_{c \in [\check{c}, \grave{c})}.$$

The induced aggregate authorization probability is

$$Q(c) = \begin{cases} 1, & c \in [c_L, \acute{c}), \\ 1 - \rho_L d, & c \in [\acute{c}, \check{c}), \\ \rho_H, & c \in [\check{c}, \grave{c}), \\ \rho_H d, & c \in [\grave{c}, \dot{c}), \\ 0, & c \in [\dot{c}, c_H], \end{cases}$$

which is non-increasing. The aggregate expected payment m can be recovered as follows:

$$\begin{aligned}
m(c) &= \acute{c} + \rho_L \left(1 - \frac{c'_d}{\theta_L}\right)(\check{c} - \acute{c}) + \rho_H \left[\grave{c} - \acute{c} + \frac{c'_d}{\theta_L}(\dot{c} - \grave{c})\right], && c \in [c_L, \acute{c}), \\
m(c) &= \left[\rho_L \left(1 - \frac{c'_d}{\theta_L}\right) + \rho_H\right]\check{c} + \rho_H \left[\grave{c} - \check{c} + \frac{c'_d}{\theta_L}(\dot{c} - \grave{c})\right], && c \in [\acute{c}, \check{c}), \\
m(c) &= \rho_L(\theta_L - c'_d) + \rho_H \left[\grave{c} + \frac{c'_d}{\theta_L}(\dot{c} - \grave{c})\right], && c \in [\check{c}, \grave{c}), \\
m(c) &= \frac{c'_d \rho_H}{\theta_L}\dot{c}, && c \in [\grave{c}, \dot{c}), \\
m(c) &= 0, && c \in [\dot{c}, c_H].
\end{aligned}$$

This matches the proposed mechanism. Since it solves (EC.19) for every cost c and satisfies the monotonicity constraint (EC.12), it is globally optimal for the auxiliary problem. □

Proof of Proposition 9. Let $q_i(c)$, $\tau_i(c)$, $\gamma_i(c)$, and $m(c)$ denote one of the optimal mechanisms characterized in Propositions 6–8. In all of these mechanisms,

$$\gamma_H(c) = 0, \quad \forall c \in [c_L, c_H],$$

and

$$\gamma_L(c) = [\theta_L q_H(c) - c'_d \tau_H(c) - \theta_L q_L(c) + c_d \tau_L(c)]^+.$$

Thus $\gamma_L(c)$ is the minimal payment that makes truthful class-L reporting weakly optimal.

We construct a mechanism for the original side-payment game that uses the same authorization and documentation rules $q_i(c)$ and $\tau_i(c)$. If $\rho_L = 0$, there are no class-L work orders, so the construction is straightforward. Suppose $\rho_L > 0$. Define the work-order-class-contingent expected net transfers to the supplier by

$$m^I_H(c) = 0, \qquad m^I_L(c) = \frac{m(c)}{\rho_L}, \quad \forall c \in [c_L, c_H].$$

Then

$$\rho_L m_L^I(c) + \rho_H m_H^I(c) = m(c),$$

so the aggregate expected net transfer to the supplier is preserved on the equilibrium path.

Fix a supplier with fulfillment cost c, and suppose it reports $\hat{c}$. If the supplier induces a class-L field evaluator to report truthfully, the least costly payment is $\gamma_L(\hat{c})$, paid conditional on an L report. Under the proposed transfer split, the supplier's expected payoff from this continuation is

$$\rho_L \left[m_L^I(\hat{c}) - cq_L(\hat{c}) - \gamma_L(\hat{c})\right] + \rho_H \left[m_H^I(\hat{c}) - cq_H(\hat{c})\right] = m(\hat{c}) - c\left[\rho_L q_L(\hat{c}) + \rho_H q_H(\hat{c})\right] - \rho_L \gamma_L(\hat{c}).$$

Because $\gamma_H = 0$, this is exactly the supplier's payoff from reporting $\hat{c}$ in the auxiliary direct problem. Therefore, supplier incentive compatibility in (52) implies

$$m(c) - c\left[\rho_L q_L(c) + \rho_H q_H(c)\right] - \rho_L \gamma_L(c) \geq m(\hat{c}) - c\left[\rho_L q_L(\hat{c}) + \rho_H q_H(\hat{c})\right] - \rho_L \gamma_L(\hat{c}),$$

for all $(c, \hat{c}) \in [c_L, c_H]^2$.

Now consider a deviation in which the supplier reports $\hat{c}$ and induces class-L work orders to be reported as class H. The minimum payment needed for this deviation is $\tilde{\gamma}_H(\hat{c})$, defined in (3). Since all work orders are then reported as class H, the supplier's payoff from this deviation is

$$m_H^I(\hat{c}) - cq_H(\hat{c}) - \tilde{\gamma}_H(\hat{c}) = -cq_H(\hat{c}) - \tilde{\gamma}_H(\hat{c}) \leq 0.$$

Supplier participation in the auxiliary problem implies

$$m(c) - c\left[\rho_L q_L(c) + \rho_H q_H(c)\right] - \rho_L \gamma_L(c) \geq 0.$$

Hence, no supplier type can gain by inducing class-L work orders to be reported as class H. Because any payment above the minimum required to induce a given report only lowers the supplier's payoff, it is sufficient to consider the two minimal payment schemes above: γ_L for truthful class-L reporting and $\tilde{\gamma}_H$ for inflated class-H reporting.

It remains to verify field-evaluator incentives. Given the supplier's payment $\gamma_L(c)$, a class-L field evaluator weakly prefers truthful reporting because

$$\theta_L q_L(c) - c_d \tau_L(c) + \gamma_L(c) \geq \theta_L q_H(c) - c_d' \tau_H(c).$$

A class-H field evaluator cannot misreport by assumption, and its participation constraint is unchanged because $\gamma_H(c) = 0$. Class-L participation is also preserved because the field evaluator receives the same payoff as in the auxiliary mechanism.

Finally, the aggregate expected net transfer to the supplier is preserved by construction, the authorization and documentation rules are unchanged, and the supplier makes the same expected incentive payment as in the auxiliary mechanism. Therefore, the administrator's payoff is the same as in the auxiliary problem. The characterized auxiliary mechanism achieves $\bar{U}$ in the original side-payment problem. Since Proposition 5 establishes $\bar{U} \geq U^*$, we conclude that

$$U^* = \bar{U}.$$

This completes the proof. □

EC.5. Proofs for Section 6

Proof of Proposition 10. The proof follows directly from the mechanisms in Proposition 6. We consider the three cases separately.

First consider the low-θ_L separating regime in Proposition 6(1). If

$$\hat{c}_d = \beta c_d, \qquad \hat{c}'_d = \beta c'_d, \qquad \hat{c}_p = \beta c_p,$$

and the mechanism remains in the same regime, then the thresholds

$$\acute{c} = \sup\left\{c : c'_d \rho_L w(c;\theta_L) + c_p \rho_H \theta_L \geq 0\right\}$$

and

$$\check{c} = \sup\left\{c : c'_d w(c;\theta_H) - c_p \theta_L \geq 0\right\}$$

are unchanged because both defining inequalities are multiplied by β. The documentation probability in the separating interval becomes

$$\hat{\tau}_H(c) = \frac{\theta_L}{\beta c'_d},$$

and the review cost becomes βc_p, so the expected review cost

$$\hat{c}_p \hat{\tau}_H(c) = \beta c_p \frac{\theta_L}{\beta c'_d} = \frac{c_p \theta_L}{c'_d}$$

is unchanged. Similarly, the class-H field evaluator's documentation cost is unchanged:

$$\hat{c}_d \hat{\tau}_H(c) = \beta c_d \frac{\theta_L}{\beta c'_d} = \frac{c_d \theta_L}{c'_d}.$$

The authorization rule, aggregate authorization probability, and supplier transfer are unchanged. Hence all payoffs are unchanged.

Now suppose only the inflated-report documentation cost c'_d varies, holding c_d and c_p fixed. Let $\acute{c}(c'_d)$, $\check{c}(c'_d)$, and $U^*(c'_d)$ denote the corresponding thresholds and optimal objective value.

In the low-θ_L separating regime, the administrator's expected payoff is

$$U^*(c'_d) = \rho_L \int_{c_L}^{\acute{c}(c'_d)} w(c;\theta_L)\,\mathrm{d}F(c) + \rho_H \int_{c_L}^{\check{c}(c'_d)} w(c;\theta_H)\,\mathrm{d}F(c) - \rho_H \frac{c_p \theta_L}{c'_d}\left[F(\check{c}(c'_d)) - F(\acute{c}(c'_d))\right].$$

When the thresholds are interior, differentiating with respect to c_d' gives

$$\begin{aligned}\frac{\mathrm{d}U^*(c_d')}{\mathrm{d}c_d'} &= \left[\rho_L w(\acute{c}(c_d');\theta_L) + \rho_H \frac{c_p\theta_L}{c_d'}\right] f(\acute{c}(c_d'))\frac{\mathrm{d}\acute{c}(c_d')}{\mathrm{d}c_d'} \\ &\quad + \rho_H\left[w(\check{c}(c_d');\theta_H) - \frac{c_p\theta_L}{c_d'}\right] f(\check{c}(c_d'))\frac{\mathrm{d}\check{c}(c_d')}{\mathrm{d}c_d'} \\ &\quad + \rho_H\frac{c_p\theta_L}{(c_d')^2}\left[F(\check{c}(c_d')) - F(\acute{c}(c_d'))\right].\end{aligned}$$

The first two terms vanish by the definitions of the thresholds:

$$c_d'\rho_L w(\acute{c}(c_d');\theta_L) + c_p\rho_H\theta_L = 0, \qquad c_d' w(\check{c}(c_d');\theta_H) - c_p\theta_L = 0.$$

Therefore,

$$\frac{\mathrm{d}U^*(c_d')}{\mathrm{d}c_d'} = \rho_H\frac{c_p\theta_L}{(c_d')^2}\left[F(\check{c}(c_d')) - F(\acute{c}(c_d'))\right] \geq 0.$$

Hence, decreasing c_d' weakly decreases the administrator's expected payoff. The calculation above assumes interior cutoffs. Boundary cutoffs can be handled by the corresponding one-sided derivative; if a cutoff is locally pinned at c_L or c_H, its derivative term simply drops out.

Next consider the intermediate-θ_L separating regime in Proposition 6(2). Under a proportional cost change, the thresholds $\acute{c}, \check{c}, \hat{c}$ remain unchanged, while

$$d(\beta) = \frac{\beta c_d'}{\theta_L}.$$

Using the reduced objective, the optimal objective value can be written as

$$\begin{aligned}U^*(\beta) &= \int_{c_L}^{\acute{c}} \left[\rho_L w(c;\theta_L) + \rho_H w(c;\theta_H)\right] \mathrm{d}F(c) \\ &\quad + \int_{\acute{c}}^{\check{c}} \left[\rho_L\left(1 - \frac{\beta c_d'}{\theta_L}\right) w(c;\theta_L) + \rho_H w(c;\theta_H) - \beta\rho_H c_p\right] \mathrm{d}F(c) \\ &\quad + \int_{\check{c}}^{\hat{c}} \rho_H\left[\frac{\beta c_d'}{\theta_L} w(c;\theta_H) - \beta c_p\right] \mathrm{d}F(c).\end{aligned}$$

Differentiating with respect to β gives

$$\frac{\mathrm{d}U^*(\beta)}{\mathrm{d}\beta} = \int_{\acute{c}}^{\check{c}} \left[-\frac{\rho_L c_d'}{\theta_L} w(c;\theta_L) - \rho_H c_p\right] \mathrm{d}F(c) + \int_{\check{c}}^{\hat{c}} \rho_H\left[\frac{c_d'}{\theta_L} w(c;\theta_H) - c_p\right] \mathrm{d}F(c).$$

On $[\acute{c}, \check{c})$, the definition of $\acute{c}$ and monotonicity of $w(c;\theta_L)$ imply

$$c_d'\rho_L w(c;\theta_L) + c_p\rho_H\theta_L \leq 0,$$

and hence the first integrand is nonnegative. On $[\check{c}, \hat{c})$, the definition of $\hat{c}$ implies

$$c_d' w(c;\theta_H) - c_p\theta_L \geq 0,$$

so the second integrand is also nonnegative. Therefore

$$\frac{\mathrm{d}U^*(\beta)}{\mathrm{d}\beta} \geq 0.$$

Thus, a proportional reduction in documentation-related costs, corresponding to a decrease in β, weakly decreases the administrator's expected payoff in this regime.

The class-L payoff is

$$\mathcal{V}_L(\beta) = \theta_L F(\acute{c}) + (\theta_L - \beta c_d')\left[F(\check{c}) - F(\acute{c})\right],$$

which is decreasing in β. Therefore, decreasing β weakly increases $\mathcal{V}_L$. The class-H field evaluator's expected payoff is

$$\mathcal{V}_H(\beta) = \theta_H F(\acute{c}) + (\theta_H - \beta c_d)\left[F(\check{c}) - F(\acute{c})\right] + \beta\left(\frac{c_d'\theta_H}{\theta_L} - c_d\right)\left[F(\grave{c}) - F(\check{c})\right].$$

Differentiating with respect to β yields

$$\frac{\mathrm{d}\mathcal{V}_H(\beta)}{\mathrm{d}\beta} = -c_d\left[F(\check{c}) - F(\acute{c})\right] + \left(\frac{c_d'\theta_H}{\theta_L} - c_d\right)\left[F(\grave{c}) - F(\check{c})\right].$$

The first term is nonpositive, while the second term is nonnegative since $c_d'\theta_H/\theta_L \geq c_d$. Hence, the sign depends on the relative sizes of the intervals $[\acute{c}, \check{c})$ and $[\check{c}, \grave{c})$, as well as on the filing cost c_d.

To see that either sign can arise, consider a uniform fulfillment-cost distribution on $[0,1]$ with $\alpha = 0$, so $w(c;\theta) = \theta - 2c$. Let

$$\theta_L = 0.4, \quad \theta_H = 0.9, \quad \rho_L = 0.25, \quad \rho_H = 0.75, \quad c_d' = 0.2, \quad c_p = 0.02.$$

These parameters satisfy the conditions of the intermediate separating regime, and the thresholds are

$$\acute{c} = 0.26, \qquad \check{c} = 0.3875, \qquad \grave{c} = 0.43.$$

If $c_d = 0.15$, then

$$\frac{\mathrm{d}\mathcal{V}_H}{\mathrm{d}\beta} = -0.15(0.3875 - 0.26) + \left(\frac{0.2 \cdot 0.9}{0.4} - 0.15\right)(0.43 - 0.3875) = -0.006375 < 0.$$

Thus, in this case, a proportional reduction in documentation-related costs, corresponding to a decrease in β, increases the class-H field evaluator's payoff. If instead $c_d = 0.05$, with all other parameters unchanged, then

$$\frac{\mathrm{d}\mathcal{V}_H}{\mathrm{d}\beta} = -0.05(0.3875 - 0.26) + \left(\frac{0.2 \cdot 0.9}{0.4} - 0.05\right)(0.43 - 0.3875) = 0.010625 > 0.$$

Thus, in this case, the same proportional cost reduction decreases the class-H field evaluator's payoff, confirming the effect on $\mathcal{V}_H$ is ambiguous.

Now suppose only the inflated-report documentation cost c_d' varies, holding c_d and c_p fixed. Let $\acute{c}(c_d')$, $\grave{c}(c_d')$, and $U^*(c_d')$ denote the corresponding thresholds and optimal objective value; the threshold $\check{c}$ is independent of c_d'. A reduction in the cost of substantiating an inflated report corresponds to a decrease in c_d'.

In the intermediate separating regime, the administrator's expected payoff can be written as

$$\begin{aligned}U^*(c_d') = \rho_L &\left[\int_{c_L}^{\acute{c}(c_d')} w(c;\theta_L)\,\mathrm{d}F(c) + \left(1 - \frac{c_d'}{\theta_L}\right)\int_{\acute{c}(c_d')}^{\check{c}} w(c;\theta_L)\,\mathrm{d}F(c)\right] \\ &+ \rho_H\left[\int_{c_L}^{\check{c}} w(c;\theta_H)\,\mathrm{d}F(c) + \frac{c_d'}{\theta_L}\int_{\check{c}}^{\grave{c}(c_d')} w(c;\theta_H)\,\mathrm{d}F(c) - c_p\left(F(\grave{c}(c_d')) - F(\acute{c}(c_d'))\right)\right].\end{aligned}$$

When the thresholds are interior, differentiating with respect to c_d' gives

$$\begin{aligned}\frac{\mathrm{d}U^*(c_d')}{\mathrm{d}c_d'} = &\left[\frac{c_d'\rho_L}{\theta_L} w(\acute{c}(c_d');\theta_L) + c_p\rho_H\right] f(\acute{c}(c_d'))\frac{\mathrm{d}\acute{c}(c_d')}{\mathrm{d}c_d'} \\ &+ \rho_H\left[\frac{c_d'}{\theta_L} w(\grave{c}(c_d');\theta_H) - c_p\right] f(\grave{c}(c_d'))\frac{\mathrm{d}\grave{c}(c_d')}{\mathrm{d}c_d'} \\ &- \frac{\rho_L}{\theta_L}\int_{\acute{c}(c_d')}^{\check{c}} w(c;\theta_L)\,\mathrm{d}F(c) + \frac{\rho_H}{\theta_L}\int_{\check{c}}^{\grave{c}(c_d')} w(c;\theta_H)\,\mathrm{d}F(c).\end{aligned}$$

The first two terms vanish by the cutoff equations

$$c_d'\rho_L w(\acute{c}(c_d');\theta_L) + c_p\rho_H\theta_L = 0, \qquad c_d' w(\grave{c}(c_d');\theta_H) - c_p\theta_L = 0.$$

Hence

$$\frac{\mathrm{d}U^*(c_d')}{\mathrm{d}c_d'} = \int_{\acute{c}(c_d')}^{\check{c}}\left[-\frac{\rho_L}{\theta_L} w(c;\theta_L)\right]\mathrm{d}F(c) + \int_{\check{c}}^{\grave{c}(c_d')}\left[\frac{\rho_H}{\theta_L} w(c;\theta_H)\right]\mathrm{d}F(c).$$

On $[\acute{c}(c_d'), \check{c})$, $w(c;\theta_L) \leq 0$, and on $[\check{c}, \grave{c}(c_d'))$, $w(c;\theta_H) \geq 0$. Therefore,

$$\frac{\mathrm{d}U^*(c_d')}{\mathrm{d}c_d'} \geq 0.$$

That is, decreasing c_d' weakly decreases the administrator's expected payoff. The calculation above assumes interior cutoffs. Boundary cutoffs can be handled by the corresponding one-sided derivative; if a cutoff is locally pinned at c_L or c_H, its derivative term simply drops out.

The class-L field evaluator's expected payoff is

$$\mathcal{V}_L(c_d') = \theta_L F(\acute{c}(c_d')) + (\theta_L - c_d')\left[F(\check{c}) - F(\acute{c}(c_d'))\right].$$

Differentiating gives

$$\frac{\mathrm{d}\mathcal{V}_L(c_d')}{\mathrm{d}c_d'} = c_d' f(\acute{c}(c_d'))\frac{\mathrm{d}\acute{c}(c_d')}{\mathrm{d}c_d'} - \left[F(\check{c}) - F(\acute{c}(c_d'))\right].$$

The cutoff equation for $\acute{c}(c_d')$ implies

$$\frac{\mathrm{d}\acute{c}(c_d')}{\mathrm{d}c_d'} \leq 0.$$

Hence, we have
$$\frac{\mathrm{d}\mathcal{V}_L(c_d')}{\mathrm{d}c_d'} \le 0,$$
and a reduction in c_d' weakly increases the payoff of a class-L field evaluator.

The class-H field evaluator's expected payoff is

$$\mathcal{V}_H(c_d') = \theta_H F(\acute{c}(c_d')) + (\theta_H - c_d)\left[F(\check{c}) - F(\acute{c}(c_d'))\right] + \left(\frac{c_d'\theta_H}{\theta_L} - c_d\right)\left[F(\grave{c}(c_d')) - F(\check{c})\right].$$

Differentiating gives

$$\frac{\mathrm{d}\mathcal{V}_H(c_d')}{\mathrm{d}c_d'} = c_d f(\acute{c}(c_d'))\frac{\mathrm{d}\acute{c}(c_d')}{\mathrm{d}c_d'} + \frac{\theta_H}{\theta_L}\left[F(\grave{c}(c_d')) - F(\check{c})\right] + \left(\frac{c_d'\theta_H}{\theta_L} - c_d\right) f(\grave{c}(c_d'))\frac{\mathrm{d}\grave{c}(c_d')}{\mathrm{d}c_d'}.$$

The first term is nonpositive, while the latter two terms are nonnegative when the corresponding intervals are nonempty. Therefore the sign is not determined in general.

To see that either sign can arise, consider a uniform fulfillment-cost distribution on $[0,1]$ with $\alpha = 0$, so $w(c;\theta) = \theta - 2c$. Let

$$\theta_L = 0.4, \quad \theta_H = 0.675, \quad \rho_L = 0.2, \quad \rho_H = 0.8, \quad c_p = 0.02, \quad c_d = 0.10.$$

These parameters satisfy the conditions of the intermediate separating regime for both $c_d' = 0.2$ and $c_d' = 0.3$. The thresholds are

$$\acute{c}(c_d') = 0.2 + \frac{0.016}{c_d'}, \qquad \check{c} = 0.31, \qquad \grave{c}(c_d') = 0.3375 - \frac{0.004}{c_d'}.$$

Substituting these expressions into the derivative of the class-H field evaluator's payoff gives

$$\frac{\mathrm{d}\mathcal{V}_H(c_d')}{\mathrm{d}c_d'} = \frac{297}{6400} - \frac{1}{500(c_d')^2}.$$

When $c_d' = 0.2$,
$$\frac{\mathrm{d}\mathcal{V}_H(c_d')}{\mathrm{d}c_d'} = \frac{297}{6400} - \frac{1}{500(0.2)^2} = -0.00359375 < 0.$$
Thus, at $c_d' = 0.2$, reducing c_d' increases the payoff of a class-H field evaluator. By contrast, when $c_d' = 0.3$,
$$\frac{\mathrm{d}\mathcal{V}_H(c_d')}{\mathrm{d}c_d'} = \frac{297}{6400} - \frac{1}{500(0.3)^2} \approx 0.02418 > 0.$$
Thus, at $c_d' = 0.3$, reducing c_d' decreases the payoff of a class-H field evaluator. Therefore, the effect of reducing c_d' on $\mathcal{V}_H$ is ambiguous.

Finally, in the pooling regime in Proposition 6(3), the optimal mechanism uses no documentation and no supplier-financed payments:

$$\tau_L(c) = \tau_H(c) = \gamma_L(c) = \gamma_H(c) = 0.$$

The pooling threshold depends only on

$$\rho_L w(c;\theta_L) + \rho_H w(c;\theta_H),$$

which is independent of c_d, c'_d, and c_p. Hence both cost changes leave the mechanism and all payoffs unchanged. $\square$

Proof of Proposition 11. We begin with the first part of the proposition. Under Proposition 7, the optimal mechanism satisfies

$$\tau_L(c) = \tau_H(c) = \gamma_H(c) = 0, \qquad \gamma_L(c) = \theta_L \mathbb{1}_{c \in [\bar{c}, \acute{c})},$$

where

$$\bar{c} = \sup\{c : w(c;\theta_L) + \theta_L \geq 0\}$$

and

$$\acute{c} = \sup\{c : -\rho_L\theta_L + \rho_H w(c;\theta_H) \geq 0\}.$$

Neither threshold depends on c_d, c'_d, or c_p, and the mechanism uses no documentation. Therefore, as long as the parameters remain in the same regime, either cost change has no local effect on the optimal mechanism or on the equilibrium payoffs of the participants.

Next, we prove the second part of the proposition, where the mechanism is the one characterized in Proposition 8. First consider the proportional change

$$\hat{c}_d = \beta c_d, \qquad \hat{c}'_d = \beta c'_d, \qquad \hat{c}_p = \beta c_p.$$

The thresholds $\acute{c}$ and $\dot{c}$ depend on the ratio c_p/c'_d, while $\check{c}$ and $\grave{c}$ do not depend on documentation costs. Therefore, all four thresholds remain unchanged under proportional scaling, as long as the parameters remain in the same regime. In this regime, the partial-authorization probability is $d(\beta) = \beta c'_d/\theta_L$, and the incentive payment on $[\check{c}, \grave{c})$ is $\theta_L - \beta c'_d$.

Using the optimal mechanism in Proposition 7, the administrator's expected payoff can be written as

$$\begin{aligned}
U^*(\beta) = &\int_{c_L}^{\acute{c}} [\rho_L w(c;\theta_L) + \rho_H w(c;\theta_H)] \, \mathrm{d}F(c) \\
&+ \int_{\acute{c}}^{\check{c}} \left[\rho_L \left(1 - \frac{\beta c'_d}{\theta_L}\right) w(c;\theta_L) + \rho_H w(c;\theta_H) - \beta\rho_H c_p\right] \mathrm{d}F(c) \\
&+ \int_{\check{c}}^{\grave{c}} [\rho_H w(c;\theta_H) - \beta\rho_H c_p - \rho_L(\theta_L - \beta c'_d)] \, \mathrm{d}F(c) \\
&+ \int_{\grave{c}}^{\dot{c}} \rho_H \left[\frac{\beta c'_d}{\theta_L} w(c;\theta_H) - \beta c_p\right] \mathrm{d}F(c).
\end{aligned}$$

Differentiating with respect to β yields

$$\frac{\mathrm{d}U^*(\beta)}{\mathrm{d}\beta} = \int_{\acute{c}}^{\check{c}} \left[-\frac{\rho_L c_d'}{\theta_L} w(c;\theta_L) - \rho_H c_p \right] \mathrm{d}F(c) + \int_{\check{c}}^{\grave{c}} [\rho_L c_d' - \rho_H c_p] \, \mathrm{d}F(c) + \int_{\grave{c}}^{\dot{c}} \rho_H \left[\frac{c_d'}{\theta_L} w(c;\theta_H) - c_p \right] \mathrm{d}F(c).$$

On $[\acute{c}, \check{c})$, the definition of $\acute{c}$ and the monotonicity of $w(c;\theta_L)$ imply

$$c_d' \rho_L w(c;\theta_L) + c_p \rho_H \theta_L \leq 0,$$

so the first integrand is nonnegative. On $[\check{c}, \grave{c})$, the second integrand is nonnegative by the given condition

$$c_d' \rho_L - c_p \rho_H \geq 0.$$

On $[\grave{c}, \dot{c})$, the definition of $\dot{c}$ implies

$$c_d' w(c;\theta_H) - c_p \theta_L \geq 0,$$

so the third integrand is also nonnegative. Hence

$$\frac{\mathrm{d}U^*(\beta)}{\mathrm{d}\beta} \geq 0,$$

implying that a proportional reduction in documentation-related costs, corresponding to a decrease in β, weakly decreases the administrator's expected payoff.

The class-L field evaluator's payoff is

$$\mathcal{V}_L(\beta) = \theta_L F(\acute{c}) + (\theta_L - \beta c_d') \left[F(\grave{c}) - F(\acute{c}) \right],$$

which is decreasing in β. Hence a proportional cost reduction weakly increases $\mathcal{V}_L$. The class-H field evaluator's expected payoff is

$$\mathcal{V}_H(\beta) = \theta_H F(\acute{c}) + (\theta_H - \beta c_d) \left[F(\grave{c}) - F(\acute{c}) \right] + \beta \left(\frac{c_d' \theta_H}{\theta_L} - c_d \right) \left[F(\dot{c}) - F(\grave{c}) \right].$$

Differentiating with respect to β gives

$$\frac{\mathrm{d}\mathcal{V}_H}{\mathrm{d}\beta} = -c_d \left[F(\grave{c}) - F(\acute{c}) \right] + \left(\frac{c_d' \theta_H}{\theta_L} - c_d \right) \left[F(\dot{c}) - F(\grave{c}) \right].$$

The first term is nonpositive, whereas the second term is nonnegative because $c_d' \geq c_d$ and $\theta_H \geq \theta_L$. Hence the sign depends on the relative lengths of the intervals $[\acute{c}, \grave{c})$ and $[\grave{c}, \dot{c})$, and is not determined in general.

To see that both signs can arise, consider a uniform fulfillment-cost distribution on $[0,1]$ with $\alpha = 0$. Then $w(c;\theta) = \theta - 2c$. Let

$$\theta_L = 0.4, \quad \theta_H = 1.2, \quad \rho_L = 0.3, \quad \rho_H = 0.7, \quad c_d' = 0.2, \quad c_p = 0.05.$$

These parameters satisfy the conditions of Proposition 8, and the relevant thresholds are

$$\acute{c}=0.3167, \quad \check{c}=0.4, \quad \grave{c}=0.5143, \quad \dot{c}=0.5500.$$

If $c_d=0.05$, then

$$\frac{\mathrm{d}\mathcal{V}_H}{\mathrm{d}\beta}=-0.05(0.5143-0.3167)+\left(\frac{0.2\cdot 1.2}{0.4}-0.05\right)(0.5500-0.5143)\approx 0.0098>0.$$

Thus, in this case, a proportional cost reduction, corresponding to a decrease in β, decreases the class-H field evaluator's payoff. If instead $c_d=0.18$, with all other parameters unchanged, then

$$\frac{\mathrm{d}\mathcal{V}_H}{\mathrm{d}\beta}=-0.18(0.5143-0.3167)+\left(\frac{0.2\cdot 1.2}{0.4}-0.18\right)(0.5500-0.5143)\approx -0.0206<0.$$

Thus, in this case, the same proportional cost reduction increases the class-H field evaluator's payoff, confirming that the effect on $\mathcal{V}_H$ is ambiguous.

Now consider a reduction in c_d' alone, holding c_d and c_p fixed. Let $\acute{c}(c_d')$, $\dot{c}(c_d')$, and $U^*(c_d')$ denote the corresponding thresholds and optimal objective value. The thresholds $\check{c}$ and $\grave{c}$ are independent of c_d'. Using the optimal mechanism in Proposition 8, the administrator's expected payoff is

$$\begin{aligned}U^*(c_d')=&\int_{c_L}^{\acute{c}(c_d')}[\rho_L w(c;\theta_L)+\rho_H w(c;\theta_H)]\,\mathrm{d}F(c)+\int_{\acute{c}(c_d')}^{\check{c}}\left[\rho_L\left(1-\frac{c_d'}{\theta_L}\right)w(c;\theta_L)+\rho_H w(c;\theta_H)-\rho_H c_p\right]\mathrm{d}F(c)\\&+\int_{\check{c}}^{\grave{c}}[\rho_H w(c;\theta_H)-\rho_H c_p-\rho_L(\theta_L-c_d')]\,\mathrm{d}F(c)+\int_{\grave{c}}^{\dot{c}(c_d')}\rho_H\left[\frac{c_d'}{\theta_L}w(c;\theta_H)-c_p\right]\mathrm{d}F(c).\end{aligned}$$

When the thresholds are interior, differentiating with respect to c_d' gives

$$\begin{aligned}\frac{\mathrm{d}U^*(c_d')}{\mathrm{d}c_d'}=&\left[\frac{c_d'\rho_L}{\theta_L}w(\acute{c}(c_d');\theta_L)+\rho_H c_p\right]f(\acute{c}(c_d'))\frac{\mathrm{d}\acute{c}(c_d')}{\mathrm{d}c_d'}+\rho_H\left[\frac{c_d'}{\theta_L}w(\dot{c}(c_d');\theta_H)-c_p\right]f(\dot{c}(c_d'))\frac{\mathrm{d}\dot{c}(c_d')}{\mathrm{d}c_d'}\\&+\int_{\acute{c}(c_d')}^{\check{c}}\left[-\frac{\rho_L}{\theta_L}w(c;\theta_L)\right]\mathrm{d}F(c)+\int_{\check{c}}^{\grave{c}}\rho_L\,\mathrm{d}F(c)+\int_{\grave{c}}^{\dot{c}(c_d')}\left[\frac{\rho_H}{\theta_L}w(c;\theta_H)\right]\mathrm{d}F(c).\end{aligned}$$

The first two terms vanish by the cutoff equations

$$c_d'\rho_L w(\acute{c}(c_d');\theta_L)+c_p\rho_H\theta_L=0, \qquad c_d' w(\dot{c}(c_d');\theta_H)-c_p\theta_L=0.$$

Hence

$$\frac{\mathrm{d}U^*(c_d')}{\mathrm{d}c_d'}=\int_{\acute{c}(c_d')}^{\check{c}}\left[-\frac{\rho_L}{\theta_L}w(c;\theta_L)\right]\mathrm{d}F(c)+\int_{\check{c}}^{\grave{c}}\rho_L\,\mathrm{d}F(c)+\int_{\grave{c}}^{\dot{c}(c_d')}\left[\frac{\rho_H}{\theta_L}w(c;\theta_H)\right]\mathrm{d}F(c).$$

The first integrand is nonnegative on $[\acute{c}(c_d'),\check{c})$, the second term is nonnegative, and the third integrand is nonnegative on $[\grave{c},\dot{c}(c_d'))$. Therefore

$$\frac{\mathrm{d}U^*(c_d')}{\mathrm{d}c_d'}\geq 0,$$

meaning that decreasing c_d' weakly decreases the administrator's objective value.

The class-L field evaluator's payoff is

$$\mathcal{V}_L(c_d') = \theta_L F(\acute{c}(c_d')) + (\theta_L - c_d')\,[F(\grave{c}) - F(\acute{c}(c_d'))]\,.$$

Differentiating with respect to c_d' yields

$$\frac{\mathrm{d}\mathcal{V}_L(c_d')}{\mathrm{d}c_d'} = c_d' f(\acute{c}(c_d'))\acute{c}'(c_d') - [F(\grave{c}) - F(\acute{c}(c_d'))]\,.$$

Since the cutoff equation for $\acute{c}(c_d')$ implies $\acute{c}'(c_d') \le 0$, we have

$$\frac{\mathrm{d}\mathcal{V}_L(c_d')}{\mathrm{d}c_d'} \le 0.$$

Thus a reduction in c_d' weakly increases $\mathcal{V}_L$.

Finally, the class-H field evaluator's expected payoff is

$$\mathcal{V}_H(c_d') = \theta_H F(\acute{c}(c_d')) + (\theta_H - c_d)\,[F(\check{c}) - F(\acute{c}(c_d'))] + \left(\frac{c_d'\theta_H}{\theta_L} - c_d\right)[F(\dot{c}(c_d')) - F(\check{c})]\,,$$

where $\check{c}$ and $\grave{c}$ are independent of c_d', while $\acute{c}$ and $\dot{c}$ depend on c_d'. Differentiating with respect to c_d' yields

$$\frac{\mathrm{d}\mathcal{V}_H}{\mathrm{d}c_d'} = c_d f(\acute{c}(c_d'))\frac{\mathrm{d}\acute{c}(c_d')}{\mathrm{d}c_d'} + \frac{\theta_H}{\theta_L}\,[F(\dot{c}(c_d')) - F(\check{c})] + \left(\frac{c_d'\theta_H}{\theta_L} - c_d\right) f(\dot{c}(c_d'))\frac{\mathrm{d}\dot{c}(c_d')}{\mathrm{d}c_d'}.$$

The first term is nonpositive, while the latter two terms are nonnegative when the corresponding intervals are nonempty. Therefore, the sign is not determined in general.

To see that either sign can arise, consider a uniform fulfillment-cost distribution on $[0,1]$ with $\alpha = 0$, so that $w(c;\theta) = \theta - 2c$. Let

$$\theta_L = 0.4, \quad \theta_H = 1, \quad \rho_L = 0.2, \quad \rho_H = 0.8, \quad c_p = 0.04, \quad c_d = 0.15.$$

These parameters satisfy the conditions of Proposition 8 for both $c_d' = 0.2$ and $c_d' = 0.3$. The thresholds are

$$\acute{c}(c_d') = 0.2 + \frac{0.032}{c_d'}, \qquad \check{c} = 0.4, \qquad \grave{c} = 0.45, \qquad \dot{c}(c_d') = 0.5 - \frac{0.008}{c_d'}.$$

Hence

$$\frac{\mathrm{d}\mathcal{V}_H}{\mathrm{d}c_d'} = 0.125 - \frac{0.04c_d}{(c_d')^2}.$$

When $c_d' = 0.2$,

$$\frac{\mathrm{d}\mathcal{V}_H}{\mathrm{d}c_d'} = 0.125 - \frac{0.04(0.15)}{0.2^2} = -0.025 < 0.$$

Thus, at $c_d' = 0.2$, a reduction in c_d' increases the class-H field evaluator's payoff. By contrast, when $c_d' = 0.3$,

$$\frac{\mathrm{d}\mathcal{V}_H}{\mathrm{d}c_d'} = 0.125 - \frac{0.04(0.15)}{0.3^2} \approx 0.0583 > 0.$$

Thus, at $c_d' = 0.3$, a reduction in c_d' decreases the class-H field evaluator's payoff, confirming the effect of reducing c_d' on $\mathcal{V}_H$ is ambiguous. □