TRUTHFUL DATA PROCUREMENT VIA EXPERIMENTAL DESIGN

Anonymous authors
Paper under double-blind review

ABSTRACT

Building a model for a specific task, whether reading medical images or training a warehouse robot from video, often requires buying data from multiple sellers. How should a buyer pay for data she has not yet seen? While the problem can be viewed as a reverse auction, the nature of data introduces three idiosyncratic difficulties. A buyer should evaluate datasets and select sellers without fully accessing data, since once she has seen it she has no reason to pay (Arrow's inspection paradox). A selected seller can misreport labels. And even if sellers were honest, finding the optimal set of datasets is computationally intractable.

We study data procurement for supervised learning, where a buyer purchases datasets from strategic sellers with private costs, and we design a mechanism, the Greedy Auction for Data Elicitation via Surrogate (GRADES), which addresses those three difficulties. In the OLS V-optimal design setting of Lu et al. (2024), GRADESis cost-truthful and individually rational in expectation over the label noise, makes truthful label delivery a dominant strategy, and achieves a bi-criteria welfare guarantee. We further extend GRADESto generalized linear models and to Bayesian decision problems with informational substitutes. Finally, we show empirically that it matches the welfare of an existing non-strategic procurement method.

1 INTRODUCTION

Machine learning models are increasingly adapted to specific applications, which requires acquiring data for each application from many holders. Building a model to read medical images needs data from multiple hospitals Medical Imaging and Data Resource Center (2026); Advanced Research Projects Agency for Health (2024). Training a video model licenses footage from many creators Bloomberg News (2025). Robotics developers pay contributors to record household chores ABC News (2026). In each case, the buyer (called principal) decides which sellers to buy from and how much to pay them. The value of a seller's data depends on the buyer's task, and the buyer cannot fully assess it before the purchase. How should a buyer pay for data she has not yet seen?

We study this question as procurement for supervised learning Agarwal et al. (2019); Castro Fernandez (2023). A principal has a value function $v(S)$—the prediction quality achieved by training on datasets $\{D_i\}_{i\in S}$—and seeks to select a set of sellers maximizing her value net of their costs, where c_i is seller i's private cost.

At a high level, this can be viewed as a reverse auction (procurement) over N single-parameter agents (Vickrey, 1961; Archer & Tardos, 2001) (see Section 2.1 for details). Given a known value function, sellers (mis)report their cost, the mechanism selects k winner to maximizing the welfare, and each selected seller is paid and delivers her data. This is a standard reverse-auction design dating back to Myerson (1981) However, the connection is nominal. It assumes that $v(S)$ is known, that sellers cannot manipulate the data they deliver, and that the welfare-maximizing allocation is efficiently computable. Data procurement violates all three assumptions:

- **Arrow's inspection paradox** (Arrow, 1962) The value $v(S)$ depends on the sellers' data. If the principal can directly access it during allocation, the principal has no incentive to pay afterward. The mechanism must partition each seller's data into two components: a *report*

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

shared freely at Step 1 to inform allocation, and a *deliverable* purchased at Step 3 that the principal still needs after seeing the report. The report must be *sufficient for allocation but insufficient for learning*.

- **Delivery moral hazard.** Once selected and paid, what prevents sellers from delivering fabricated labels? This is an *information elicitation* problem: the mechanism must incorporate additional incentives to elicit truthful data delivery.
- **Computational intractability.** The first two challenges are unique to information goods. A third is common to combinatorial auctions: maximizing the welfare $v(S) - \sum_{i \in S} c_i$ is NP-hard for general set functions, and even for submodular v admits no constant multiplicative approximation (Sviridenko et al., 2017). Addressing this requires understanding structural properties of v to design approximation algorithms with provable guarantees.

Our results These difficulties suggest that a data market needs more than one score. The principal must score datasets by their usefulness before seeing their labels, score cost reports against competing sellers, and score delivered labels. In this paper, we provide a framework that coordinates these scores.

We propose Greedy Auction for Data Elicitation via Surrogate (GRADES). The mechanism first computes a surrogate value from what sellers reveal when they bid, without fully accessing their data, and runs the distorted greedy algorithm of Harshaw et al. Harshaw et al. (2019) on this surrogate. It pays each winner her critical bid plus a data quality bonus that measures how much her labels improve prediction. The surrogate addresses Arrow's paradox by letting the principal allocate without fully observing the data she pays for. The data quality bonus incentivizes winners to deliver truthful labels. The distorted greedy algorithm provides an approximation guarantee.

Our main theoretical result applies GRADES to least squares regression with V-optimal design, the setting of Lu et al. (2024) We show that GRADES is dominant-strategy truthful in costs and individually rational, and that truthful label delivery is a best response for every winner. Theorems 2.2 and 2.6 Moreover, GRADES achieves a $(1-e^{-\delta}, 1)$ bi-criteria welfare guarantee in polynomial time, where δ can be as large as the surrogate's submodularity ratio. We also provide an almost matching hardness result (Theorem 2.3).

We then generalize the framework beyond least squares (Theorem 3.2). On the valuation side, Section 3.1 extends surrogate valuation to generalized linear models and to Bayesian decision problems in which datasets are informational substitutes. On the data quality side, Section 3.2 shows that the bonus can be built from any truthful score in information elicitation, such as a proper scoring rule or peer prediction.

Finally, we compare GRADES with adapted single- and multi-step DAVED (Lu et al., 2024), Data-OOB Kwon & Zou (2023a) and random selection in Section 4. The comparison covers synthetic Gaussian regression and four real datasets (MIMIC-III, RSNA Bone Age, Fitzpatrick17K and DrugLib), with 300 single-record sellers per market and two cost laws. At capacity $k = 25$, GRADES attains the highest average welfare in nine of ten dataset.(Table 1) Its selection step is 260–420 times faster than multi-step DAVED in Fig. 2. Finally, the test error (MSE) is comparable in Fig. 1.

RELATED WORK

Information can be scored in many ways, largely because different jobs call for different scores, and a data market needs several at once. We review three types, one for each component of GRADES.

Three class of information scores The first type treats information as an ordinary good, priced by the cost of acquiring it. Following the Pandora's-box model of Weitzman (1978). With strategic sources the cost is private, and acquisition becomes single-parameter mechanism design (Myerson, 1981; Roth & Schoenebeck, 2012; Ghosh et al., 2014; Chen et al., 2018). GRADES uses its critical-bid to compensate the cost.

The second type values information by what it does for a particular decision (Lindley, 1956; Raiffa & Schlaifer, 2000; Howard, 1966). Examples include the reduction in expected test error (V-objective) (Chaloner & Verdinelli, 1995b; Lu et al., 2024), in expected squared distance from the

true parameter (A-objective) (Chaloner & Verdinelli, 1995b; Bian et al., 2017), and in the log-determinant of the posterior covariance (D-objective) (Chaloner & Verdinelli, 1995b; Horel et al., 2014). These scores depend on the decision problem: the V-objective, for instance, depends on the test design, so different test points value the same data differently. Finally, Hu et al. (2026) instead score each seller by a single-parameter function.

Algorithmically, GRADES is closest to DAVED (Lu et al., 2024), which optimizes a continuous relaxation of the V-objective by Frank–Wolfe herding. Their algorithm runs until the selected data exhaust the budget which may not achieve polynomial running time. GRADES instead runs the distorted greedy algorithm of Harshaw et al. (2019) on the V-objective, in polynomial time, with a bi-criteria welfare guarantee (Theorem 2.2) that is nearly tight (Theorem 2.3). Our simulation in Fig. 2 compares the running time between DAVED and GRADES.

Finally, date quality score reward truthful data over any other manipulation, which does not depend on the decision problem. A strictly proper score makes truthful reporting optimal whatever the report is used for (Glenn et al., 1950; Savage, 1971; Gneiting & Raftery, 2007). Peer prediction replaces ground truth with other reports (Miller et al., 2005; Prelec, 2004), market scoring rules pay for marginal improvement to a shared hypothesis (Abernethy & Frongillo, 2011). GRADES introduce such a score for label incentives.

The valuation of information is a set function $v : 2^{[N]} \to \mathbb{R}$, and hence a high-dimensional object. A line of work therefore aggregates v into a scalar capturing seller i's "contribution": Data Shapley(Ghorbani & Zou, 2019), Leave-One-Out (Koh & Liang, 2017),Data-OOB (Kwon & Zou, 2023b), Data Banzhaf (Wang & Jia, 2023), and related variants. Their guarantees concern how a given v is divided and depend on the choice of v.

From welfare algorithms to procurement mechanisms. There is a long line of research on general procurement mechanisms. Ours relates most to those have approximation submodular valuation.

Three objectives are standard. Budget-feasible mechanisms (Singer, 2010) maximize value subject to a hard cap on total payments, and recent variants target welfare under the same cap (Cui et al., 2026; Eden et al., 2026). All of them assume submodular values. Maximizing Surplus, value minus payments, prior-free is difficulty due to frugality lower bounds (Karlin & Kempe, 2005), and maximizing it in the Bayesian setting requires cost priors (Myerson, 1981; Chen et al., 2018). Finally, welfare does not admits constant approximation, so guarantees are bi-criteria and depend on the structure of g. For submodular g the bound $(1 - 1/e, 1)$ is tight (Sviridenko et al., 2017). For γ-weakly submodular g (Das & Kempe, 2018), distorted greedy attains $(1-e^{-\gamma}, 1)$ (Harshaw et al., 2019), with extensions (Wang et al., 2021; Geng et al., 2022; Gong et al., 2023). For experimental design, bounds on the submodularity ratio (Bian et al., 2017; Chamon & Ribeiro, 2017) control greedy's approximation ratio. Theorem 2.5 sharpens the Woodbury bound of Harshaw et al. (2019, Claim 10) to block leverage.

2 DATA PROCUREMENT FOR ORDINARY LEAST SQUARES

We start with the setting of DAVED (Lu et al., 2024), data acquisition as V-optimal experimental design, and extend it to a general framework in Section 3.

2.1 THE PROCUREMENT GAME

A principal has a prediction task, and buys labeled data from N potential sellers. Seller $i \in [N] = \{1, \ldots, N\}$ holds a dataset $D_i = (\boldsymbol{X}_i, \boldsymbol{y}_i)$ of n_i labeled points, with features $\boldsymbol{X}_i \in \mathbb{R}^{n_i \times d}$ and labels $\boldsymbol{y}_i \in \mathbb{R}^{n_i}$, and incurs a private *cost* $c_i > 0$ to deliver it (effort, opportunity cost, or privacy loss). We write $\boldsymbol{c} = (c_1, \ldots, c_N)$. The principal holds her own labeled data $D_0 = (\boldsymbol{X}_0, \boldsymbol{y}_0)$ of n_0 points and a *capacity* $k \in \mathbb{N}$, the largest number of sellers she will buy from.

A *mechanism* $(\pi, \boldsymbol{p})$ is an allocation rule π, which maps the sellers' reports to a set $S \subseteq [N]$ with $|S| \leq k$, and a payment rule $\boldsymbol{p} = (p_1, \ldots, p_N)$, which maps reports and delivered data to payments, with $p_i = 0$ for $i \notin S$. The game has three stages. Given the principal's mechanism $(\pi, \boldsymbol{p})$, k, D_0, and a prediction task specified in Theorem 2.1 and generalized in Section 3.1,

1. **Report.** Each seller i publishes her features $\boldsymbol{X}_i$[1] and reports a cost $\hat{c}_i$, possibly different from c_i.
2. **Allocate.** The mechanism selects the winners $S = \pi(\hat{\boldsymbol{c}}, \boldsymbol{X})$ from the cost reports $\hat{\boldsymbol{c}} = (\hat{c}_1, \ldots, \hat{c}_N)$ and the features $\boldsymbol{X} = (\boldsymbol{X}_1, \ldots, \boldsymbol{X}_N)$ alone.
3. **Deliver and pay.** Each winner $i \in S$ delivers labels $\hat{\boldsymbol{y}}_i \in \mathbb{R}^{n_i}$, possibly different from $\boldsymbol{y}_i$,[2] and receives p_i. The principal trains on the delivered data $\hat{D}_S = D_0 \cup \{(\boldsymbol{X}_i, \hat{\boldsymbol{y}}_i)\}_{i \in S}$, and gets $\boldsymbol{y}^{\text{test}}$

We use hats to mark reports: $\hat{c}_i$, $\hat{\boldsymbol{y}}_i$ and $\hat{D}_S$. Quantities the principal *estimates* from data carry a check instead, as in $\check{\theta}_S$ for the fit on D_S. A pure *strategy* for seller i is a cost report $\hat{c}_i(c_i)$ at stage 1 and a delivery rule $\hat{\boldsymbol{y}}_i(\boldsymbol{y}_i, S)$ at stage 3, and her utility is quasi-linear, $u_i = (p_i - c_i)\mathbf{1}[i \in S]$, the cost being incurred only upon selection and delivery.

The principal's objective is the *welfare*

$$W(S; \boldsymbol{c}) \coloneqq v(S) - c(S) \tag{1}$$

where the *valuation* $v : 2^{[N]} \to \mathbb{R}$ is the improvement due to the truthful data from sellers in S, and $c(S) = \sum_{i \in S} c_i$ is the private total cost of the selected sellers. As sellers are strategic and can misreport both costs and labels, we introduce two solution concepts for them.

1. **Cost incentive (DSIC and IR).** Truthful cost reporting $\hat{c}_i = c_i$ is a dominant strategy for every seller, and participation is individually rational, $p_i - c_i \geq 0$ if $i \in S$. Here we additional take expectation over the labels.
2. **Label incentive.** Given a common prior, truthful delivery $\hat{\boldsymbol{y}}_i = \boldsymbol{y}_i$ is a best response for every winner under a common prior μ specified before Theorem 2.6.
3. **Bi-criteria welfare approximation.** The allocation is computable in polynomial time and is (α, β)-*approximate*[3] for $\alpha, \beta \in (0, 1]$ if $W(S; \boldsymbol{c}) \geq \max_{S' \subseteq [N] : |S'| \leq k} \alpha v(S') - \beta c(S')$.

Value and Surrogate

Definition 2.1. *Labels follow $y = x^\top \theta + \epsilon$ with an unknown parameter $\theta \in \mathbb{R}^d$ and independent Gaussian noise with $\mathbb{E}[\epsilon] = 0$ and $\mathbb{E}[\epsilon^2] = \sigma^2$, for the principal's, the sellers', and the test data. The principal has an unlabeled* test design *$\boldsymbol{X}^{\text{test}} \in \mathbb{R}^{m \times d}$ to predict, and uses the least-squares fit $\check{\theta}_S \coloneqq f_{\text{LS}}(D_S)$ on the pooled truthful data D_S.*

The principal's *risk* is her expected test error, $\ell(S) \coloneqq \frac{1}{m} \sum_{q=1}^{m} \mathbb{E}\big[(x_q^{\text{test}})^\top \check{\theta}_S - y_q^{\text{test}}\big]^2$. The *realized test loss* of an estimate $\check{\theta}$, and of the fit on a dataset T, are $\bar{\ell}(\check{\theta}) \coloneqq \frac{1}{m} \sum_{q=1}^{m} \big((x_q^{\text{test}})^\top \check{\theta} - y_q^{\text{test}}\big)^2$ and $\bar{\ell}(T) \coloneqq \bar{\ell}\big(f_{\text{LS}}(T)\big)$, so that $\ell(S) = \mathbb{E}[\bar{\ell}(D_S)]$. The *true valuation* is the risk reduction from acquiring S,

$$v(S) \coloneqq \ell(\emptyset) - \ell(S). \tag{2}$$

In general the valuation is not available at stage 2, as the risk of a learner depends on the labels. The allocation therefore runs on a *surrogate* computed from the features alone. We define the information matrix of the principal's features together with those of a subset S of sellers as $\boldsymbol{I}(S) \coloneqq \boldsymbol{X}_0^\top \boldsymbol{X}_0 + \sum_{i \in S} \boldsymbol{X}_i^\top \boldsymbol{X}_i$, and the *V-optimal surrogate* with the *task weight* $\boldsymbol{A} \coloneqq \frac{1}{m} (\boldsymbol{X}^{\text{test}})^\top \boldsymbol{X}^{\text{test}}$ is

$$\tilde{v}(S) \coloneqq \sigma^2 \big[\operatorname{tr}(\boldsymbol{A}\boldsymbol{I}(\emptyset)^{-1}) - \operatorname{tr}(\boldsymbol{A}\boldsymbol{I}(S)^{-1})\big]. \tag{3}$$

We assume throughout that $\boldsymbol{X}_0$ and $\boldsymbol{X}^{\text{test}}$ have full rank d.

2.2 The Grades Mechanism

We introduce *Greedy Auction for Data Elicitation via Surrogate* (Grades) which consists of three components: The *allocation* is a distorted greedy (Harshaw et al., 2019) run on the surrogate $\tilde{v}$.

[1] Here we assumes that features are reported truthfully,e.g., (raw imaging modalities)

[2] We assume her labels $\boldsymbol{y}_i$ are realized only after allocation, and that c_i is independent of θ and of the label noise, so that her information at bidding carries no signal about her quality bonus.

[3] A multiplicative guarantee ($\beta = \alpha$) is unattainable even for submodular v (Sviridenko et al., 2017).

The *cost payment* is the classical Myerson's payment (Myerson, 1981; Archer & Tardos, 2001). Finally, the *quality bonus* is the recentered marginal improvement of the delivered labels defined in Theorem 2.6. We further generalize them in Section 3.2.

The GRADES Mechanism (Greedy Auction for Data Elicitation via Surrogate).
Input: capacity k, distortion $\delta \in (0, 1]$, a surrogate $\tilde{v}(S)$, and a *quality score* b_i evaluated at delivery (instantiated in Theorem 2.6).

1. **Report.** Each seller i submits features $\boldsymbol{X}_i$ and cost $\hat{c}_i$.
2. **Allocate.** Initialize $S \leftarrow \emptyset$. For $t = 1, \ldots, k$, let

$$i_t \in \arg\max_{i \notin S} \Big(1 - \tfrac{\delta}{k}\Big)^{k-t} \big[\tilde{v}(S \cup \{i\}) - \tilde{v}(S)\big] - \hat{c}_i.$$

   If the maximum score is positive, set $S \leftarrow S \cup \{i_t\}$. Otherwise leave S unchanged (the round is skipped). Output $S_{\text{greedy}} = S$.
3. **Deliver and pay.** Sellers in S_{greedy} deliver labels. Seller $i \in S_{\text{greedy}}$ receives

$$p_i = \tau_i + b_i, \tag{4}$$

   where $\tau_i = \sup\{\hat{c}_i : i \in S_{\text{greedy}}(\hat{c}_i, \hat{c}_{-i})\}$ is the critical bid.

Theorem 2.2 (Cost incentive and welfare)**.** *Given the linear model of Theorem 2.1, there is* $\gamma \geq \frac{1}{1+h_{\max}} > 0$*, and* $h_{\max} \coloneqq \max_{i \in [N]} \lambda_{\max}\big(\boldsymbol{X}_i \boldsymbol{I}(\emptyset)^{-1} \boldsymbol{X}_i^\top\big)$ *so that for every capacity* k*, distortion* $0 < \delta \leq \gamma$*, and quality score* $\boldsymbol{b} = (b_i)_{i \in [N]}$ *with* $\mathbb{E}[b_i \mid \boldsymbol{X}, \hat{\boldsymbol{c}}] = 0$ *at truthful delivery,* GRADES *run on* $\tilde{v}$ *in Eq.* (3) *satisfies the cost incentive, the allocation is a bi-criteria welfare approximation with* $(\alpha, \beta) = (1 - e^{-\delta}, 1)$*, and runs in* $O(Nk^2(d^3 + d^2 \max_i n_i))$ *time.*

The first challenge is showing approximation guarantee even without cost incentive. The following results shows that our approximation result is almost tight for $\beta = 1$.

Theorem 2.3 (Hardness)**.** *For every constant* $\epsilon > 0$*, it is NP hard to find an* $(1 - 1/e + \epsilon)$*-approximation for* $\tilde{v}$ *even when* $n_i = 1$ *for all* $i \in [N]$*. Moreover, for all* $\beta \in (0, 1]$*, it is NP-hard to have* $(1 - e^{-\beta} + \epsilon, \beta)$*-approximation for the welfare in Eq.* (1)*.*

To address this, we combine the distorted greedy algorithm of Harshaw et al. (2019) with the meta-mechanism of Deng et al. (2025). In each round, a seller's score is non-increasing in her own bid, and the other sellers' scores do not depend on her bid, so lowering her bid can never exclude her and the allocation is monotone. With each winner paid her critical bid, Myerson's lemma yields the cost incentive. Finally, we show that the surrogate $\tilde{v}$ of Eq. (3) coincides with the true valuation v of Eq. (2) and is γ-approximate submodular (Theorem 2.5), to achieve bi-criteria welfare guarantee for the true valuation.

Definition 2.4 (Submodularity ratio (Das & Kempe, 2018))**.** *The* submodularity ratio *of a monotone set function* $g : 2^{[N]} \to \mathbb{R}$ *is the largest* $\gamma \in [0, 1]$ *such that for all* L *and* S *with* $S \cap L = \emptyset$*,* $|S| \leq k$*,*

$$\sum_{j \in S} \big[g(L \cup \{j\}) - g(L)\big] \geq \gamma \cdot \big[g(L \cup S) - g(L)\big].$$

Lemma 2.5 (γ-submodularity of the V-optimal surrogate)**.** *For disjoint subsets* $L, S \subseteq [N]$ *with* $|S| \leq k$*, let* $\boldsymbol{P} \coloneqq \boldsymbol{I}(L)^{-1}$ *and define* $h_j \coloneqq \lambda_{\max}\big(\boldsymbol{X}_j \boldsymbol{P} \boldsymbol{X}_j^\top\big)$ *for seller* $j \in S$*. For every positive semidefinite task weight* $\boldsymbol{A}$*, the submodularity ratio of* $\tilde{v}$ *satisfies*

$$\gamma \geq \min_{L,S} \frac{1}{1 + \max_{j \in S} h_j}.$$

For the quality score b_i of Eq. (4), in contrast to Theorem 2.2, the label incentive needs to specify the sellers' beliefs about the parameter θ (common prior). Here we assume uninformative common prior $p(\theta) \propto 1$. Specifically, for every $S \subseteq [N]$, the sellers' common belief over θ conditioned on the pooled data D_S is square-integrable and has conditional mean equal to the least-squares fit, $\mathbb{E}[\theta \mid D_S] = f_{\text{LS}}(D_S)$.

Theorem 2.6 (Label incentive)**.** *If the sellers' common prior μ is uninformative, then paying winner i the quality score*

$$b_i = \bar{\ell}(D_0) - \bar{\ell}\big(D_0 \cup \{(\boldsymbol{X}_i, \hat{\boldsymbol{y}}_i)\}\big) - \tilde{v}(\{i\}), \tag{5}$$

satisfies $\mathbb{E}[b_i \mid \boldsymbol{X}, \hat{\boldsymbol{c}}] = 0$ at truthful delivery, and GRADES *satisfies the label incentive.*

3 DATA PROCUREMENT FRAMEWORK

Although we introduced GRADES for OLS, its design extends to a general framework for data procurement, given an appropriate surrogate $\tilde{v}$ and quality score $\boldsymbol{b}$. In this section, we give sufficient conditions on $\tilde{v}$ and $\boldsymbol{b}$ so that the cost incentive, the label incentive, and the welfare guarantee hold. We defer further discussion in Section C.

Definition 3.1 (Admissible surrogate and quality score)**.** *Given a valuation $v : 2^{[N]} \to \mathbb{R}$ and constants $\rho \geq 1$, $\gamma \in (0, 1]$ and $\epsilon \geq 0$, the surrogate $\tilde{v}$ is* admissible with (ρ, γ) *if*

1. *$\tilde{v}$ is computable from D_0 and the stage-1 reports, monotone, and non-negative.*
2. *$\tilde{v}$ ρ-approximates v, that is, $\tilde{v}(S)/\rho \leq v(S) \leq \rho\tilde{v}(S)$ for all $S \subseteq [N]$.*
3. *$\tilde{v}$ has submodularity ratio at least γ (Theorem 2.4).*

The quality score $\boldsymbol{b} = (b_i)_{i\in[N]}$ is admissible with ϵ *if at truthful delivery, $\big|\mathbb{E}[b_i \mid \boldsymbol{X}, \hat{\boldsymbol{c}}]\big| \leq \epsilon$ for every winner i.*

A mechanism satisfies the *ϵ-cost incentive* if for every seller i, true cost c_i, report $\hat{c}_i$, and profile $\hat{c}_{-i}$, $\mathbb{E}\big[u_i((c_i, \hat{c}_{-i}); c_i)\big] \geq \mathbb{E}\big[u_i((\hat{c}_i, \hat{c}_{-i}); c_i)\big] - \epsilon$ and $\mathbb{E}\big[u_i((c_i, \hat{c}_{-i}), \boldsymbol{X}; c_i)\big] \geq -\epsilon$.

Theorem 3.2 (Meta-theorem)**.** *For any value function v, admissible $\tilde{v}$ with (ρ, γ) and $\boldsymbol{b}$ with ϵ, capacity k and distortion $0 < \delta \leq \gamma$,* GRADES *satisfies the 2ϵ-cost incentive, and under truthful bids*

$$\rho v(S_{\text{greedy}}) - c(S_{\text{greedy}}) \geq \max_{S' \subseteq [N] : |S'| \leq k} \rho^{-1}(1 - e^{-\delta})v(S') - c(S').$$

Moreover, GRADES *makes $O(Nk^2)$ oracle calls to $\tilde{v}$ and evaluates b_i once per winner.*

Note that $\rho = 1$ reduces to the $(1 - e^{-\delta}, 1)$-approximation of Theorem 2.2.

3.1 VALUE SURROGATES

To avoid Arrow's paradox, we explore the design value surrogate $\tilde{v}$ for the real valuation v as Eqs. (2) and (3). We introduce two general family for admissible $\tilde{v}$.

Fisher information surrogate The Gaussian model with uninformative prior of Section 2 generalizes to a canonical exponential family, $p(y \mid x, \theta) = h(y) \exp\{yx^\top\theta - \psi(x^\top\theta)\}$, with the conjugate prior μ_0 induced by the principal's data D_0. When the principal's loss is the squared error $\|\theta - \check{\theta}\|^2$ of her estimate $\check{\theta}$, the valuation is the reduction in Bayes risk from acquiring the data of the sellers in S, Chaloner & Verdinelli (1995a)

$$v(S) = \operatorname{tr} \operatorname{Cov}_{\mu_0}(\theta) - \mathbb{E}\big[\operatorname{tr} \operatorname{Cov}(\theta \mid D_S)\big], \tag{6}$$

with the expectation over the labels under the prior predictive of μ_0. Eq. (6) generally has no closed form, and one may use the inverse Fisher information. Specifically, for a design $\boldsymbol{X}$ write $\boldsymbol{W}(\boldsymbol{X}) \coloneqq \operatorname{diag}(\psi''(x^\top\check{\theta}_0))_{x\in\boldsymbol{X}}$ and set $\boldsymbol{I}(S) \coloneqq \boldsymbol{X}_0^\top \boldsymbol{W}(\boldsymbol{X}_0)\boldsymbol{X}_0 + \sum_{i\in S} \boldsymbol{X}_i^\top \boldsymbol{W}(\boldsymbol{X}_i)\boldsymbol{X}_i$. We can set surrogate as

$$\tilde{v}(S) \coloneqq \operatorname{tr} \boldsymbol{I}(\emptyset)^{-1} - \operatorname{tr} \boldsymbol{I}(S)^{-1}. \tag{7}$$

This is the V-optimal surrogate Eq. (3) with $\sigma^2 = 1$, $\boldsymbol{A} = \mathbb{I}_d$. It is therefore computable from D_0 and the stage-1 reports, monotone, non-negative, and approximated submodular. By Theorem 2.5. We leave bounding ρ open.

Value of information and informational substitutes. More generally, a Bayesian decision problem $(\Theta, \mathcal{A}, \ell, \mu)$ consists of an outcome $\theta \in \Theta$, an action $a \in \mathcal{A}$, and a loss $\ell : \Theta \times \mathcal{A} \to \mathbb{R}$, together with a joint distribution μ over the outcome and all sellers' data. After observing data the principal plays the Bayes act minimizing her expected loss, and the valuation of a set of sellers is the improvement over expected loss:

$$v_\ell(S) := \min_{a \in \mathcal{A}} \mathbb{E}\big[\ell(\theta, a)\big] - \mathbb{E}\Big[\min_{a \in \mathcal{A}} \mathbb{E}\big[\ell(\theta, a) \mid D_S\big]\Big], \tag{8}$$

the value of information of D_S (Howard, 1966). The valuation of Eqs. (2) and (6) are special cases. In particular, for Eq. (2), the action is a prediction of the test labels, the loss is the squared test loss.

In particular, if v_ℓ is known in the allocation stage and *informational substitutes* (Chen & Waggoner, 2017)— v_ℓ is submodular on $2^{[N]}$— the principal can run GRADES with constant approximation.

3.2 QUALITY SCORES

The delivery stage is an information elicitation problem which ensures that truthful delivery optimizes each selected seller's expected score.

Definition 3.3 (Label-truthful)**.** *A quality score $\boldsymbol{b}$ is* label-truthful *if for every winner i and every report $\hat{\boldsymbol{y}}_i$, when the other winners deliver truthfully,*

$$\mathbb{E}\big[b_i(\hat{\boldsymbol{y}}_i, \boldsymbol{y}_{S\setminus i}) \mid D_i\big] \le \mathbb{E}\big[b_i(\boldsymbol{y}_i, \boldsymbol{y}_{S\setminus i}) \mid D_i\big],$$

with equality only if $f(\hat{D}_S) = f(D_S)$, where f is the principal's learner.

We give two classical constructions of the quality score $\boldsymbol{b}$: a proper scoring rule when the principal holds labeled test data at the payment stage, and peer prediction when she does not.

Let PS be a strictly proper scoring rule for predictive distributions of the test labels. Let $\check{p}_0$ be the principal's baseline prediction and $\check{p}(\hat{\boldsymbol{y}}_i)$ her prediction after adding seller i's delivery. We pay seller i the improvement of her delivery over the baseline, recentered by its mean at truthful delivery:

$$b_i(\hat{\boldsymbol{y}}_S) = G_i - \mathbb{E}[G_i \mid \boldsymbol{X}, \hat{\boldsymbol{c}}], \qquad G_i := \mathrm{PS}\big(\check{p}(\hat{\boldsymbol{y}}_i), \boldsymbol{y}^{\mathrm{test}}\big) - \mathrm{PS}\big(\check{p}_0, \boldsymbol{y}^{\mathrm{test}}\big).$$

Without test labels, the principal can instead score each delivery against the other winners' deliveries, as in peer prediction (Miller et al., 2005; Prelec, 2004). Truthful delivery is then a Bayes–Nash equilibrium rather than a dominant strategy. Let s_i be a peer-prediction score under which truthful delivery is an equilibrium, and suppose its mean at truthful delivery is at most $M > 0$ in absolute value. Then $b_i = \frac{\epsilon}{M} s_i$ is admissible for any $\epsilon > 0$, and positive rescaling preserves the equilibrium.

4 EXPERIMENTS

We evaluate the quality of allocation of our GRADES and other baseline (DAVED(Lu et al., 2024), Data-OOB(Kwon & Zou, 2023a), and Random) and compare the welfare in Table 1, valuation in Fig. 1, and running time in Fig. 2.

4.1 EXPERIMENTAL SETUP

We retain Gaussian regression, MIMIC (Johnson et al., 2016; Pollard et al., 2026), RSNA Bone Age (Halabi et al., 2019), Fitzpatrick17K (available-image subset)(Groh et al., 2021), and DrugLib(Gräßer et al., 2018). Each of 20 market resamples has $N = 300$ single-record sellers, $n_i = 1$, and three disjoint test designs of $m = 32$ records. Given testing data, we have $\boldsymbol{A} = (\boldsymbol{X}^{\mathrm{test}})^\top \boldsymbol{X}^{\mathrm{test}}/m$, and define

$$\boldsymbol{I}_\eta(S) = \eta \mathbb{I}_d + \sum_{i \in S} \boldsymbol{X}_i^\top \boldsymbol{X}_i, \widetilde{v}_\eta(S) = 30 \left(1 - \frac{\mathrm{tr}\left(\boldsymbol{A}\boldsymbol{I}_\eta(S)^{-1}\right)}{\mathrm{tr}\left(\boldsymbol{A}\boldsymbol{I}_\eta(\emptyset)^{-1}\right)}\right), \widetilde{W}_\eta(S; c) = \widetilde{v}_\eta(S) - c(S).$$

The matrix $\eta \mathbb{I}_d$ can be viewed as a regularizer, and $\tilde{v}_\eta$ is a scaled version of Eq. (3). Each seller's cost use saved uniform integer types $r_i \in \{1, \ldots, 5\}$, transformed by $g(r) = \sqrt{r}$ or r^2 and divided by the market mean of $g(r_i)$.

Table 1: Welfare $\widetilde{W}_1(S;c)$ (higher is better) at $k = 25$ and $\eta = 1$ under GRADES-defined acquisition-cost caps. Each dataset has subcolumns for square-root and square costs. Bold marks the best displayed mean among the eight acquisition rules (ties at displayed precision).

	Gaussian		MIMIC		RSNA		Fitzpatrick[†]		DrugLib	
Method	$\sqrt{r}$	r^2	$\sqrt{r}$	r^2	$\sqrt{r}$	r^2	$\sqrt{r}$	r^2	$\sqrt{r}$	r^2
GRADES ($\delta = 1$)	**4.716**	**15.331**	**15.653**	24.092	**20.244**	**27.964**	**20.024**	**28.023**	**20.315**	**28.031**
Greedy-B ($\delta = 0$)	4.698	15.005	15.492	23.778	19.634	25.803	19.289	25.206	19.734	25.159
Design-greedy-B	2.056	3.952	12.539	12.372	16.129	15.218	15.428	13.777	15.914	15.651
DAVED single (adapted)	0.809	1.661	3.008	4.822	10.251	5.383	10.325	4.995	11.424	4.477
DAVED multi (adapted)	4.507	15.311	14.918	**24.110**	18.475	27.834	18.523	27.905	18.423	27.879
Data-OOB	-1.184	0.310	2.942	3.010	10.954	3.856	9.342	2.814	9.676	3.098
Random-B	-1.114	0.508	6.401	3.336	10.365	4.197	9.097	2.721	9.307	3.704

Means are over 20 trials after averaging the three test tasks within each market. All baseline' cost budget C equals to cost of GRADES.

Each trial in Table 1 proceeds as follows. We run GRADES with $\delta = 1$ and capacity k, and let C be the total cost of the sellers it selects. Each baseline then selects at most k sellers with total cost at most C. Greedy runs the allocation step of GRADES with $\delta = 0$, and stops once it has selected k sellers or its next selection would exceed C. Data-OOB and Random add sellers under the same two limits. For every method, we fit ridge regression with penalty η on the selected sellers' data and evaluate it on the test set.

$$\check{\theta}_{S,\eta} = \boldsymbol{I}_\eta(S)^{-1}\boldsymbol{X}_S^\top \boldsymbol{y}_S, \qquad \bar{\ell}_\eta(S) = \frac{1}{m}\|\boldsymbol{X}^{\text{test}}\check{\theta}_{S,\eta} - \boldsymbol{y}^{\text{test}}\|_2^2.$$

We fix $\eta = 1$, sweep $k = 1, \ldots, 30$, and use $k = 25$ for endpoint tables. Three task outcomes are averaged within each market before calculating uncertainty.

4.2 COMPETITIVE ACQUISITION PERFORMANCE AND SELECTION TIME

Table 1 compares normalized design welfare $\widetilde{W}_1$ at $k = 25$ and $\eta = 1$. GRADES has higher mean design welfare than adapted multi-step DAVED in nine of the ten dataset–cost conditions. The exception is MIMIC under squared costs, where the means are 24.092 for GRADES and 24.110 for DAVED. GRADES also has higher mean design welfare than Greedy-B ($\delta = 0$) and Design-greedy-B in all ten conditions.

Figure 1 shows the impact the capacity k on the Gaussian setting. GRADES's MSE on the testing data generally decreases as k increases and follows a similar trajectory to adapted multi-step DAVED's. Its design welfare levels off at larger capacities under square-root costs, whereas it continues to increase over the measured range under squared costs.

Figure 2 compares running times in the Gaussian setting. GRADES is more than 100 times faster than adapted multi-step DAVED in every tested setting. Section D.3 provide additional details

Under the stated forward resource protocol, GRADES provides favorable cost-sensitive design allocation relative to adapted multi-step DAVED and substantially lower measured selection overhead, while realized prediction remains task- and cost-regime-dependent. We therefore keep the stage-1 design objective in the main text and defer the full MSE, empirical V, $V - C$, and surrogate-versus-realized-value analysis to Appendix D.

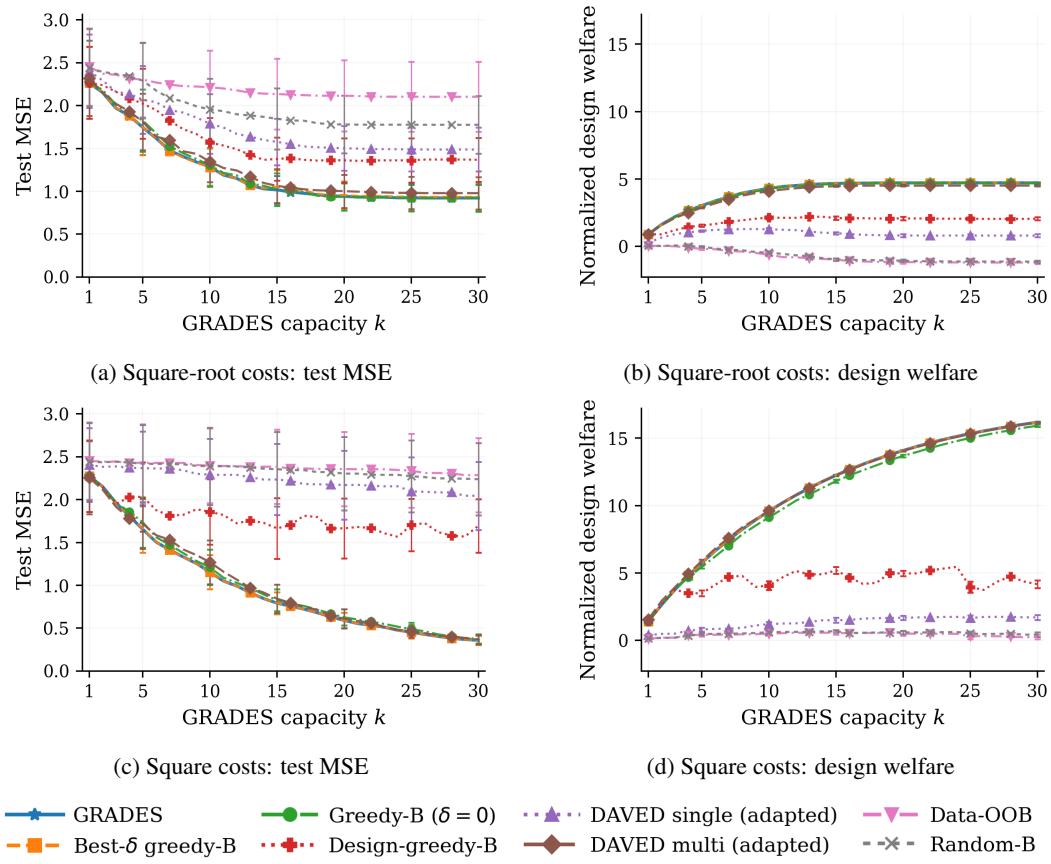

(a) Square-root costs: test MSE

(b) Square-root costs: design welfare

(c) Square costs: test MSE

(d) Square costs: design welfare

Figure 1: Gaussian acquisition under GRADES-defined cost caps, $\eta = 1$. Lines connect all 30 measured capacities; bars show one standard error over 20 trials after task averaging, at $k = 1$ and multiples of five. The axis is GRADES's capacity, not a common count; expenditures may differ. Design welfare is $\widetilde{W}_1$, and MSE is unscaled. Method styles are shared by all acquisition and timing figures.

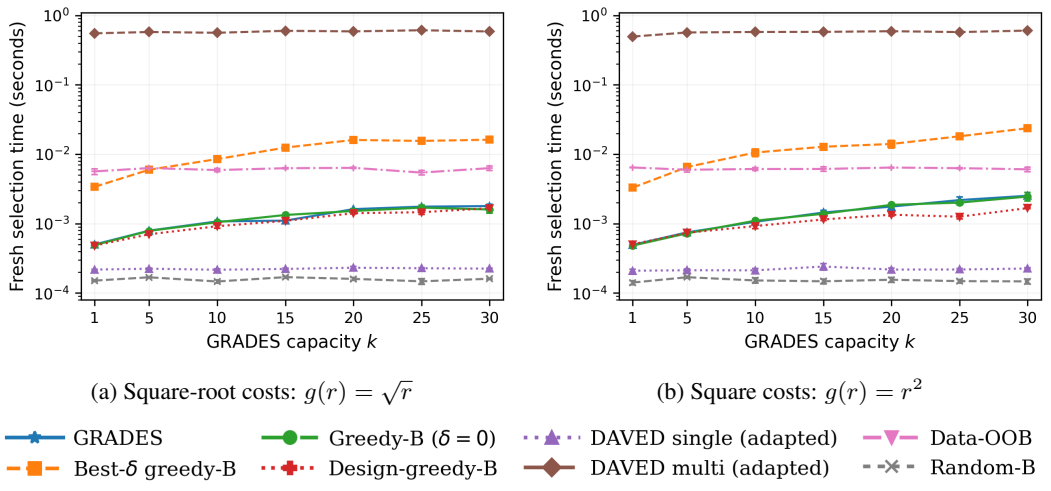

(a) Square-root costs: $g(r) = \sqrt{r}$

(b) Square costs: $g(r) = r^2$

Figure 2: Gaussian fresh selection runtime at $\eta = 1$ under GRADES-defined acquisition-cost caps. The time axis is logarithmic. Points show mean technical-repeat medians, with one standard error over five trials. The horizontal axis is GRADES's capacity k. DAVED rows are shared-regularizer adaptations.

AI USE STATEMENT

Large language models (LLMs) were used in this work as general-purpose assistants. Specifically, they supported adjusting phrasing for clarity, polishing grammar, shortening and restructuring text, and reformatting LaTeX content, including tables. LLMs were also used to assist with code editing and debugging, and with the visualization and presentation of experimental results, including generating and refining plotting code, figure formatting, labels, and layouts. All AI-assisted outputs were reviewed and verified by the authors. The research questions, conceptual framework, theoretical contributions, experimental results, and final conclusions are those of the authors. The authors take full responsibility for the final content of this work.

ETHICS STATEMENT

(This section is **recommended** and does not count toward the page limit.) This work uses existing benchmark datasets and does not involve new human-subject data collection. Potential concerns related to privacy, bias, and representation should be considered in real-world deployment.

REPRODUCIBILITY STATEMENT

(This section is **recommended** and does not count toward the page limit.)

Complete proofs, experimental settings, and implementation details are provided in the appendix and supplementary materials. Code for reproducing the reported results will be released upon acceptance.

REFERENCES

ABC News. Egocentric training's gig workers are helping robots learn how to do your dishes. ABC News (Australia), September 20, 2026.

Jacob D Abernethy and Rafael Frongillo. A collaborative mechanism for crowdsourcing prediction problems. *Advances in neural information processing systems*, 24, 2011.

Advanced Research Projects Agency for Health. ARPA-H announces medical imaging data partnership with FDA. https://arpa-h.gov/news-and-events/arpa-h-announces-medical-imaging-data-partnership-fda, 2024. Accessed September 2026.

Anish Agarwal, Munther Dahleh, and Tuhin Sarkar. A marketplace for data: An algorithmic solution. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pp. 701–726, 2019.

Nivasini Ananthakrishnan, Alireza Fallah, and Michael I. Jordan. Signaling in data markets via free samples, 2026. URL https://arxiv.org/abs/2602.16919.

Aaron Archer and Éva Tardos. Truthful mechanisms for one-parameter agents. In *Proceedings 42nd IEEE Symposium on Foundations of Computer Science*, pp. 482–491. IEEE, 2001.

Kenneth Joseph Arrow. Economic welfare and the allocation of resources for invention. In *Readings in industrial economics: Volume two: Private enterprise and state intervention*, pp. 219–236. Springer, 1962.

Andrew An Bian, Joachim M Buhmann, Andreas Krause, and Sebastian Tschiatschek. Guarantees for greedy maximization of non-submodular functions with applications. In *International conference on machine learning*, pp. 498–507. PMLR, 2017.

Bloomberg News. YouTubers are selling their unused video footage to AI companies. Bloomberg, January 10, 2025. https://news.bloomberglaw.com/ip-law/youtubers-are-selling-their-unused-video-footage-to-ai-companies, 2025.

Raul Castro Fernandez. Data-sharing markets: model, protocol, and algorithms to incentivize the formation of data-sharing consortia. *Proceedings of the ACM on Management of Data*, 1(2):1–25, 2023.

Kathryn Chaloner and Isabella Verdinelli. Bayesian experimental design: A review. *Statistical Science*, 10(3):273–304, 1995a. ISSN 08834237. URL http://www.jstor.org/stable/2246015.

Kathryn Chaloner and Isabella Verdinelli. Bayesian experimental design: A review. *Statistical science*, pp. 273–304, 1995b.

Luiz Chamon and Alejandro Ribeiro. Approximate supermodularity bounds for experimental design. *Advances in Neural Information Processing Systems*, 30, 2017.

Yiling Chen and Bo Waggoner. Informational substitutes, 2017. URL https://arxiv.org/abs/1703.08636.

Yiling Chen, Nicole Immorlica, Brendan Lucier, Vasilis Syrgkanis, and Juba Ziani. Optimal data acquisition for statistical estimation. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pp. 27–44, 2018.

Shuang Cui, He Huang, Yu-e Sun, and Chen Xue. Budget-feasible mechanisms for submodular welfare maximization in procurement auctions. *arXiv preprint arXiv:2605.00411*, 2026.

Abhimanyu Das and David Kempe. Approximate submodularity and its applications: Subset selection, sparse approximation and dictionary selection. *Journal of Machine Learning Research*, 19 (3):1–34, 2018.

Yuan Deng, Amin Karbasi, Vahab Mirrokni, Renato Paes Leme, Grigoris Velegkas, and Song Zuo. Procurement auctions via approximately optimal submodular optimization. In *Proceedings of the 42nd International Conference on Machine Learning*, ICML'25. JMLR.org, 2025.

Shahar Dobzinski and Jan Vondrák. On the hardness of welfare maximization in combinatorial auctions with submodular valuations. *arXiv preprint arXiv:1202.2792*, 2012.

Alon Eden, Kira Goldner, Eldar Kerner, and Thodoris Tsilivis. From welfare to utility: Generalized objectives in budget-feasible procurement. *arXiv preprint arXiv:2607.00101*, 2026.

Uriel Feige. A threshold of ln n for approximating set cover. *Journal of the ACM (JACM)*, 45(4): 634–652, 1998.

Mengxue Geng, Shufang Gong, Bin Liu, and Weili Wu. Bicriteria algorithms for maximizing the difference between submodular function and linear function under noise. In *International Conference on Algorithmic Applications in Management*, pp. 133–143. Springer, 2022.

Amirata Ghorbani and James Zou. Data shapley: Equitable valuation of data for machine learning. In *International conference on machine learning*, pp. 2242–2251. PMLR, 2019.

Arpita Ghosh, Katrina Ligett, Aaron Roth, and Grant Schoenebeck. Buying private data without verification. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pp. 931–948, 2014.

W Brier Glenn et al. Verification of forecasts expressed in terms of probability. *Monthly weather review*, 78(1):1–3, 1950.

Tilmann Gneiting and Adrian E Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American statistical Association*, 102(477):359–378, 2007.

Shufang Gong, Bin Liu, Mengxue Geng, and Qizhi Fang. Algorithms for maximizing monotone submodular function minus modular function under noise. *Journal of Combinatorial Optimization*, 45(4):96, 2023.

Felix Gräßer, Surya Kallumadi, Hagen Malberg, and Sebastian Zaunseder. Aspect-based sentiment analysis of drug reviews applying cross-domain and cross-data learning. In *Proceedings of the 2018 International Conference on Digital Health*, pp. 121–125, New York, NY, USA, 2018. Association for Computing Machinery. doi: 10.1145/3194658.3194677. URL https://doi.org/10.1145/3194658.3194677.

Matthew Groh, Caleb Harris, Luis Soenksen, Felix Lau, Rachel Han, Aerin Kim, Arash Koochek, and Omar Badri. Evaluating deep neural networks trained on clinical images in dermatology with the fitzpatrick 17k dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pp. 1820–1828, June 2021.

Safwan S. Halabi, Luciano M. Prevedello, Jayashree Kalpathy-Cramer, Artem B. Mamonov, Alexander Bilbily, Mark Cicero, Ian Pan, Lucas Araújo Pereira, Rafael Teixeira Sousa, Nitamar Abdala, Felipe Campos Kitamura, Hans H. Thodberg, Leon Chen, George Shih, Katherine Andriole, Marc D. Kohli, Bradley J. Erickson, and Adam E. Flanders. The rsna pediatric bone age machine learning challenge. *Radiology*, 290(2):498–503, 2019. doi: 10.1148/radiol.2018180736. URL https://doi.org/10.1148/radiol.2018180736.

Chris Harshaw, Moran Feldman, Justin Ward, and Amin Karbasi. Submodular maximization beyond non-negativity: Guarantees, fast algorithms, and applications. In *International Conference on Machine Learning*, pp. 2634–2643. PMLR, 2019.

Thibaut Horel, Stratis Ioannidis, and S Muthukrishnan. Budget feasible mechanisms for experimental design. In *Latin American Symposium on Theoretical Informatics*, pp. 719–730. Springer, 2014.

Ronald A Howard. Information value theory. *IEEE Transactions on systems science and cybernetics*, 2(1):22–26, 1966.

Yuchen Hu, Martin J. Wainwright, and Stephen Bates. Buying data of unknown quality: Statistical information procurement auctions, 2026. URL https://arxiv.org/abs/2604.08821.

Alistair Johnson, Tom Pollard, and Roger Mark. MIMIC-III Clinical Database. *PhysioNet*, September 2016. doi: 10.13026/C2XW26. URL https://doi.org/10.13026/C2XW26. Version 1.4.

Feiyang Kang, Hoang Anh Just, Anit Kumar Sahu, and Ruoxi Jia. Performance scaling via optimal transport: Enabling data selection from partially revealed sources. *Advances in Neural Information Processing Systems*, 36:61341–61363, 2023.

Anna R Karlin and David Kempe. Beyond vcg: Frugality of truthful mechanisms. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)*, pp. 615–624. IEEE, 2005.

Pang Wei Koh and Percy Liang. Understanding black-box predictions via influence functions. In *International conference on machine learning*, pp. 1885–1894. Pmlr, 2017.

Yongchan Kwon and James Zou. Data-oob: Out-of-bag estimate as a simple and efficient data value. In *International conference on machine learning*, pp. 18135–18152. PMLR, 2023a.

Yongchan Kwon and James Zou. Data-OOB: Out-of-bag estimate as a simple and efficient data value. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 18135–18152. PMLR, 2023b. URL https://proceedings.mlr.press/v202/kwon23e.html.

Dennis V Lindley. On a measure of the information provided by an experiment. *The Annals of Mathematical Statistics*, 27(4):986–1005, 1956.

Charles Lu, Baihe Huang, Sai Praneeth Karimireddy, Praneeth Vepakomma, Michael I. Jordan, and Ramesh Raskar. Data acquisition via experimental design for data markets. In Amir Globersons, Lester Mackey, Danielle Belgrave, Angela Fan, Ulrich Paquet, Jakub M. Tomczak, and Cheng Zhang (eds.), *Advances in Neural Information Processing Systems 37: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*, 2024. URL http://papers.nips.cc/paper_files/paper/2024/hash/d5e8326bbec25e1c608787d24488521b-Abstract-Conference.html.

Medical Imaging and Data Resource Center. MIDRC: Medical imaging and data resource center. https://www.midrc.org, 2026. Accessed September 2026.

Nolan Miller, Paul Resnick, and Richard Zeckhauser. Eliciting informative feedback: The peer-prediction method. *Management Science*, 51(9):1359–1373, 2005.

Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.

Tom Pollard, Benjamin E. Moody, Li-wei H. Lehman, Brian J. Gow, Chrystinne Fernandes, Chen Xie, Alistair Johnson, Roger G. Mark, and Thomas Heldt. PhysioNet as a global platform for biomedical research. *Nature Health*, 1(8):792–795, August 2026. ISSN 3005-0693. doi: 10.1038/s44360-026-00096-z. URL https://doi.org/10.1038/s44360-026-00096-z.

Drazen Prelec. A bayesian truth serum for subjective data. *science*, 306(5695):462–466, 2004.

Howard Raiffa and Robert Schlaifer. *Applied statistical decision theory*. John Wiley & Sons, 2000.

Aaron Roth and Grant Schoenebeck. Conducting truthful surveys, cheaply. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pp. 826–843, 2012.

Leonard J Savage. Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association*, 66(336):783–801, 1971.

Yaron Singer. Budget feasible mechanisms. In *2010 IEEE 51st Annual Symposium on foundations of computer science*, pp. 765–774. IEEE, 2010.

Maxim Sviridenko, Jan Vondrák, and Justin Ward. Optimal approximation for submodular and supermodular optimization with bounded curvature. *Mathematics of Operations Research*, 42(4): 1197–1218, 2017.

William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 16(1):8–37, 1961.

Jiachen T Wang and Ruoxi Jia. Data banzhaf: A robust data valuation framework for machine learning. In *International conference on artificial intelligence and statistics*, pp. 6388–6421. PMLR, 2023.

Yijing Wang, Yicheng Xu, and Xiaoguang Yang. On maximizing the difference between an approximately submodular function and a linear function subject to a matroid constraint. In *International Conference on Combinatorial Optimization and Applications*, pp. 75–85. Springer, 2021.

Martin Weitzman. *Optimal search for the best alternative*, volume 78. Department of Energy, 1978.

Xinyi Xu, Zhaoxuan Wu, Chuan Sheng Foo, and Bryan Kian Hsiang Low. Validation free and replication robust volume-based data valuation. *Advances in Neural Information Processing Systems*, 34:10837–10848, 2021.

A PRELIMINARIES

We write $\phi_t \coloneqq (1-\delta/k)^{k-t}$ for the distortion factor, and, for marginal gains, we carry the function into the bar notation: for any set function g and disjoint $T, L \subseteq [N]$,

$$g(T \mid L) \coloneqq g(L \cup T) - g(L),$$

so that $\tilde{v}(i \mid L)$ is the surrogate gain of seller i at L and $v(i \mid L)$ the gain of the true valuation. A proper scoring rule is written PS, as in Section 3.2, to keep S for sets of sellers. Four symbols carry a fixed meaning paper-wide. π is the allocation rule of Section 2.1 and never a prior. Priors and posteriors over θ are μ, with μ_0 the prior induced by D_0. $\boldsymbol{I}$ is an information matrix, $\boldsymbol{I}(S)$ in Eq. (3) and its Fisher analogue in Eq. (7). Mutual information is written MI, not I. The two accents are likewise disjoint: a hat marks a report ($\hat{c}_i$, $\hat{\boldsymbol{y}}_i$, $\hat{D}_S$), a check an estimate ($\check{\theta}_S$, $\check{\theta}_0$, $\check{p}$). A symbol carrying both, such as $\check{\theta}(\hat{\boldsymbol{y}}_j)$, is an estimate read off a delivery.

A.1 MATRIX IDENTITIES

Lemma A.1 (Woodbury identity)**.** *Let $\boldsymbol{M} \in \mathbb{R}^{d\times d}$ be invertible and $\boldsymbol{U} \in \mathbb{R}^{d\times n}$. Then $\boldsymbol{M} + \boldsymbol{U}\boldsymbol{U}^\top$ is invertible if and only if $\mathbb{I}_n + \boldsymbol{U}^\top \boldsymbol{M}^{-1}\boldsymbol{U}$ is, in which case*

$$(\boldsymbol{M} + \boldsymbol{U}\boldsymbol{U}^\top)^{-1} = \boldsymbol{M}^{-1} - \boldsymbol{M}^{-1}\boldsymbol{U}(\mathbb{I}_n + \boldsymbol{U}^\top \boldsymbol{M}^{-1}\boldsymbol{U})^{-1}\boldsymbol{U}^\top \boldsymbol{M}^{-1}.$$

Corollary A.2 (Marginal gain of the V-optimal surrogate)**.** *For disjoint $L, T \subseteq [N]$, let $\boldsymbol{P} = \boldsymbol{I}(L)^{-1}$, let $\boldsymbol{X}_T$ stack the blocks $\boldsymbol{X}_i$ for $i \in T$, and let $n_T = \sum_{i\in T} n_i$. The V-optimal surrogate Eq.* (3) *satisfies*

$$\tilde{v}(L \cup T) - \tilde{v}(L) = \sigma^2 \operatorname{tr}\Big(\boldsymbol{X}_T \boldsymbol{P}\boldsymbol{A}\boldsymbol{P}\boldsymbol{X}_T^\top \big(\mathbb{I}_{n_T} + \boldsymbol{X}_T \boldsymbol{P}\boldsymbol{X}_T^\top\big)^{-1}\Big).$$

Proof. Apply Theorem A.1 with $\boldsymbol{M} = \boldsymbol{I}(L)$, which is invertible because $\boldsymbol{X}_0$ has full rank, and $\boldsymbol{U} = \boldsymbol{X}_T^\top$, so that $\boldsymbol{M} + \boldsymbol{U}\boldsymbol{U}^\top = \boldsymbol{I}(L \cup T)$:

$$\boldsymbol{I}(L \cup T)^{-1} = \boldsymbol{P} - \boldsymbol{P}\boldsymbol{X}_T^\top \big(\mathbb{I}_{n_T} + \boldsymbol{X}_T \boldsymbol{P}\boldsymbol{X}_T^\top\big)^{-1} \boldsymbol{X}_T \boldsymbol{P}.$$

Substituting into $\tilde{v}(L\cup T) - \tilde{v}(L) = \sigma^2[\operatorname{tr}(\boldsymbol{A}\boldsymbol{P}) - \operatorname{tr}(\boldsymbol{A}\boldsymbol{I}(L\cup T)^{-1})]$ and using the cyclic invariance of the trace yields the claim. □

A.2 SOLUTION CONCEPTS

We now state formally the solution concepts introduced in Section 2.1. Seller i's utility under the cost report profile $\hat{\boldsymbol{c}}$, the feature profile $\boldsymbol{X}$ and the label-delivery profile $\hat{\boldsymbol{y}}$ is

$$u_i(\hat{\boldsymbol{c}}, \boldsymbol{X}, \hat{\boldsymbol{y}}; c_i) = \big(p_i(\hat{\boldsymbol{c}}, \boldsymbol{X}, \hat{\boldsymbol{y}}_S) - c_i\big) \cdot \mathbf{1}[i \in \pi(\hat{\boldsymbol{c}}, \boldsymbol{X})].$$

Definition A.3 (Cost-DSIC in expectation)**.** *A mechanism $(\pi, \boldsymbol{p})$ is* cost-DSIC *(in expectation over the labels) if for every seller i, every true cost c_i, every misreport $\hat{c}_i$, every profile $\hat{c}_{-i}$ of others' reports, and every feature profile $\boldsymbol{X}$,*

$$\mathbb{E}\big[u_i\big((c_i, \hat{c}_{-i}), \boldsymbol{X}, \hat{\boldsymbol{y}}; c_i\big)\big] \geq \mathbb{E}\big[u_i\big((\hat{c}_i, \hat{c}_{-i}), \boldsymbol{X}, \hat{\boldsymbol{y}}; c_i\big)\big],$$

where the expectation is over the label.

A mechanism is individually rational (IR) *if*

$$\mathbb{E}\big[u_i\big((c_i, \hat{c}_{-i}), \boldsymbol{X}, \hat{\boldsymbol{y}}; c_i\big)\big] \geq 0,$$

where the expectation is over the label.

The ϵ-cost incentive of Section 3 is these two definitions with one slack for both.

Definition A.4 (Label-BNE truthfulness)**.** *Suppose all sellers report costs truthfully (Theorem A.3). A profile of label-delivery strategies $(\hat{\boldsymbol{y}}_i^*)_{i\in[N]}$ is a* Bayes–Nash equilibrium (BNE) *for label reporting if, for every seller i, every realization D_i, and every alternative report $\hat{\boldsymbol{y}}_i'$,*

$$\mathbb{E}\left[u_i\big(\boldsymbol{c}, \boldsymbol{X}, (\hat{\boldsymbol{y}}_i^*, \hat{\boldsymbol{y}}_{-i}^*); c_i\big)\middle| D_i\right] \geq \mathbb{E}\left[u_i\big(\boldsymbol{c}, \boldsymbol{X}, (\hat{\boldsymbol{y}}_i', \hat{\boldsymbol{y}}_{-i}^*); c_i\big)\middle| D_i\right],$$

where the expectation is taken over θ under seller i's belief. The mechanism is label-BNE truthful *if $\hat{\boldsymbol{y}}_i^* = \boldsymbol{y}_i$ for all i is a BNE.*

Together, our solution concept has two stages: cost reports form a dominant-strategy equilibrium in stage 1 and, conditional on the resulting selection, label deliveries form a Bayes–Nash equilibrium in stage 3.

A.3 Procurement auctions and submodular maximization

The greedy stage of Section 2.2 is a *per-round scoring auction*: in round t it assigns each remaining seller i a score $s_i^t(\hat{c}_i)$, computed from her own bid and the current candidate set S_{t-1}, selects the highest-scoring seller if her score is positive, and otherwise skips the round. Deng et al. (2025) convert greedy algorithms of this form into sealed-bid auctions. We restate their condition and conversion in our notation.

Definition A.5 (Per-round scores, Deng et al., 2025, Assumption 4.2)**.** *For every round t, every seller i, and every candidate set S_{t-1}, the score $s_i^t(\hat{c}_i)$ is*

1. *non-increasing in seller i's own bid $\hat{c}_i$,*
2. *negative whenever $\hat{c}_i$ is large enough, and*
3. *independent of the other sellers' bids.*

Algorithm 2 of Deng et al. Deng et al. (2025)
Input: the sellers $[N]$, a bid profile $\hat{\boldsymbol{c}}$, and a meta-algorithm $\mathcal{A} = (G)$ with scoring rule G.
Output: a set of winners $S^\star$ and a payment vector $\boldsymbol{p}$.

1. **Seed.** Generate a random seed r if $\mathcal{A}$ is randomized; otherwise set $r = 0$.
2. **Allocate.** $S^\star \leftarrow \mathcal{A}([N], \hat{\boldsymbol{c}}, r)$.
3. **Pay.** For each seller $i \in S^\star$:
   (a) Run $\mathcal{A}\big((\infty, \hat{\boldsymbol{c}}_{-i}), r\big)$ and record its intermediate solutions $S_0, S_1, \ldots, S_N$.
   (b) Set $p_i \leftarrow 0$. For $t = 1, \ldots, N$, update

$$p_i \leftarrow \max\Big(p_i,\ \sup\big\{z \geq 0 : i = \operatorname*{arg\,max}_{\ell \notin S_{t-1}} G\big(\ell, S_{t-1}, (z, \hat{\boldsymbol{c}}_{-i}), t, r\big), \\ G\big(i, S_{t-1}, (z, \hat{\boldsymbol{c}}_{-i}), t, r\big) > 0\big\}\Big),$$

   with the convention $\sup \emptyset = 0$.
4. **Return** $S^\star$ and $\boldsymbol{p}$.

Theorem A.6 (Greedy-to-auction conversion, Deng et al., 2025, Theorem 4.3)**.** *Under Theorem A.5, the above mechanism satisfies IC and IR.*

B Proofs and details for Section 2

B.1 Proof of Theorem 2.2

Proof of Theorem 2.2. To apply Theorem 3.2, we show Eqs. (2) and (3) satisfy Theorem 3.1 with $\rho = 1$ and γ defined in the statement.

First, the surrogate $\tilde{v}$ of Eq. (3) is computed from the features. Adding sellers adds the positive semidefinite terms $\boldsymbol{X}_i^\top \boldsymbol{X}_i$ to $\boldsymbol{I}(S)$, and $\boldsymbol{M} \mapsto \operatorname{tr}(\boldsymbol{A}\boldsymbol{M}^{-1})$ is Loewner-decreasing on positive definite matrices for positive semidefinite $\boldsymbol{A}$, so $\tilde{v}$ is monotone with $\tilde{v}(\emptyset) = 0$, hence non-negative.

Second, as the principal uses f_{LS}, the estimator $\check{\theta}_S = (\boldsymbol{X}_S^\top \boldsymbol{X}_S)^{-1} \boldsymbol{X}_S^\top \boldsymbol{y}_S$ is unbiased with covariance

$$\mathbb{E}\big[(\check{\theta}_S - \theta)(\check{\theta}_S - \theta)^\top\big] = \sigma^2 \boldsymbol{I}(S)^{-1}, \tag{9}$$

so the risk $\ell(S)$ of Section 2.1 evaluates in one line. Writing $y_q^{\mathrm{test}} = (x_q^{\mathrm{test}})^\top \theta + \epsilon_q$ with test noise independent of D_S, the cross term vanishes and

$$\ell(S) = \sigma^2 + \mathbb{E}\big\|\check{\theta}_S - \theta\big\|_{\boldsymbol{A}}^2 = \sigma^2 + \operatorname{tr}\Big(\boldsymbol{A}\mathbb{E}\big[(\check{\theta}_S - \theta)(\check{\theta}_S - \theta)^\top\big]\Big) = \sigma^2\big(1 + \operatorname{tr}\big(\boldsymbol{A}\boldsymbol{I}(S)^{-1}\big)\big),$$

and therefore

$$v(S) = \ell(\emptyset) - \ell(S) = \sigma^2\big[\mathrm{tr}(\boldsymbol{A}\boldsymbol{I}(\emptyset)^{-1}) - \mathrm{tr}(\boldsymbol{A}\boldsymbol{I}(S)^{-1})\big] = \tilde{v}(S).$$

Third, for the submodularity ratio, Theorem 2.5 gives $\gamma \geq \min_{L,S}(1 + \max_{j\in S} h_j)^{-1}$, where $h_j = \lambda_{\max}(\boldsymbol{X}_j\boldsymbol{P}\boldsymbol{X}_j^\top)$ is the block leverage of seller j under $\boldsymbol{P} = \boldsymbol{I}(L)^{-1}$. Since $\boldsymbol{I}(L) \succeq \boldsymbol{I}(\emptyset)$ for every $L \subseteq [N]$ (acquiring more sellers only adds the positive semidefinite terms $\boldsymbol{X}_i^\top\boldsymbol{X}_i$), we have $\boldsymbol{P} \preceq \boldsymbol{I}(\emptyset)^{-1}$, and hence

$$h_j = \lambda_{\max}\big(\boldsymbol{X}_j\boldsymbol{P}\boldsymbol{X}_j^\top\big) \leq \lambda_{\max}\big(\boldsymbol{X}_j\boldsymbol{I}(\emptyset)^{-1}\boldsymbol{X}_j^\top\big) \leq h_{\max}.$$

This yields $\gamma \geq 1/(1 + h_{\max})$, computable from the features before the auction runs.

Finally, for the running time, the greedy stage runs k rounds. Each round evaluates the marginal gain $\tilde{v}(i \mid S)$ for $O(N)$ candidate sellers. Computing $\boldsymbol{I}(S \cup \{i\}) = \boldsymbol{I}(S) + \boldsymbol{X}_i^\top\boldsymbol{X}_i$ costs $O(d^2 n_i)$, and inverting it costs $O(d^3)$, so each marginal gain takes $O(d^3 + d^2 n_{\max})$ with $n_{\max} \coloneqq \max_i n_i$. The payment stage re-runs the greedy stage once per selected seller with her bid at ∞. In round t of that run, the supremum in the payment step of Section A.3 has the closed form

$$\phi_t\tilde{v}(i \mid S_{t-1}) - \max\Big\{0, \max_{\ell \notin S_{t-1}\cup\{i\}}\big[\phi_t\tilde{v}(\ell \mid S_{t-1}) - \hat{c}_\ell\big]\Big\}, \tag{10}$$

or 0 if this is negative, which reuses the gains already computed. Therefore, the total runtime is $O(Nk^2(d^3 + d^2 n_{\max}))$. □

B.2 Proof of Theorem 2.5

Proof. Fix a pair of disjoint subsets L (the prefix) and S (the residual). Write $g_j \coloneqq \tilde{v}(L\cup\{j\}) - \tilde{v}(L)$ and $g_S \coloneqq \tilde{v}(L \cup S) - \tilde{v}(L)$. The factor σ^2 of Eq. (3) is common to both and cancels in the comparison, so we set $\sigma^2 = 1$ for this proof. By the definition of γ, it suffices to show

$$\sum_{j\in S} g_j \geq \frac{g_S}{1 + \max_{j\in S} h_j}.$$

Let $\boldsymbol{G} \coloneqq \boldsymbol{X}_S\boldsymbol{P}\boldsymbol{X}_S^\top$ and $\boldsymbol{Q} \coloneqq \boldsymbol{X}_S\boldsymbol{P}\boldsymbol{A}\boldsymbol{P}\boldsymbol{X}_S^\top$, both positive semidefinite in $\mathbb{R}^{n_S\times n_S}$ with $n_S = \sum_{j\in S} n_j$. The block structure of $\boldsymbol{X}_S$ gives $\boldsymbol{G}$ diagonal blocks $\boldsymbol{G}_{jj} = \boldsymbol{X}_j\boldsymbol{P}\boldsymbol{X}_j^\top$ and $\boldsymbol{Q}$ diagonal blocks $\boldsymbol{Q}_{jj} = \boldsymbol{X}_j\boldsymbol{P}\boldsymbol{A}\boldsymbol{P}\boldsymbol{X}_j^\top$, with $\lambda_{\max}(\boldsymbol{G}_{jj}) = h_j$. The marginal-gain identity of Theorem A.2 gives

$$g_j = \mathrm{tr}\big(\boldsymbol{Q}_{jj}(\mathbb{I}_{n_j} + \boldsymbol{G}_{jj})^{-1}\big) \text{ and } g_S = \mathrm{tr}\big(\boldsymbol{Q}(\mathbb{I}_{n_S} + \boldsymbol{G})^{-1}\big).$$

We lower bound the left-hand side. By the definition of h_j, each block satisfies $\boldsymbol{G}_{jj} \preceq h_j\mathbb{I}_{n_j}$, so $(\mathbb{I}_{n_j} + \boldsymbol{G}_{jj})^{-1} \succeq (1 + h_j)^{-1}\mathbb{I}_{n_j}$. Thus,

$$\sum_{j\in S} g_j \geq \sum_{j\in S} \frac{\mathrm{tr}(\boldsymbol{Q}_{jj})}{1 + h_j} \geq \frac{\sum_{j\in S}\mathrm{tr}(\boldsymbol{Q}_{jj})}{1 + \max_{j\in S} h_j} = \frac{\mathrm{tr}(\boldsymbol{Q})}{1 + \max_{j\in S} h_j},$$

the last step because the trace of $\boldsymbol{Q}$ is the sum of the traces of its diagonal blocks.

We upper bound g_S. Since $\boldsymbol{G} \succeq 0$ we have $(\mathbb{I}_{n_S} + \boldsymbol{G})^{-1} \preceq \mathbb{I}$, and $\boldsymbol{Q} \succeq 0$, so

$$g_S = \mathrm{tr}\big(\boldsymbol{Q}(\mathbb{I}_{n_S} + \boldsymbol{G})^{-1}\big) \leq \mathrm{tr}(\boldsymbol{Q}).$$

This step uses no condition on $\boldsymbol{A}$ beyond positive semidefiniteness. Combining the two bounds yields $\sum_{j\in S} g_j \geq \mathrm{tr}(\boldsymbol{Q})/(1 + \max_{j\in S} h_j) \geq g_S/(1 + \max_{j\in S} h_j)$. □

Claim 10 of Harshaw et al. (2019) proves the same shape of bound for Bayesian A-optimal design, $\gamma \geq (1 + s^2\lambda_{\max}(\Sigma)/\sigma^2)^{-1}$ with $s = \max_i \|x_i\|$, by the same Woodbury idea. Theorem 2.5 replaces the product of a norm and a spectral norm by the block leverage h_j, which is tighter, and covers multi-row sellers and every positive semidefinite $\boldsymbol{A}$.

B.3 Proof of Theorem 2.3

Theorem 2.3 has two clauses, proved in turn below: no polynomial-time $(1-1/e+\epsilon)$-approximation of the surrogate $\tilde{v}$ itself, and, for every $\beta \in (0,1]$, no polynomial-time $(1-e^{-\beta}+\epsilon, \beta)$-approximation of the welfare. Both run on one family of single-record instances.

For the first clause we study the *value problem*: given an instance $(\boldsymbol{X}_0, \{\boldsymbol{X}_i\}_{i\in[N]}, \boldsymbol{X}^{\text{test}}, \sigma^2, k)$ of Section 2.1, with $\boldsymbol{X}_0$ and $\boldsymbol{X}^{\text{test}}$ of full rank d, find $S \subseteq [N]$ with $|S| \le k$ maximizing $\tilde{v}(S)$ in Eq. (3). An algorithm is an *α-approximation* if it always returns such an S with $\tilde{v}(S) \ge \alpha \max_{S'\subseteq[N]:|S'|\le k} \tilde{v}(S')$. All instances below have rational entries, so $\tilde{v}(S)$ can be evaluated exactly in polynomial time.

The reduction starts from maximum k-coverage. For sets $T_1, \dots, T_N \subseteq [d]$ and $S \subseteq [N]$, write $\kappa(S) \coloneqq \left|\bigcup_{i\in S} T_i\right|$ for the number of covered elements. For $T \subseteq [d]$, $\mathbf{1}_T \in \{0,1\}^d$ is its indicator vector and $\mathbf{1} \coloneqq \mathbf{1}_{[d]}$.

Lemma B.1 (Gap maximum k-coverage (Feige, 1998))**.** *For every constant $\epsilon > 0$ it is NP-hard, given $T_1, \dots, T_N \subseteq [d]$ and $k \in \mathbb{N}$, to distinguish*

(YES) some k pairwise disjoint sets among $T_1, \dots, T_N$ have union $[d]$, from

(NO) $\kappa(S) \le (1 - e^{-|S|/k} + \epsilon)d$ for every $S \subseteq [N]$ with $|S| \le k$, in particular $\kappa(S) \le (1 - 1/e + \epsilon)d$.

Feige's reduction produces YES instances covered by k disjoint sets. See Dobzinski & Vondrák (2012, Section 4.2) for this form, whose NO bound reads $1-(1-1/k)^\ell + \epsilon$ for any $\ell \le k$ sets. Since $0 \le e^{-1/k} - (1-1/k) \le 1/(2k^2)$ and both lie in $[0,1]$, $e^{-\ell/k} - (1-1/k)^\ell \le \ell/(2k^2) \le 1/k$ for $\ell \le k$, so applying it with $\epsilon/2$ gives Theorem B.1 on instances with $k \ge 2/\epsilon$. The remaining instances are solved exactly by enumerating the $N^{O(1/\epsilon)}$ families of at most k sets.

The surrogate is not a coverage function, because $\boldsymbol{I}(\emptyset)$ regularizes it. The next lemma says that it is nevertheless sandwiched by one, which is what preserves Feige's gap while $\boldsymbol{X}_0$ and $\boldsymbol{X}^{\text{test}}$ keep full rank.

Lemma B.2 (Soft projection)**.** *Suppose $\boldsymbol{I}(\emptyset) = \lambda \mathbb{I}_d$ for some $\lambda > 0$. Then*

$$\tilde{v}(S) = \frac{\sigma^2}{\lambda} \operatorname{tr}\big(\boldsymbol{A}\boldsymbol{Q}_S\big),$$
$$\boldsymbol{Q}_S \coloneqq \mathbb{I}_d - \lambda \boldsymbol{I}(S)^{-1}.$$

Moreover $0 \preceq \boldsymbol{Q}_S \preceq \Pi_S$, where Π_S is the orthogonal projection onto the span of the rows of $\{\boldsymbol{X}_i\}_{i\in S}$, and $\operatorname{tr}\boldsymbol{Q}_S \le \operatorname{tr}\Pi_S \le d$.

Proof. The identity is Eq. (3) rearranged: since $\boldsymbol{I}(\emptyset)^{-1} = \lambda^{-1}\mathbb{I}_d$,

$$\boldsymbol{I}(\emptyset)^{-1} - \boldsymbol{I}(S)^{-1} = \lambda^{-1}\big[\mathbb{I}_d - \lambda \boldsymbol{I}(S)^{-1}\big] = \lambda^{-1}\boldsymbol{Q}_S.$$

For the bounds, let R be the row span of $\{\boldsymbol{X}_i\}_{i\in S}$, so that Π_S projects onto R. On $R^\perp$ the seller rows contribute nothing, so $\boldsymbol{I}(S)$ acts as $\lambda\mathbb{I}_d$ there and $\boldsymbol{Q}_S$ vanishes. On R, $\boldsymbol{I}(S) \succeq \lambda\mathbb{I}_d$ gives $\lambda\boldsymbol{I}(S)^{-1} \preceq \mathbb{I}_d$, hence $0 \preceq \boldsymbol{Q}_S$. Both $\boldsymbol{Q}_S$ and Π_S are therefore supported on R, with $\boldsymbol{Q}_S \preceq \Pi_S$ there, and taking traces gives $\operatorname{tr}\boldsymbol{Q}_S \le \operatorname{tr}\Pi_S = \dim R \le d$. □

Proof of Theorem 2.3, the value clause. The claim for ϵ implies it for every larger ϵ, so assume $\epsilon \le 1/4$. Fix an integer $q \ge 2/\sqrt{\epsilon}$ and set $\eta \coloneqq 1/q$ and $\lambda \coloneqq \eta^2 \le \epsilon/4$. Given $T_1, \dots, T_N \subseteq [d]$ and k as in Theorem B.1 with $\epsilon/4$ in place of ϵ, let seller i hold the single record $\boldsymbol{X}_i \coloneqq \mathbf{1}_{T_i}^\top \in \{0,1\}^{1\times d}$, so that $n_i = 1$ for every i, and take

$$\boldsymbol{X}_0 \coloneqq \eta\mathbb{I}_d,$$
$$\boldsymbol{X}^{\text{test}} \coloneqq \begin{pmatrix} \mathbf{1}^\top \\ \eta\mathbb{I}_d \end{pmatrix} \in \mathbb{R}^{(d+1)\times d}.$$

Then $m = d+1$, $\boldsymbol{I}(\emptyset) = \boldsymbol{X}_0^\top\boldsymbol{X}_0 = \lambda\mathbb{I}_d$ and $\boldsymbol{A} = \frac{1}{d+1}\big(\mathbf{1}\mathbf{1}^\top + \lambda\mathbb{I}_d\big)$, and both $\boldsymbol{X}_0$ and $\boldsymbol{X}^{\text{test}}$ have full rank d. By Theorem B.2, the normalized value

$$\nu(S) \coloneqq \frac{\lambda(d+1)}{\sigma^2}\tilde{v}(S) = \mathbf{1}^\top\boldsymbol{Q}_S\mathbf{1} + \lambda \operatorname{tr}\boldsymbol{Q}_S$$

has the same maximizers and approximation ratios as $\tilde{v}$.

Upper bound. For every S, the rows $\boldsymbol{X}_i$, $i \in S$, are supported on $C_S := \bigcup_{i\in S} T_i$, hence so is $\boldsymbol{w} := \Pi_S \mathbf{1}$. Therefore $\|\boldsymbol{w}\|^2 = \langle \boldsymbol{w}, \mathbf{1}\rangle = \langle \boldsymbol{w}, \mathbf{1}_{C_S}\rangle \le \|\boldsymbol{w}\|\sqrt{\kappa(S)}$, and by Theorem B.2

$$\nu(S) \le \mathbf{1}^\top \Pi_S \mathbf{1} + \lambda \operatorname{tr} \Pi_S = \|\boldsymbol{w}\|^2 + \lambda \operatorname{tr} \Pi_S \le \kappa(S) + \lambda d. \tag{11}$$

Lower bound. Let O index k pairwise disjoint sets with union $[d]$, and write $s_i := |T_i|$. The rows $\boldsymbol{X}_i$, $i \in O$, are pairwise orthogonal with $\|\boldsymbol{X}_i\|^2 = s_i$ and $\sum_{i\in O} \boldsymbol{X}_i^\top = \mathbf{1}$, so $\boldsymbol{Q}_O = \sum_{i\in O: s_i \ge 1} (s_i + \lambda)^{-1} \boldsymbol{X}_i^\top \boldsymbol{X}_i$ and

$$\nu(O) \ge \mathbf{1}^\top \boldsymbol{Q}_O \mathbf{1} = \sum_{i\in O: s_i \ge 1} \frac{s_i^2}{s_i + \lambda} \ge \frac{1}{1+\lambda} \sum_{i\in O} s_i = \frac{d}{1+\lambda}. \tag{12}$$

Gap. In a YES instance, Eq. (12) gives $\max_{|S|\le k} \nu(S) \ge d/(1+\lambda)$, so a $(1-1/e+\epsilon)$-approximation returns S with

$$\nu(S) \ge \frac{1 - 1/e + \epsilon}{1+\lambda} d > \Big(1 - \frac{1}{e} + \frac{\epsilon}{2}\Big) d,$$

where the strict inequality uses $\lambda \le \epsilon/4$ and $1 - 1/e + \epsilon/2 \le 1$. In a NO instance, Eq. (11) gives $\nu(S) \le (1 - 1/e + \epsilon/4 + \lambda) d \le (1 - 1/e + \epsilon/2) d$ for every $|S| \le k$. Evaluating ν on the returned set would therefore decide the NP-hard problem of Theorem B.1. □

Proof of Theorem 2.3, the welfare clause. In the model of Theorem 2.1 we have $v = \tilde{v}$, so the welfare Eq. (1) is $W(S; \boldsymbol{c}) = \tilde{v}(S) - c(S)$ and the surrogate carries it. Fix $\beta \in (0, 1]$ and write α for the value factor of the bi-criteria pair, so that the claim is that no polynomial-time allocation rule is (α, β)-approximate with $\alpha \ge 1 - e^{-\beta} + \epsilon$. The claim for ϵ implies it for every larger ϵ, and it is vacuous for $\epsilon > 1$ because $\alpha \le 1$, so assume $\epsilon \le 1$ and set $\epsilon' := \epsilon/5$. Take an instance of Theorem B.1 with ϵ' in place of ϵ and build the single-record instance of the value clause from it, now with $\lambda \le \epsilon'$. Writing $s := \lambda(d+1)/\sigma^2$ and $g := s\tilde{v} = \nu$, Eqs. (11) and (12) read

$$\begin{aligned} &g(S) \le \kappa(S) + \epsilon' d \text{ for every } S \subseteq [N], \\ &g(O) \ge (1 - \epsilon') d, \end{aligned} \tag{13}$$

where, in a YES instance, O indexes k pairwise disjoint sets with union $[d]$. The second bound uses $1/(1+\lambda) \ge 1 - \lambda$. All sellers are given the same cost, so the hardness is not an artifact of a cost profile.

Fix a rational ρ with $e^{-\beta} \le \rho \le e^{-\beta} + \epsilon'$, and give every seller the cost $c_i := \rho d/(ks) > 0$. Then $G(S) := sW(S; \boldsymbol{c}) = g(S) - \rho d|S|/k$.

YES instances. If the allocation rule is (α, β)-approximate, comparing its output S with $S' = O$ and using Eq. (13), $|O| \le k$, $\alpha \le 1$ and $\beta\rho \le \beta e^{-\beta} + \epsilon'$ gives

$$\begin{aligned} G(S) &\ge \alpha g(O) - \beta\rho d \frac{|O|}{k} \ge \big[\alpha(1 - \epsilon') - \beta\rho\big] d \\ &\ge \big[\alpha - 2\epsilon' - \beta e^{-\beta}\big] d \ge \big[1 - e^{-\beta} - \beta e^{-\beta} + 3\epsilon'\big] d, \end{aligned}$$

where the last step uses $\alpha \ge 1 - e^{-\beta} + 5\epsilon'$.

NO instances. Every S with $|S| = \ell \le k$ satisfies, by Eq. (13), Theorem B.1 and $\rho \ge e^{-\beta}$,

$$\begin{aligned} G(S) &\le \Big(1 - e^{-\ell/k} - e^{-\beta} \frac{\ell}{k} + 2\epsilon'\Big) d \le \Big(\max_{x\in[0,1]} \big[1 - e^{-x} - e^{-\beta} x\big] + 2\epsilon'\Big) d \\ &= \big[1 - e^{-\beta} - \beta e^{-\beta} + 2\epsilon'\big] d, \end{aligned}$$

since $x \mapsto 1 - e^{-x} - e^{-\beta} x$ is concave with derivative $e^{-x} - e^{-\beta}$, which vanishes at $x = \beta \in (0, 1]$.

$G(S)$ is rational and computable exactly, so comparing it with a rational number strictly between the two bounds would decide the NP-hard problem of Theorem B.1. □

Example B.3 (The V-optimal surrogate is not submodular)**.** *Let $d = 2$, $\sigma^2 = 1$, $\boldsymbol{X}_0 = \mathbb{I}_2$ and $\boldsymbol{X}^{\text{test}} = \operatorname{diag}(3, 2)$, so that $m = 2$ and $\boldsymbol{A} = \operatorname{diag}(9/2, 2)$, and let two one-row sellers hold $\boldsymbol{X}_1 = (0, 1)$ and $\boldsymbol{X}_2 = (2, 1)$. Then buying seller* 2 raises *what seller* 1 *is worth at the margin, $\tilde{v}(1 \mid \{2\}) = 34/33 > 1 = \tilde{v}(1 \mid \emptyset)$, so $\tilde{v}$ is not submodular. It is nevertheless approximately submodular: at capacity $k = 2$ its submodularity ratio is $\gamma = 143/144$.*

B.4 Proof of Theorem 2.6

Proof of Theorem 2.6. Fix a winner j, let $\boldsymbol{P}_j \coloneqq \boldsymbol{I}(\{j\})^{-1}$, write $\check{\theta}_\emptyset$ for the least-squares fit on D_0 and $\check{\theta}(\boldsymbol{z}) \coloneqq \boldsymbol{P}_j(\boldsymbol{X}_0^\top \boldsymbol{y}_0 + \boldsymbol{X}_j^\top \boldsymbol{z})$ for the fit on $D_0 \cup \{(\boldsymbol{X}_j, \boldsymbol{z})\}$. Then $b_j = \bar{\ell}(\check{\theta}_\emptyset) - \bar{\ell}(\check{\theta}(\hat{\boldsymbol{y}}_j)) - \tilde{v}(\{j\})$ involves no other seller's labels.

At truthful delivery the raw score has mean $\ell(\emptyset) - \ell(\{j\}) = v(\{j\}) = \tilde{v}(\{j\})$ by the $\rho = 1$ step of Section B.1, so $\mathbb{E}[b_j \mid \boldsymbol{X}, \hat{\boldsymbol{c}}] = 0$.

Label incentive. For any ϑ fixed given $D_0 \cup D_j$, averaging over the test noise gives $\mathbb{E}[\bar{\ell}(\vartheta) \mid \theta] = \sigma^2 + \|\theta - \vartheta\|_{\boldsymbol{A}}^2$. Under the uninformative prior, seller j's belief given $D_0 \cup D_j$ has mean $\check{\theta}(\boldsymbol{y}_j)$ and finite covariance Σ_j, so the cross term vanishes and $\mathbb{E}[\bar{\ell}(\vartheta) \mid D_0 \cup D_j] = \sigma^2 + \operatorname{tr}(\boldsymbol{A}\Sigma_j) + \|\check{\theta}(\boldsymbol{y}_j) - \vartheta\|_{\boldsymbol{A}}^2$. Applying this to $\check{\theta}_\emptyset$ and $\check{\theta}(\hat{\boldsymbol{y}}_j)$,

$$\mathbb{E}[b_j \mid D_0 \cup D_j] = \|\check{\theta}(\boldsymbol{y}_j) - \check{\theta}_\emptyset\|_{\boldsymbol{A}}^2 - \tilde{v}(\{j\}) - \left\|\boldsymbol{P}_j \boldsymbol{X}_j^\top (\hat{\boldsymbol{y}}_j - \boldsymbol{y}_j)\right\|_{\boldsymbol{A}}^2.$$

Only the last term depends on $\hat{\boldsymbol{y}}_j$. It is the $\boldsymbol{A}$-distance between the fits at the delivered and the true labels, and it does not depend on D_0. It is non-negative and vanishes at $\hat{\boldsymbol{y}}_j = \boldsymbol{y}_j$, so truthful delivery maximizes $\mathbb{E}[b_j \mid D_j]$ whatever the other winners deliver. Truthful delivery is therefore a dominant strategy, and in particular a Bayes–Nash equilibrium. Since $\boldsymbol{A} \succ 0$ and $\boldsymbol{P}_j$ is invertible, the penalty vanishes only if $\boldsymbol{X}_j^\top \hat{\boldsymbol{y}}_j = \boldsymbol{X}_j^\top \boldsymbol{y}_j$. The pooled fit $f_{\mathrm{LS}}(\hat{D}_S) = \boldsymbol{I}(S)^{-1}\big(\boldsymbol{X}_0^\top \boldsymbol{y}_0 + \sum_{i \in S} \boldsymbol{X}_i^\top \hat{\boldsymbol{y}}_i\big)$ depends on $\hat{\boldsymbol{y}}_j$ only through $\boldsymbol{X}_j^\top \hat{\boldsymbol{y}}_j$, so this forces $f_{\mathrm{LS}}(\hat{D}_S) = f_{\mathrm{LS}}(D_S)$. That is the equality clause of Theorem 3.3. □

C Details for Section 3

Proof of Theorem 3.2. Write S_t for the candidate set after round t of Step 2 of Section 2.2, with $S_0 = \emptyset$, and break ties in the $\arg\max$ lexicographically.

Cost incentive. Since $\tilde{v}$ does not depend on the bids, each per-round score $\phi_t \tilde{v}(i \mid S_{t-1}) - \hat{c}_i$ satisfies Theorem A.5, so by Theorem A.6 and Myerson's lemma the critical-bid payment τ_i is DSIC and IR. Adding b_i with $|\mathbb{E}[b_i \mid \boldsymbol{X}, \hat{\boldsymbol{c}}]| \le \epsilon$ shifts each report's expected utility by at most ϵ, which gives the 2ϵ-cost incentive, and the ϵ-cost incentive when $\mathbb{E}[b_i \mid \boldsymbol{X}, \hat{\boldsymbol{c}}]$ does not depend on $\hat{c}_i$.

Welfare. Under truthful bids, the greedy stage is the distorted greedy of Harshaw et al. (2019) run on $\tilde{v}$ with the true costs and distortion δ. By monotonicity, the inequality of Theorem 2.4 at ratio γ also holds at every $\delta \le \gamma$. Their proof invokes it only with $L = S_t$ and $S = T \setminus S_t$, where $|S| \le k$. So Harshaw et al. (2019, Thm. 3) gives, for every T with $|T| \le k$,

$$\tilde{v}(S_{\mathrm{greedy}}) - c(S_{\mathrm{greedy}}) \ge (1 - e^{-\delta})\tilde{v}(T) - c(T).$$

By Theorem 3.1, $\tilde{v}(S_{\mathrm{greedy}}) \le \rho v(S_{\mathrm{greedy}})$ and $\tilde{v}(T) \ge v(T)/\rho$ gives the bound of Theorem 3.2.

Efficiency. The greedy stage makes $O(Nk)$ oracle calls. By Theorem A.6, each of the at most k winners' critical bids requires one more run, with no binary search. This gives $O(Nk^2)$ calls in total. The quality score b_i is evaluated once per winner. □

The canonical case is the Shannon value of information Example C.1 with conditionally independent data.

Example C.1 (Shannon value of information)**.** *If the sellers' data are independent conditional on θ, let the action set $\mathcal{A}$ be the distributions over θ and ℓ be the log score. The value of information Eq.* (8)*, with the prior μ_0 induced by D_0, equals the mutual information $\mathrm{MI}(\theta; \boldsymbol{y}_S \mid D_0)$ between θ and the sellers' labels. The instance is informational substitutes because, under conditional independence, $S \mapsto \mathrm{MI}(\theta; \boldsymbol{y}_S \mid D_0)$ is monotone and submodular.*

Proposition C.2 (Admissible Bayesian surrogates)**.** *Suppose the principal knows the joint distribution of the outcome and the sellers' data, and the instance is informational substitutes. Then $\tilde{v} = v_\ell$ is admissible with $\rho = 1$ and $\gamma = 1$, and* GRADES *run at $\delta = 1$ satisfies the guarantees of Theorem 3.2 with the bi-criteria bound $(1 - 1/e, 1)$ on the true valuation.(Chen & Waggoner, 2017)*

Evaluating v_ℓ, however, often requires full knowledge of the common prior μ, and even when μ is known, v_ℓ rarely has a closed form. The objectives in Eqs. (2) and (6) avoid this by using surrogates Eqs. (3) and (7). Similar idea is also used in Xu et al. (2021); Lu et al. (2024). Alternatively, Hu et al. (2026) score sellers by a single parameter function, and Kang et al. (2023); Ananthakrishnan et al. (2026) ask sellers to provide subset of their dataset in allocation stage. To make data markets practical, we believe a better understanding is needed of valuation function v and which surrogates are admissible (Definition 3.1) the models studied here.

D Additional acquisition evidence

D.1 Reproduction parameters and allocation definitions

This appendix documents the forward comparison in Section 4. The experiment has no initially owned labeled records, uses a fixed regularizer $I_\eta(\varnothing) = \eta I_d$, and gives every comparator the monetary cap induced by native GRADES,

$$B_{G,\eta}(k) = c(S_{G,\eta}(k)).$$

The regularizer is not interpreted as initially owned data. The realized acquisitions, final fits, and expenditures are those of the forward study; reverse-budget experiments are not mixed into the results below.

Table 2: Forward-study parameters for the experimental comparison.

Parameter	Setting
Market and test tasks	$N = 300$ single-record sellers; three disjoint $m = 32$ test designs; 20 market resamples.
Initial labeled data	None; the zero-coefficient predictor is the no-acquisition reference.
Capacity	$k = 1, \ldots, 30$, with $k = 25$ for endpoint tables.
Primary regularization	$\eta = 1$; $I_\eta(S) = \eta I_d + \sum_{i \in S} X_i^\top X_i$.
Final learner	Acquired-only ridge: $\check{\theta}_{S,\eta} = I_\eta(S)^{-1} X_S^\top y_S$. The penalty includes the explicit intercept when present.
Real-data representation	Seller-fitted imputation/scaling, non-whitened PCA to at most 16 components, then one intercept coordinate.
Task weight	$A = (X^{\text{test}})^\top X^{\text{test}}/32$.
Design value	$\widetilde{v}_\eta(S) = 30(1 - f_\eta(S)/f_\eta(\varnothing))$, $f_\eta(S) = \operatorname{tr}(A I_\eta(S)^{-1})$.
Design welfare	$\widetilde{W}_\eta(S; c) = \widetilde{v}_\eta(S) - c(S)$.
Cost construction	Saved $r_i \sim \text{Unif}\{1, \ldots, 5\}$; $g(r) = \sqrt{r}$ or r^2, normalized separately to mean seller cost one.
Native GRADES	$\delta = 1$; positive distorted marginal-gain-minus-cost rule; capacity k and no external spending cap.
Comparator resource convention	Monetary cap $B_G = c(S_G)$; greedy controls also obey $\lvert S\rvert \le k$; ranked methods have no additional count cap.
Multi-step DAVED	Cost-constrained fractional design, at most 2,000 Frank–Wolfe updates, followed by ranking and discrete purchasing.
Numerical purchasing	Inclusive cost-feasible prefix; stop at the first unaffordable record.
Statistical unit	Market after averaging its three test tasks; pointwise paired t intervals use 19 degrees of freedom.

All methods use the same seller/test splits. Real-data preprocessing is fitted on seller covariates only.Test features are available to the task-specific design criteria, whereas test labels are reserved for evaluation. The feature-based selectors—GRADES, the greedy controls, and the adapted DAVED selectors—do not access seller labels before purchase. Data-OOB is different: its ranking uses seller labels through out-of-bag prediction errors, but it never accesses the test labels. Thus the comparison includes both stage-1-feasible feature-only rules and a label-access benchmark.

Let $P_S = I_\eta(S)^{-1}$. The Woodbury identity gives

$$\widetilde{v}_\eta(i \mid S) = \frac{30}{f_\eta(\varnothing)} \frac{x_i^\top P_S A P_S x_i}{1 + x_i^\top P_S x_i}.$$

GRADES recomputes this marginal gain from the current accepted set and uses the round-t score

$$\phi_t \widetilde{v}_\eta(i \mid S) - c_i, \qquad \phi_t = \left(1 - \frac{\delta}{k}\right)^{k-t}.$$

A round is accepted only when the largest score is positive. Because $\phi_t \leq 1$, every accepted seller also has $\widetilde{v}_\eta(i \mid S) - c_i > 0$. This is a statement about the declared design objective; it is not a statement about realized test-MSE improvement.

Adapted multi-step DAVED minimizes

$$F_\eta(z) = \operatorname{tr}\left[A\left(\eta I_d + \sum_i z_i x_i x_i^\top\right)^{-1}\right], \qquad z \geq 0, \quad c^\top z \leq B_G.$$

The final fractional weights are ranked and converted to an unweighted, cost-feasible acquisition. Adapted single-step DAVED ranks the zero-acquisition gradient score. Greedy-B sets $\delta = 0$, Design-greedy-B removes the explicit cost penalty from the greedy score, and Best-δ greedy-B reports the best final design welfare over the declared distortion grid. Random-B uses a fixed pseudorandom ranking.

The experiment deliberately retains the V-optimal matrix structure, but its fixed ηI_d baseline is a regularizer rather than a full-rank labeled D_0, and its final learner is acquired-only ridge. Consequently, the theoretical identity $v(S) = \widetilde{v}(S)$ for expected OLS risk reduction is not asserted for this experimental control. We therefore report design welfare, realized prediction, and empirical prediction value separately.

For an outcome Y, let

$$\bar{Y}_{h,u} = \frac{1}{3}\sum_{a=1}^{3} Y_{h,a,u}$$

denote the average over the three test tasks for method u in market resample h. Displayed means and standard errors use the resulting 20 market-level observations. For paired comparisons between GRADES and adapted multi-step DAVED, let $d_h = \bar{Y}_{h,G} - \bar{Y}_{h,D}$ and report $\bar{d} \pm t_{0.975,19} s_d/\sqrt{20}$. The intervals are pointwise and are not equivalence tests.

For an acquired set S,

$$\bar{\ell}_\eta(S) = \frac{1}{m}\left\|X^{\text{test}}\check{\theta}_{S,\eta} - y^{\text{test}}\right\|_2^2.$$

Table 3 is the prediction table moved from the main text, augmented with one standard error and the no-acquisition reference.

MSE itself is a loss, so lower is better. To express realized prediction on the same "higher is better" direction as a value, we use one fixed reporting reference for each dataset,

$$L_{\mathcal{D}}^{\text{ref}} = \frac{1}{60}\sum_{h=1}^{20}\sum_{a=1}^{3} \bar{\ell}_{1;h,a}(\varnothing),$$

and define

$$V_{h,a}(S) = \frac{30}{L_{\mathcal{D}}^{\text{ref}}}\left[\bar{\ell}_{1;h,a}(\varnothing) - \bar{\ell}_{1;h,a}(S)\right], \qquad W_{h,a}^{\text{pred}}(S;c) = V_{h,a}(S) - c(S). \tag{14}$$

The normalization is used only for evaluation; it does not enter allocation. It fixes a common multiplier within each dataset so that method ordering by mean MSE is not changed by method-specific normalization. The scale 30 is a declared reporting convention matching the design value scale; it is not an estimate of the theoretical noise variance.

Table 4 separates the stage-1 design value from the realized prediction value. Each slash-separated pair is GRADES / adapted multi-step DAVED.

he MIMIC result makes the distinction especially transparent. Table 5 shows that Data-OOB can achieve higher empirical $V - C$ despite much lower design welfare. These metrics capture different aspects of the mechanism: allocation must score sellers before their purchased labels are revealed, so

Table 3: Test MSE at $k = 25$ and $\eta = 1$: mean $\pm$ one standard error over 20 markets after within-market test-task averaging. Lower is better. Bold includes the no-acquisition reference.

Method	Gaussian	MIMIC	RSNA	Fitzpatrick†	DrugLib
Square-root costs: $g(r) = \sqrt{r}$					
No acquisition	2.4547 ± 0.4503	219.78 ± 11.90	$\mathbf{17644.84} \pm 237.01$	$\mathbf{13.4113} \pm 0.1844$	$\mathbf{57.1667} \pm 0.8865$
Random-B	1.7768 ± 0.3357	184.49 ± 17.22	32366.20 ± 1544.98	23.4749 ± 1.0886	112.5088 ± 8.7646
Data-OOB	2.1035 ± 0.4089	$\mathbf{170.81} \pm 11.11$	31965.07 ± 1346.56	21.3262 ± 0.7231	105.7125 ± 5.2272
Design-greedy-B	1.3609 ± 0.2561	218.76 ± 31.22	25158.18 ± 1021.19	19.6502 ± 0.5288	83.7899 ± 4.2881
DAVED single (adapted)	1.4897 ± 0.2559	265.70 ± 15.21	20286.87 ± 403.25	16.5210 ± 0.4579	68.3846 ± 2.0075
DAVED multi (adapted)	0.9815 ± 0.1890	203.72 ± 18.03	27129.21 ± 1041.01	22.7918 ± 1.0231	106.3381 ± 9.1366
Greedy-B ($\delta = 0$)	$\mathbf{0.9238} \pm 0.1592$	184.96 ± 15.75	27659.87 ± 1174.07	22.0434 ± 0.8781	107.5224 ± 5.9639
GRADES ($\delta = 1$)	0.9241 ± 0.1531	172.94 ± 14.73	27259.76 ± 1244.78	22.3763 ± 0.7008	114.2275 ± 6.8253
Square costs: $g(r) = r^2$					
No acquisition	2.4547 ± 0.4503	219.78 ± 11.90	17644.84 ± 237.01	13.4113 ± 0.1844	57.1667 ± 0.8865
Random-B	2.2720 ± 0.4191	205.32 ± 20.68	21242.54 ± 827.05	14.6547 ± 0.2556	68.0747 ± 1.9165
Data-OOB	2.3384 ± 0.4313	186.71 ± 12.63	21374.13 ± 916.93	15.3193 ± 0.2987	73.1730 ± 4.6677
Design-greedy-B	1.7034 ± 0.3060	205.05 ± 22.81	24048.72 ± 826.97	17.1515 ± 0.6420	83.4452 ± 3.0499
DAVED single (adapted)	2.0955 ± 0.3997	252.43 ± 17.00	18118.06 ± 297.12	13.9345 ± 0.2821	59.1314 ± 1.2065
DAVED multi (adapted)	0.4632 ± 0.0754	$\mathbf{173.33} \pm 17.78$	7353.30 ± 1068.07	7.0350 ± 1.3165	43.3723 ± 3.5062
Greedy-B ($\delta = 0$)	0.4892 ± 0.0774	179.75 ± 17.59	20138.51 ± 1946.05	18.1441 ± 0.8570	89.7283 ± 5.8087
GRADES ($\delta = 1$)	$\mathbf{0.4547} \pm 0.0735$	177.13 ± 18.10	$\mathbf{4467.57} \pm 409.05$	$\mathbf{4.3715} \pm 0.3965$	$\mathbf{40.8242} \pm 3.2883$

Table 4: Value–cost accounting at $k = 25$, $\eta = 1$. Each pair is GRADES / adapted multi-step DAVED. $\widetilde{v}_1$ and $\widetilde{W}_1$ are the feature-based design value and design welfare used to evaluate allocation. V and $V - C$ use the realized ridge test loss through Eq. equation 14. Values are rounded means; the empirical columns are evaluation metrics and do not enter selection.

Dataset	Costs	$\widetilde{v}_1$ G/D	C G/D	$\widetilde{W}_1$ G/D	V G/D	$V - C$ G/D
Gaussian	$\sqrt{r}$	12.506 / 12.241	7.790 / 7.734	4.716 / 4.507	18.706 / 18.005	10.916 / 10.271
Gaussian	r^2	17.643 / 17.623	2.313 / 2.313	15.331 / 15.311	24.443 / 24.339	22.130 / 22.026
MIMIC	$\sqrt{r}$	22.355 / 21.454	6.702 / 6.536	15.653 / 14.918	6.394 / 2.192	-0.308 / -4.344
MIMIC	r^2	26.423 / 26.441	2.331 / 2.331	24.092 / 24.110	5.822 / 6.340	3.491 / 4.009
RSNA	$\sqrt{r}$	28.252 / 26.250	8.008 / 7.776	20.244 / 18.475	-16.347 / -16.125	-24.355 / -23.901
RSNA	r^2	29.671 / 29.540	1.707 / 1.706	27.964 / 27.834	22.404 / 17.498	20.697 / 15.792
Fitzpatrick†	$\sqrt{r}$	28.558 / 26.811	8.534 / 8.288	20.024 / 18.523	-20.054 / -20.983	-28.588 / -29.271
Fitzpatrick†	r^2	29.696 / 29.577	1.672 / 1.672	28.023 / 27.905	20.221 / 14.263	18.549 / 12.591
DrugLib	$\sqrt{r}$	28.470 / 26.300	8.154 / 7.877	20.315 / 18.423	-29.944 / -25.804	-38.098 / -33.681
DrugLib	r^2	29.734 / 29.583	1.703 / 1.703	28.031 / 27.879	8.576 / 7.239	6.873 / 5.536

GRADES uses the feature-based $\widetilde{v}_\eta$, whereas realized V requires the delivered labels and held-out test labels. This information constraint is central to the inspection problem and also determines how the positive-score rule should be interpreted. For every seller i accepted at round t, $\phi_t \widetilde{v}_\eta(i \mid S) - c_i > 0$ implies $\widetilde{v}_\eta(i \mid S) - c_i > 0$, since $\phi_t \leq 1$ and the marginal design gain is nonnegative. Thus each acceptance increases *design welfare*, but need not increase realized $V - C$. Although the V-optimal surrogate equals expected risk reduction under the full-rank initial-data OLS model of Section 2, the present no-initial-data ridge experiment does not satisfy those assumptions, so we do not invoke that identity to equate $\widetilde{v}_\eta$ with realized V. Instead, we ask whether optimizing a design criterion computable before label acquisition yields competitive prediction after delivery. The RSNA results under square-root costs further illustrate the distinction: GRADES achieves positive design welfare (20.244) but negative empirical $V - C$ (approximately -24.355), with mean MSE exceeding the no-acquisition reference. We therefore report both metrics rather than treating either as a substitute for the other.

Table 5: MIMIC, square-root costs, $k = 25$, $\eta = 1$. Data-OOB has slightly lower realized MSE and expenditure than GRADES, so it also has larger empirical $V - C$, while its feature-based design value and design welfare are much smaller. Data-OOB is a label-access benchmark: its seller ranking uses seller labels before purchase.

Method	MSE ↓	$\widetilde{v}_1$	C	$\widetilde{v}_1 - C$ ↑	V ↑	$V - C$ ↑
GRADES	172.943	22.355	6.702	15.653	6.394	−0.308
DAVED multi	203.722	21.454	6.536	14.918	2.192	−4.344
Data-OOB	**170.809**	9.094	**6.152**	2.942	**6.684**	**0.532**

Table 6 reports the paired MSE and design-welfare comparisons that underlie the main-text summary.

Table 6: Paired endpoint differences, GRADES minus adapted multi-step DAVED, at $k = 25$, $\eta = 1$. Negative MSE differences and positive design-welfare differences favor GRADES. Intervals are pointwise paired 95% t intervals over 20 markets. †Available-image subset.

Dataset	Costs	$\mathrm{MSE}_G - \mathrm{MSE}_D$ [95% CI]	$\widetilde{W}_1(S_G) - \widetilde{W}_1(S_D)$ [95% CI]
Gaussian	$\sqrt{r}$	−0.0574 [−0.1544, 0.0395]	+0.2085 [0.1732, 0.2438]
Gaussian	r^2	−0.0086 [−0.0322, 0.0151]	+0.0201 [−0.0223, 0.0625]
MIMIC	$\sqrt{r}$	−30.78 [−63.78, 2.22]	+0.7350 [0.5759, 0.8940]
MIMIC	r^2	+3.80 [−7.41, 15.02]	−0.0185 [−0.0488, 0.0118]
RSNA	$\sqrt{r}$	+130.55 [−2647.59, 2908.69]	+1.7693 [1.4772, 2.0614]
RSNA	r^2	−2885.72 [−5225.56, −545.89]	+0.1293 [0.0390, 0.2196]
Fitzpatrick†	$\sqrt{r}$	−0.4156 [−2.7553, 1.9241]	+1.5010 [1.2914, 1.7106]
Fitzpatrick†	r^2	−2.6635 [−5.4455, 0.1185]	+0.1186 [0.0467, 0.1904]
DrugLib	$\sqrt{r}$	+7.8894 [−14.3475, 30.1263]	+1.8920 [1.5794, 2.2045]
DrugLib	r^2	−2.5481 [−12.9313, 7.8351]	+0.1519 [0.0738, 0.2299]

GRADES has higher mean design welfare than adapted multi-step DAVED in nine of ten conditions, and eight paired design-welfare intervals favor GRADES. It has lower mean MSE in seven of ten conditions; only the RSNA square-cost MSE interval favors GRADES, while the other nine contain zero. Using Eq. equation 14, GRADES has higher mean $V - C$ in seven of ten conditions, but this realized-value comparison is less conclusive than the design-welfare comparison. In particular, the no-acquisition mean MSE is below every displayed square-root-cost acquisition mean on RSNA, Fitzpatrick, and DrugLib. The results therefore support favorable stage-1 design allocation without implying uniform realized prediction improvement.

D.2 Capacity sweeps on the remaining datasets

Figures 3–6 extend the Gaussian main-text figure under the same $\eta = 1$ forward protocol. The top row uses square-root costs, the bottom row square costs; columns show MSE and normalized design welfare. All 30 measured capacities are connected without smoothing. Bars show one market-level

standard error at $k = 1, 5, 10, 15, 20, 25, 30$. The two cost laws share a vertical range for each dataset and metric. Colors, markers, and line patterns are the same as in the Gaussian figure. The no-acquisition MSE is stated in each caption, and the captions' endpoint inferences use the paired intervals above, not overlap of the plotted marginal error bars.

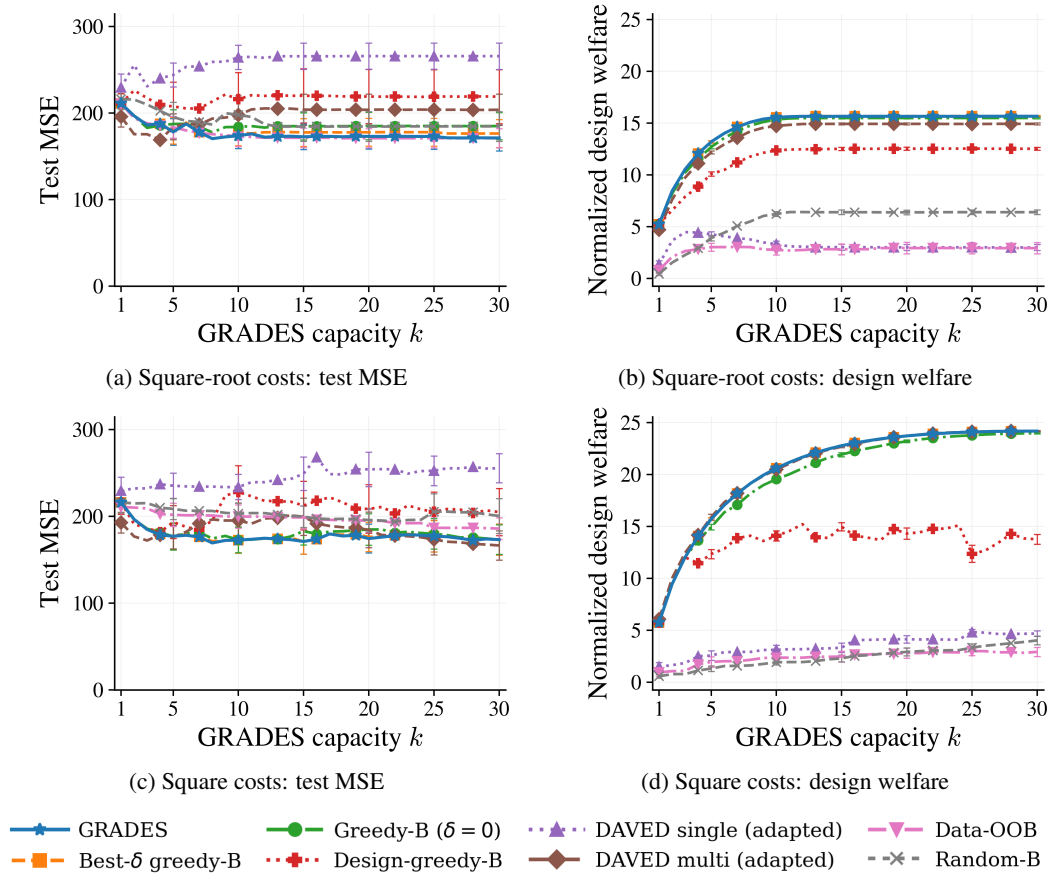

(a) Square-root costs: test MSE

(b) Square-root costs: design welfare

(c) Square costs: test MSE

(d) Square costs: design welfare

Figure 3: **MIMIC: GRADES-defined acquisition-cost caps, $\eta = 1$.** Lines show means and bars show one standard error, with the same aggregation and panel order as described above. The horizontal axis is GRADES's capacity k, not a common acquired count. At $k = 25$, the paired welfare interval against adapted multi-step DAVED favors GRADES under square-root costs, but contains zero under square costs. Both paired MSE intervals contain zero. The no-acquisition mean MSE is 219.78.

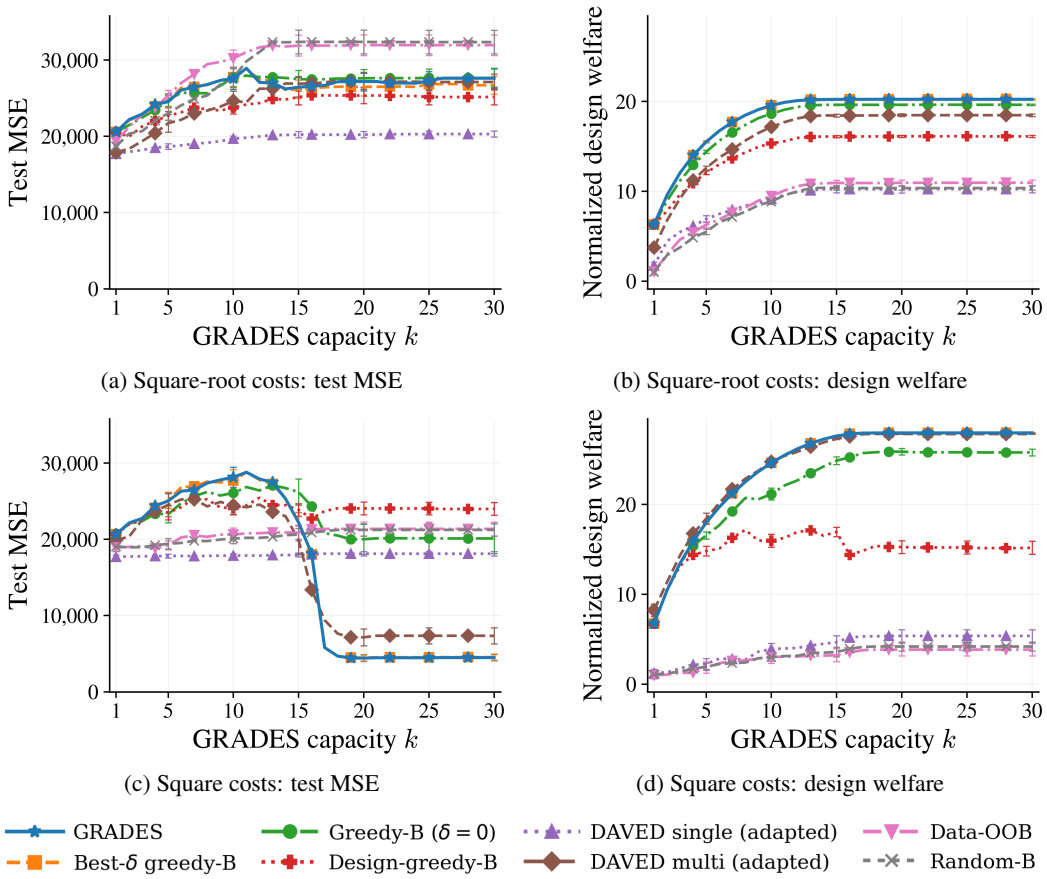

(a) Square-root costs: test MSE

(b) Square-root costs: design welfare

(c) Square costs: test MSE

(d) Square costs: design welfare

Figure 4: **RSNA Bone Age: GRADES-defined acquisition-cost caps, $\eta = 1$.** Lines show means and bars show one standard error, with the same aggregation and panel order as described above. The horizontal axis is GRADES's capacity k, not a common acquired count. At $k = 25$, paired welfare intervals against adapted multi-step DAVED favor GRADES under both cost laws; only the square-cost paired MSE interval favors GRADES. The no-acquisition mean MSE is 17644.84, below every displayed square-root-cost acquisition mean at this endpoint.

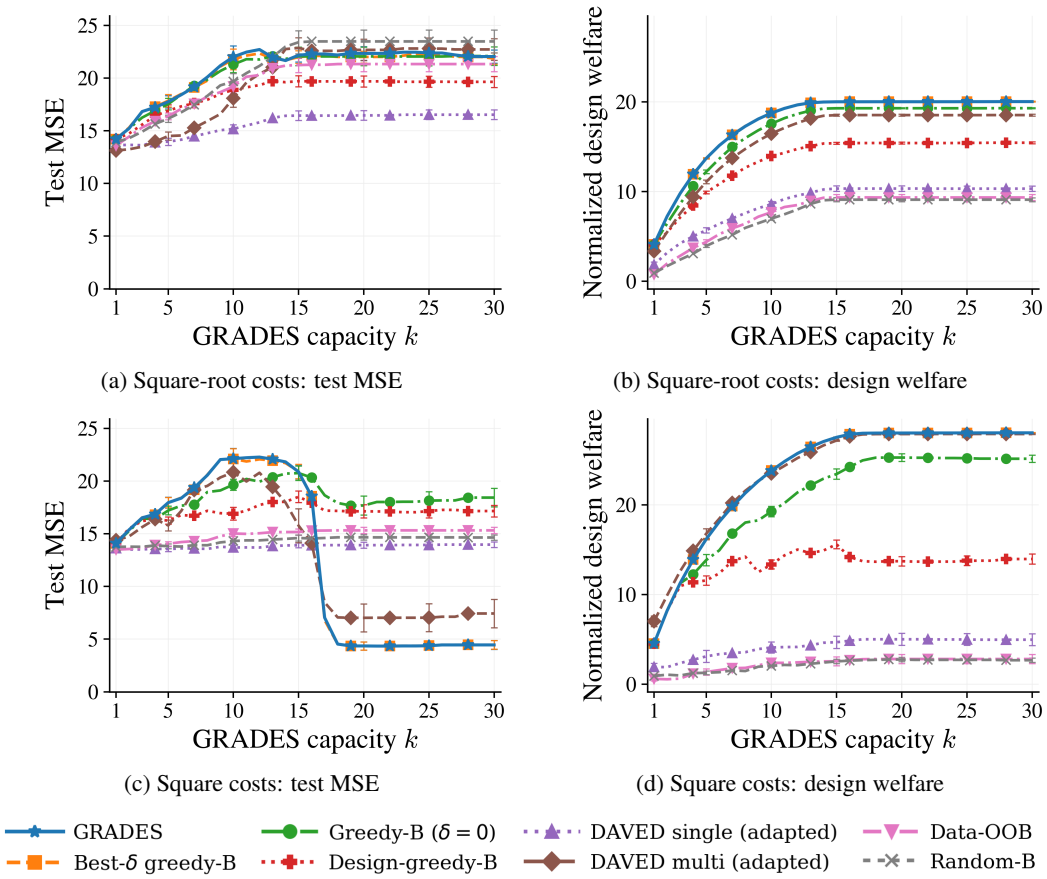

(a) Square-root costs: test MSE

(b) Square-root costs: design welfare

(c) Square costs: test MSE

(d) Square costs: design welfare

Figure 5: **Fitzpatrick17K, available-image subset: GRADES-defined acquisition-cost caps,** $\eta = 1$**.** Lines show means and bars show one standard error, with the same aggregation and panel order as described above. The horizontal axis is GRADES's capacity k, not a common acquired count. At $k = 25$, paired welfare intervals against adapted multi-step DAVED favor GRADES under both cost laws, while both paired MSE intervals contain zero. The no-acquisition mean MSE is 13.4113, below every displayed square-root-cost acquisition mean at this endpoint.

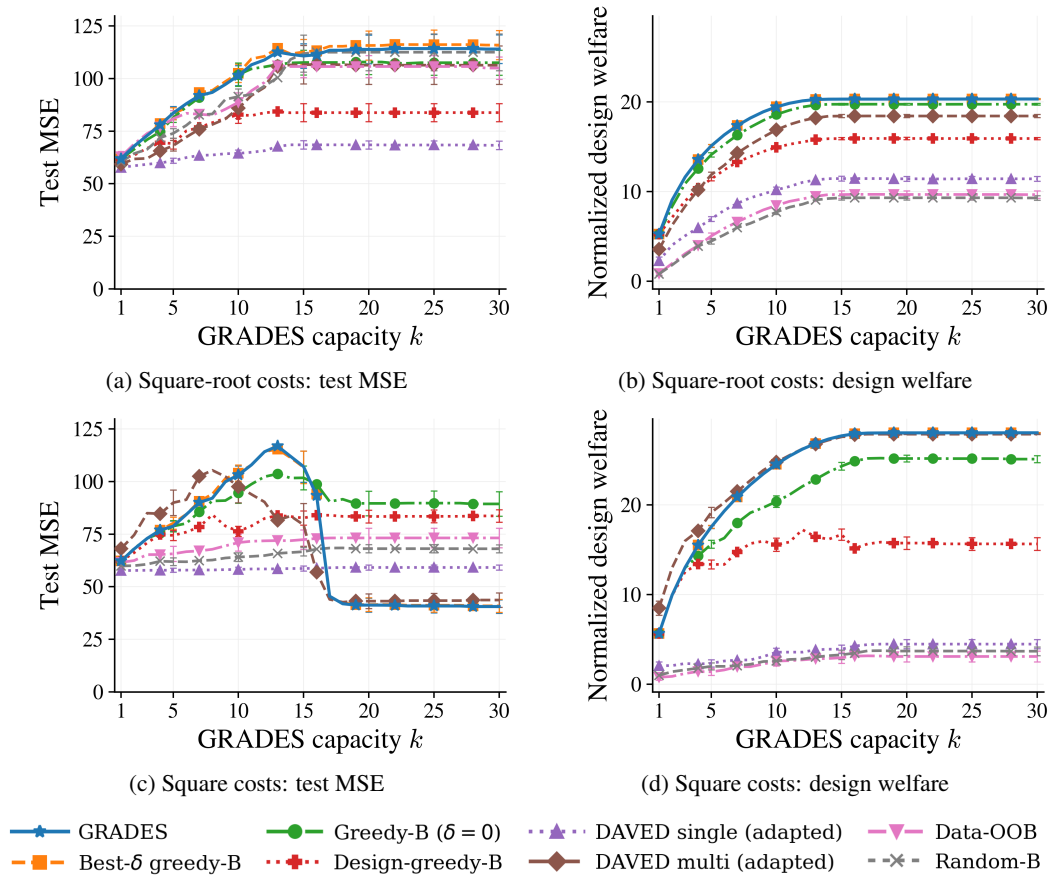

(a) Square-root costs: test MSE

(b) Square-root costs: design welfare

(c) Square costs: test MSE

(d) Square costs: design welfare

Figure 6: **DrugLib: GRADES-defined acquisition-cost caps, $\eta = 1$.** Lines show means and bars show one standard error, with the same aggregation and panel order as described above. The horizontal axis is GRADES's capacity k, not a common acquired count. At $k = 25$, paired welfare intervals against adapted multi-step DAVED favor GRADES under both cost laws, while both paired MSE intervals contain zero. The no-acquisition mean MSE is 57.1667, below every displayed square-root-cost acquisition mean at this endpoint.

The curves describe how the tested allocation rules behave as GRADES's capacity and the induced spending cap change. They do not imply monotonic MSE improvement or a growing prediction advantage over DAVED. In particular, sharp error changes and plateaus at intermediate capacities reflect the selected records and fitted models, not evidence that the surrogate and realized prediction value are the same quantity.

D.3 SELECTION RUNNING TIME

Fresh selection timing uses market seeds $0, \ldots, 4$, test-task index zero, and

$$k \in \{1, 5, 10, 15, 20, 25, 30\}.$$

Each task has one untimed warm-up followed by three recorded repetitions. The median of the three repetitions is taken within each task, and the five market-level medians are summarized by their mean and one standard error. Technical repetitions are not additional independent observations.

Timing includes selection initialization, method-specific scoring, all executed Frank–Wolfe updates, and discrete purchasing. For Best-δ, it includes all eleven distortion-grid trials; for Data-OOB, it includes its ensemble scoring. Timing excludes preprocessing, disk I/O, checkpointing, construction of the already-known comparator cap B_G, final prediction fitting, and mechanism payments. Thus these are selection-call times rather than end-to-end market runtime.

Table 7: Fresh selection time at $k = 25$, $\eta = 1$. Entries are mean $\pm$ one standard error over five market-level technical-repeat medians. The ratio is the DAVED-multi mean divided by the GRADES mean; it is not a mean of per-market ratios or an equal-accuracy speedup. †Available-image subset.

Dataset	Costs	GRADES (ms)	DAVED multi (ms)	Ratio
Gaussian	$\sqrt{r}$	1.758 ± 0.022	610.5 ± 19.4	347.2
Gaussian	r^2	2.186 ± 0.229	574.7 ± 21.9	262.9
MIMIC	$\sqrt{r}$	1.764 ± 0.022	683.6 ± 4.1	387.5
MIMIC	r^2	2.319 ± 0.015	696.7 ± 7.4	300.4
RSNA	$\sqrt{r}$	1.917 ± 0.034	803.2 ± 22.5	419.1
RSNA	r^2	2.076 ± 0.029	817.2 ± 10.3	393.7
Fitzpatrick†	$\sqrt{r}$	1.865 ± 0.053	676.9 ± 8.5	363.0
Fitzpatrick†	r^2	1.834 ± 0.146	703.4 ± 14.4	383.7
DrugLib	$\sqrt{r}$	2.005 ± 0.100	705.7 ± 14.3	352.0
DrugLib	r^2	2.103 ± 0.040	713.4 ± 14.3	339.3

At $k = 25$, GRADES's mean selection time is 1.76–2.32 ms, compared with 0.575–0.817 s for adapted multi-step DAVED. The corresponding ratios of means are approximately 263–419. All recorded endpoint DAVED calls execute the configured 2,000 Frank–Wolfe updates. Accordingly, this benchmark compares the reported finite-iteration implementations; it does not measure time to a common certified optimization error or a common test MSE.

The benchmark also does not time the full GRADES mechanism. The theoretical $O(Nk^2)$ oracle bound includes counterfactual reruns needed for critical payments, whereas the experiment times one allocation call. No empirical full-mechanism speedup follows from the selection-only numbers above.

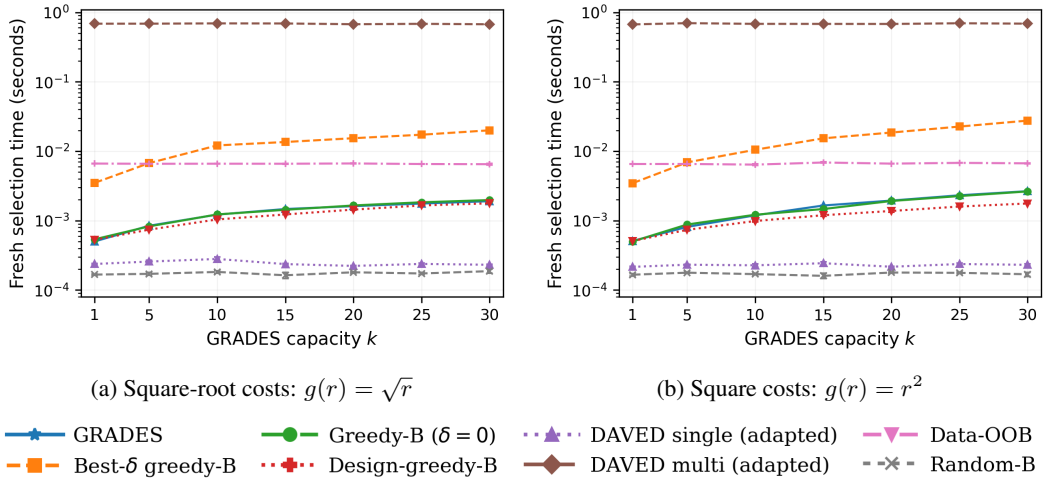

Figure 7: MIMIC fresh selection runtime at $\eta = 1$ under GRADES-defined acquisition-cost caps. The time axis is logarithmic; points and bars show means and one standard error over five market-level technical-repeat medians. The horizontal axis is GRADES's capacity, not a common acquired count. All eight primary methods are retained.

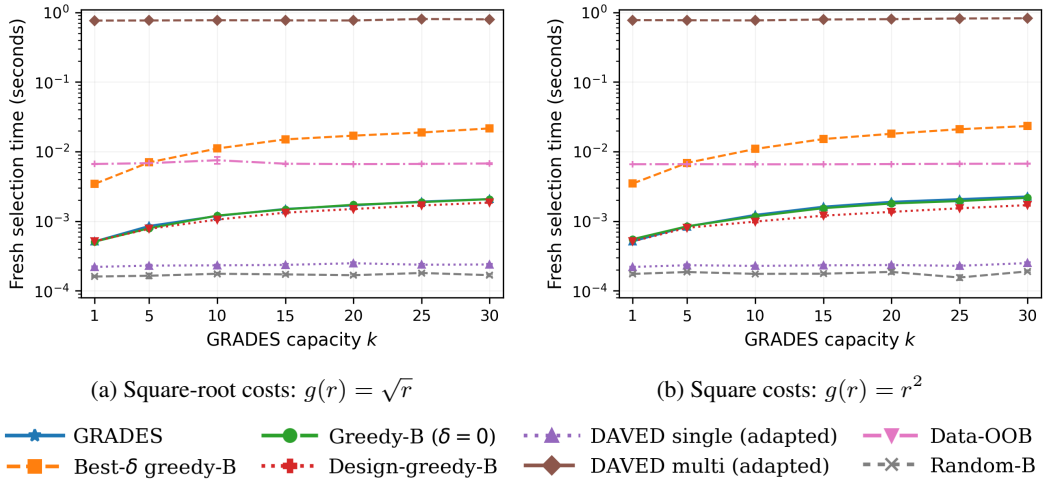

Figure 8: RSNA fresh selection runtime at $\eta = 1$ under GRADES-defined acquisition-cost caps. The time axis is logarithmic; points and bars show means and one standard error over five market-level technical-repeat medians. The horizontal axis is GRADES's capacity, not a common acquired count. All eight primary methods are retained.

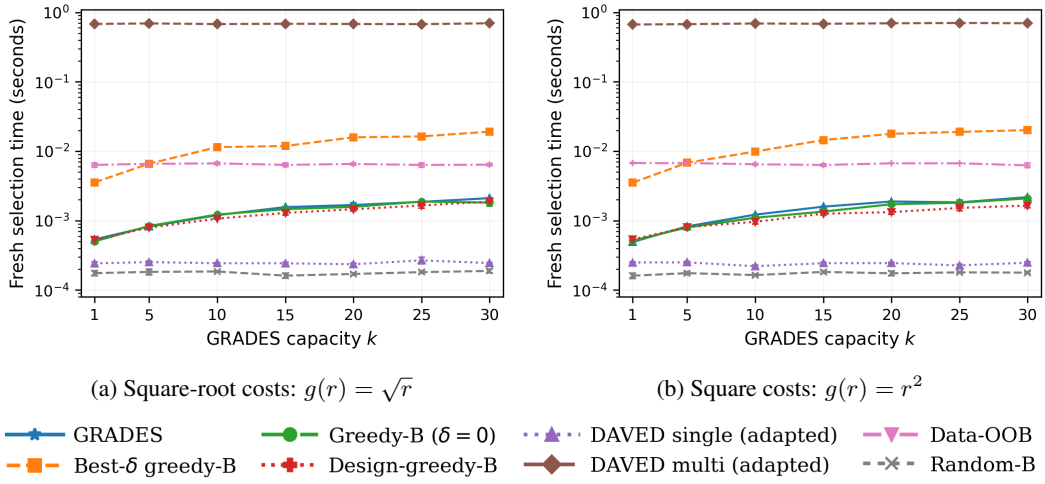

Figure 9: Fitzpatrick† fresh selection runtime at $\eta = 1$ under GRADES-defined acquisition-cost caps. The time axis is logarithmic; points and bars show means and one standard error over five market-level technical-repeat medians. The horizontal axis is GRADES's capacity, not a common acquired count. †Available-image subset.

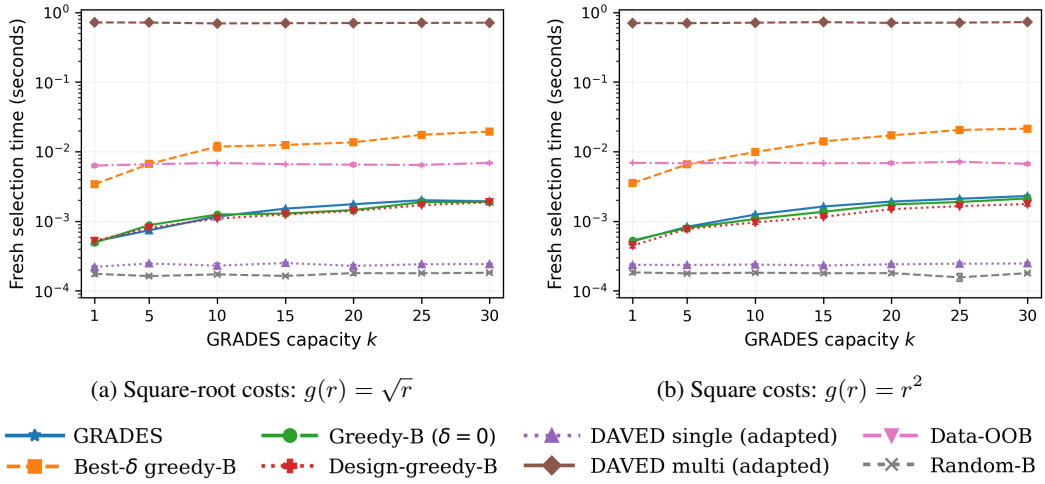

Figure 10: DrugLib fresh selection runtime at $\eta = 1$ under GRADES-defined acquisition-cost caps. The time axis is logarithmic; points and bars show means and one standard error over five market-level technical-repeat medians. The horizontal axis is GRADES's capacity, not a common acquired count. All eight primary methods are retained.