

# Online Linear Regression with Paid Stochastic Features

Nadav Merlis<sup>1</sup>, Kyoungseok Jang<sup>2</sup>, Nicolò Cesa-Bianchi<sup>3,4</sup>

<sup>1</sup> Technion – Israel Institute of Technology

<sup>2</sup> Chung-Ang University

<sup>3</sup> University of Milan

<sup>4</sup> Politecnico di Milano

nmerlis@technion.ac.il, ksjang@cau.ac.kr, nicolo.cesa-bianchi@unimi.it

## Abstract

We study an online linear regression setting in which the observed feature vectors are corrupted by noise and the learner can pay to reduce the noise level. In practice, this may happen for several reasons: for example, because features can be measured more accurately using more expensive equipment, or because data providers can be incentivized to release less private features. Assuming feature vectors are drawn i.i.d. from a fixed but unknown distribution, we measure the learner’s regret against the linear predictor minimizing a notion of loss that combines the prediction error and payment. We first study the case in which the mapping between payments and noise covariance is known and prove order-optimal regret bounds in the interaction length (up to log-factors). We then derive order-optimal bounds also when the noise covariance is unknown and prove that the regret rate is worse than the case of known covariances. Our analysis leverages matrix martingale concentration, showing that the empirical loss uniformly converges to the expected one for all payments and linear predictors.

**Extended version** — <https://arxiv.org/abs/2511.08073>

## 1 Introduction

In online linear regression, a learner sequentially observes i.i.d. feature vectors  $x_t \in \mathbb{R}^d$  and aims to correctly predict outputs  $y_t = x_t^\top \theta^*$ , knowing neither  $\theta^*$  nor the distribution from which the feature vectors are drawn. When measured by the quadratic prediction error, this is potentially the most canonical online regression setting, and several extensions to more complex losses and regression models have been studied in the past. Although in the vast majority of works the features  $x_t$  are assumed to be perfectly observable, in practice, we usually only have access to a noisy version of  $x_t$ . For example, features obtained from experiments often have measurement errors, and contextual information on users only relies on a random sample of their behavior. Sometimes the features are even intentionally corrupted (e.g., via locally differentially private mechanisms) to protect user privacy.

Motivated by scenarios where the learner can reduce the noise intensity by allocating more resources or more work, we study a variant of online linear regression where the features are stochastic and corrupted by additive noise, and the learner

may decide to pay to reduce the noise level so that a higher payment corresponds to a smaller noise covariance (w.r.t. the positive definite order). When conducting experiments, this payment can be related to renting better lab equipment or hiring more experienced staff, but also to extending the measurement duration and/or averaging more raw samples. With user data, this noise reduction can be achieved by collecting more statistics before making a decision. In privacy-protected settings, the learner may offer payments to compensate users for limited privacy loss (either directly or by offering free services)—see, e.g., (Wang, Ying, and Zhang 2018). While reducing the feature noise naturally facilitates prediction, it also causes some penalty to the learner (either time, money, or given services). Thus, the learner should optimally balance the prediction error and the cost due to payments.

We measure a predictor’s performance by a loss combining the prediction error with the payment. The goal of the learner is to minimize regret, defined as the difference between the learner’s expected cumulative loss and the cumulative loss of the best linear predictor, i.e., the one that optimally balances prediction errors and payments.

We first tackle the setting where the learner knows how the noise covariance changes as a function of the payment. This can be the case, for instance, when the noise is generated by a known differentially private mechanism. We show how to use this knowledge to effectively utilize samples collected at one payment level to estimate the optimal prediction at any other payment level. This information sharing leads to a major acceleration in the learning process and obviates the need for payment exploration. In particular, we devise an empirical loss estimator that uniformly converges to the expected one across all potential payment levels and predictions (including payments that were never offered by the algorithm). We utilize this result to prove that an algorithm that greedily minimizes this empirical loss to determine the payment level and linear predictor achieves a regret bound of  $\tilde{O}(\sqrt{T})$  after  $T$  prediction rounds. We complement this result with a lower bound, proving that the dependence on  $T$  in this bound can only be improved by logarithmic factors. Note that the optimal rate in this setting is exponentially worse than the  $\Theta(\ln T)$  bound for the noise-free case.

We then move on to the case where the noise covariances are not known in advance. This could happen, for example, when noise reduction is the result of an extended data collec-

tion period, especially if the noise is temporally correlated. In this scenario, we can no longer use our uniform estimation scheme and must incorporate a payment exploration mechanism into our algorithm. Inspired by UCB-like algorithms (Auer 2002), we limit the potential payments to a discretized grid and introduce an optimistic loss estimator, which encourages the algorithm to choose less-played payment levels. We show that by carefully tuning the grid size, this approach yields a regret bound of  $\tilde{O}(T^{2/3})$ . We prove that, similarly to the case of known covariances, the dependence on  $T$  in this bound can only be improved by logarithmic factors. Our lower bound construction carefully designs a set of noise covariance profiles to induce losses that imitate hard instances in Lipschitz bandits (Kleinberg 2004).

**Outline.** The rest of the paper is organized as follows. We first position our work in light of the existing literature in Section 2. We then formalize our setting and describe our assumptions in Section 3. In Section 4, we present our loss estimator when the noise covariance profile is known in advance; we show how to use this estimator to derive an order-optimal algorithm. Then, Section 5 extends these results to situations when the noise covariance profile is not given to the learner. We conclude the paper with a summary and a discussion on the many potential future directions in Section 6.

## 2 Related Work

**Errors-in-variables models.** Linear regression with noisy covariates (also known as errors-in-variables, or EIV) is a classical topic in statistics—see, e.g., (Fuller 2009) and (Agarwal and Singh 2021) for recent applications to differential privacy. Most results in EIV concern the problem of estimating  $\theta^*$  under different norms, typically in a high-dimensional setting (large  $d$ ) under assumptions of sparsity ( $\theta^*$  has few non-zero components) or low-rank (the covariates admit a low-dimensional representation), see (Rosenbaum and Tsybakov 2010; Loh and Wainwright 2011; Chen and Caramanis 2013; Rosenbaum and Tsybakov 2013; Kaul and Koul 2015; Belloni, Rosenbaum, and Tsybakov 2017; Datta and Zou 2017) and references therein. A related line of work studies EIV under Bayesian assumptions (Reilly and Patino-Leal 1981; Ungarala and Bakshi 2000; Figueroa-Zúñiga et al. 2022). Agarwal et al. (2023) consider EIV in an estimation error task where covariates are chosen adaptively (as in active learning or in multi-armed bandits). Closer to our work is the paper by Agarwal et al. (2019), where they study a linear regression model with fixed design and missing or corrupted variables. Unlike us, they bound the prediction error in a transductive setting, where the covariates on which predictions are evaluated are available at training time.

**Local DP.** A prominent situation where features are noisy is when they are protected by a local differentially private mechanism (Kasiviswanathan et al. 2011). Two well-known approaches to create this protection are to add the raw features either Laplace or Gaussian noise (Yang et al. 2024). Many works study how to learn when features are private, including in the context of linear regression (Duchi, Jordan, and Wainwright 2013; Smith, Thakurta, and Upadhyay 2017;

Wang and Xu 2019; Fukuchi, Tran, and Sakuma 2017). However, their goal is mostly to estimate the true parameters of the model given the noisy features; in contrast, we want to learn how to use online noisy data to output good predictions. As we will show, the best predictors are not the true parameters, but rather depend on the noise level.

**Online learning with noisy features.** Closely related to ours is the work by Cesa-Bianchi, Shalev-Shwartz, and Shamir (2011), who studied online linear regression with square loss with noisy features and known/unknown covariances. However, there are several crucial differences: they do not have a notion of cost, and their notion of regret compares to the best linear predictor on the noise-free features. Moreover, their focus is on a setting in which the learner may obtain more than one noisy realization of the same feature vector. A second closely related work is the one by Van Der Hoeven et al. (2023). They prove a  $d^2(\ln T)\sqrt{T}$  regret bound in a binary classification setting in which: the learner can pay to reduce the noise level, the covariates (called experts in their setting) are binary, and noise and payments apply independently on each coordinate. However, because of their focus on classification with independent feature noise, their techniques are not directly applicable to our regression setting with correlated noise. Other works assume noisy features, but in the partial feedback setting, include (Kim et al. 2023) and (Zheng et al. 2020).

**Privacy and incentives.** Another aspect of our work is the learner’s capability to pay for noise reduction. In the privacy literature, this is often addressed as a problem of mechanism design: an analyst designs a procedure that incentivizes individuals to sell their private data, aiming to obtain the most accurate estimation at a fixed payment budget (or paying a minimal cost for a target accuracy level) (Ghosh and Roth 2011; Nissim, Orlandi, and Smorodinsky 2012; Nissim, Vadhan, and Xiao 2014; Wang, Ying, and Zhang 2018; Fallah et al. 2024). Hsu et al. (2014) also studied the problem in which an individual is offered a payment to participate in a study and is willing to join it only if the payment covers its privacy loss. They show how an analyst running the experiment should tune the payment and privacy levels according to various considerations, including accuracy. All these works focus on payments that minimize the estimation error of the true parameters. Instead, our goal is to choose payments that balance the prediction error and the feature costs for every new incoming user. Also, our payment model is different: in our setting, the noise level is fixed (and potentially unknown) given the payment.

## 3 Setting

We study the problem of online linear regressions when the learner observes noisy features and can pay to reduce the noise level. At each round  $t \in [T]$ , the environment privately generates the true features  $x_t \in \mathbb{R}^d$  independently from a fixed unknown distribution  $D_x$  and the output  $y_t = x_t^\top \theta^*$  for some unknown  $\theta^* \in \mathbb{R}^d$ .<sup>1</sup> The learner decides to pay a cost

<sup>1</sup>All results can be easily extended to noisy linear outputs, as we discuss in Remark 3 and Appendix B.2.

(‘payment’)  $c_t \in [0, 1]$  and in return observes noisy features  $\hat{x}_t(c_t) = x_t + n_t(c_t)$ , where  $n_t(c) \sim D_n(c)$  is a zero-mean random variable independent of  $x_t$ . For brevity, we often omit the dependence on the cost  $c_t$  and denote the noisy features by  $\hat{x}_t$ . Given this observation, the learner tries to predict  $\hat{y}_t(\hat{x}_t)$  so that it is as close as possible (in the squared distance) to  $y_t$ , based on the noisy features  $\hat{x}_t$  and the data observed in the previous rounds. After the prediction, the learner observes the real value  $y_t$ . We focus on the class of linear predictors  $\hat{y}_t(\hat{x}_t) = \hat{x}_t^\top \nu$ . We denote  $\mathbb{E}[x_t] = \bar{x}$ , and also denote the covariance matrices of samples from  $D_x$  and  $D_n(c)$  by  $\Sigma_x = \mathbb{E}[x_t x_t^\top]$  and  $\Sigma_n(c) = \mathbb{E}[n_t(c) n_t(c)^\top]$ , respectively. Lastly, we denote the covariance of  $\hat{x}_t$  given a cost  $c$  by  $\Sigma_{\hat{x}}(c) = \mathbb{E}[\hat{x}_t(c) \hat{x}_t(c)^\top] = \Sigma_x + \Sigma_n(c)$ .<sup>2</sup>

**Additional Notations.** We define the quadratic form of a vector  $x \in \mathbb{R}^d$  w.r.t. a positive (semi-)definite matrix  $A \in \mathbb{R}^{d \times d}$  as  $\|x\|_A^2 = x^\top A x$ . We also denote by  $\|x\|$ , the  $\ell_2$  norm of a vector  $x \in \mathbb{R}^d$ , and by  $\|A\|_{op} = \sup_{x \neq 0} \frac{\|Ax\|}{\|x\|}$ , the operator norm of a matrix  $A \in \mathbb{R}^{d \times d}$ . We use  $A \succ 0$  to denote positive definite matrices, and similarly write  $A \succeq 0$  if  $A$  is positive semi-definite and  $B \succeq A$  when  $B - A \succeq 0$ . We denote the observed history of the decision process up to the beginning of round  $t + 1$  (including the choice of payment but not the features) by  $I_t^o = \{c_1, \hat{x}_1, \hat{y}_1, y_1, \dots, c_t, \hat{x}_t, \hat{y}_t, y_t, c_{t+1}\}$  and the entire past data by  $I_t = I_t^o \cup \{x_s, n_s\}_{s \in [t]}$ .

**Regularity assumptions.** We assume that increasing the payment decreases the noise level, i.e., if  $c_1 \leq c_2$ , it holds that  $\Sigma_n(c_2) \preceq \Sigma_n(c_1)$ . Intuitively, it is equivalent to requiring that the variance of the noise measured in any arbitrary direction cannot rise when we increase the payment. We also assume that  $\Sigma_{\hat{x}}(1) \succ 0$ ; by the covariance monotonicity assumption, this also implies that  $\Sigma_{\hat{x}}(c) \succ 0$  for all  $c \in [0, 1]$ . We further assume that both the features  $x_t$  and the noise  $n_t(c)$  are conditionally  $R^2$ -subgaussian random vectors, that is, for any  $u \in \mathbb{R}^d$ , we have  $\mathbb{E}\left[e^{u^\top (x_t - \bar{x})} | I_{t-1}\right] \leq e^{\frac{\|u\|^2 R^2}{2}}$  and  $\mathbb{E}\left[e^{u^\top n_t} | I_{t-1}\right] \leq e^{\frac{\|u\|^2 R^2}{2}}$ . Finally, we assume that both  $\|\theta^*\|$  and  $\|\bar{x}\|$  are upper bounded by a known  $S > 0$ .

**Optimality Criterion.** We compare ourselves to a linear regressor which pays a cost that optimally balances the prediction error and the payments for the features. In particular, for a given  $\lambda > 0$ , we define the loss for a fixed cost level  $c$  and linear predictor  $\nu$  as

$$\ell(c, \nu) = \mathbb{E}_{x \sim D_x, n \sim D_n(c)} \left[ \left( (x + n)^\top \nu - x^\top \theta^* \right)^2 \right] + \lambda c.$$

With a slight abuse of notation, for any given  $S > 0$ , we define the optimal linear predictor and loss given cost  $c$  as

$$\nu^*(c) = \arg \min_{\nu: \|\nu\| \leq S} \{\ell(c, \nu)\}, \quad \ell^*(c) = \min_{\nu: \|\nu\| \leq S} \{\ell(c, \nu)\},$$

and denote the optimal cost and loss by

$$c^* = \arg \min_{c \in [0, 1]} \{\ell^*(c)\}, \quad \ell^* = \min_{\substack{c \in [0, 1], \\ \nu: \|\nu\| \leq S}} \{\ell(c, \nu)\} = \min_{c \in [0, 1]} \{\ell^*(c)\}.$$

<sup>2</sup>Since  $x_t, \hat{x}_t$  are not zero mean,  $\Sigma_x, \Sigma_{\hat{x}}(c)$  are sometimes called their autocorrelation.

We remark that  $\nu^*(c), \ell^*(c)$  and  $\ell^*$  all depend on  $S$ ; we omit this dependence to simplify notations.

**Remark 1.** *Our choice to limit ourselves to linear predictions is due to both computational and statistical constraints. Specifically, a nonlinear optimal predictor requires computing  $\mathbb{E}[x_t | \hat{x}_t]^\top \theta^*$ . This estimator heavily depends on the entire feature and noise distributions, rendering the estimation problem extremely complicated and usually requires intractable computations. We adopt the solution of Van Der Hoeven et al. (2023) to this issue, and instead compare ourselves to the natural tractable benchmark. We also note that for some distributions, linear predictors are the best possible benchmark, as we later prove through our lower bounds.*

The loss and the optimal linear predictor can also be specified using the covariance matrices as follows:

**Claim 1.** *It holds that*

$$\ell(c, \nu) = \nu^\top \Sigma_{\hat{x}}(c) \nu - 2\nu^\top \Sigma_x \theta^* + \theta^{*\top} \Sigma_x \theta^* + \lambda c.$$

*Moreover, let  $\bar{\nu}(c) = (\Sigma_{\hat{x}}(c))^{-1} \Sigma_x \theta^*$ . If  $\|\bar{\nu}(c)\| \leq S$  then  $\nu^*(c) = \bar{\nu}(c)$  and*

$$\ell^*(c) = \theta^{*\top} \Sigma_x \theta^* - \theta^{*T} \Sigma_x \Sigma_{\hat{x}}(c)^{-1} \Sigma_x \theta^* + \lambda c.$$

Since we do not assume that the noise covariance is Lipschitz in the payment, the losses might not be continuous. Nonetheless, the monotonicity assumption induces a near-Lipschitz behavior:

**Claim 2.** *The loss  $\ell^*(c)$  is  $\lambda$ -one-sided Lipschitz: for any  $0 \leq c_1 \leq c_2 \leq 1$ , it holds that  $\ell^*(c_2) \leq \ell^*(c_1) + \lambda(c_2 - c_1)$ .*

In words, the loss cannot significantly increase from a slight rise in the payment. Intuitively, this holds since increasing the cost always reduces the prediction error; therefore, the loss increment is always upper bounded by the increase in the payment term  $\lambda c$ . This only linearly affects the loss, thus implying one-sided Lipschitzness. As a consequence, our algorithms could limit themselves to a finite grid of costs, knowing that even if the optimal cost is not on the grid, there is a slightly higher payment of similar loss (Tullii et al. 2024). The proofs for both claims can be found at Appendix A.

**Remark 2.** *Monotonicity of the noise covariances and Lipschitzness of the noise covariances are two incomparable assumptions – neither implies the other. Nonetheless, we rely on the monotonicity only to prove Claim 2, as well as arguing that  $\Sigma_{\hat{x}}(c) \succ 0$  for all  $c \in [0, 1]$ . Since Claim 2 also immediately holds if  $\Sigma_n(c)$  is Lipschitz in  $c$ , under the assumption that  $\Sigma_{\hat{x}}(c) \succ 0$ , all our positive results would also hold for Lipschitz covariances.*

### 3.1 Regret Definition

As the objective, algorithms are similarly measured by a loss that combines their prediction error and feature payments. We focus on minimizing the expected cumulative loss, also equivalent to minimizing the regret:

$$\text{Reg}(T) = \mathbb{E} \left[ \sum_{t=1}^T \left( (\hat{y}_t(\hat{x}_t) - y_t)^2 + \lambda c_t \right) \right] - T \ell^*.$$

The regret can be decomposed as follows (see Appendix A for the proof).

**Lemma 1.** For any arbitrary history-dependent prediction rule  $\hat{y}_t : \mathbb{R}^d \mapsto \mathbb{R}$  and any arbitrary sequence  $\nu_1, \dots, \nu_T \in \mathbb{R}^d$ , both potentially depend on  $I_{t-1}^o$ , it holds that

$$\begin{aligned} \text{Reg}(T) &= \mathbb{E} \left[ \sum_{t=1}^T \left( (\hat{y}_t(\hat{x}_t) - y_t)^2 - \left( (\hat{x}_t^\top \nu_t - y_t)^2 \right) \right) \right] \\ &\quad + \mathbb{E} \left[ \sum_{t=1}^T (\ell(c_t, \nu_t) - \ell^*) \right]. \end{aligned}$$

In particular, if the algorithm only outputs linear predictions  $\hat{y}_t(\hat{x}_t) = \hat{x}_t^\top \nu_t$ , then

$$\text{Reg}(T) = \mathbb{E} \left[ \sum_{t=1}^T (\ell(c_t, \nu_t) - \ell^*) \right].$$

The first term encapsulates the regret of the prediction rule compared to some dynamic linear predictor, while the second term measures this linear predictor compared to the best linear predictor. In the rest of the paper, we restrict ourselves to linear predictions (as discussed in Remark 1), and so we focus on bounding the loss difference.

## 4 Online Linear Regression with Known Noise Covariances

In this section, we start from the favorable setting in which the learner has perfect knowledge of the noise covariances  $\Sigma_n(c)$  for all  $c \in [0, 1]$ . In particular, to output a good prediction at the right cost, the learner only misses the feature covariance  $\Sigma_x$  and the linear map  $\theta^*$ . Since neither of these quantities depends on the payment  $c_t$ , our aim is to *share information* between different cost observations, thus improving the regret. To this end, we take a closer look at the expected loss  $\ell(c, \nu)$  in Claim 1, noticing that only a single term depends on the noise covariance: a term of the form  $\nu^\top \Sigma_{\hat{x}}(c) \nu$ . Given noisy independent feature samples at cost  $c$ , this term can be estimated empirically via  $\hat{\Sigma}_{\hat{x}}(c) = \sum_{s=1}^t \hat{x}_s(c) \hat{x}_s(c)^\top$ , whose expectation is  $t \Sigma_{\hat{x}}(c)$ . However, we only observe the noisy features for payments we played  $\{c_s\}_{s \leq t}$ , so we cannot directly use this estimator for other unplayed payments. To circumvent this, we utilize our knowledge of the noise covariance and include a correction term that *shifts* the noise covariance from  $c_s$  to  $c$ . In particular, we suggest estimating the covariance with

$$\hat{\Sigma}_{\hat{x}}(c) = \sum_{s=1}^t (\hat{x}_s(c_s) \hat{x}_s(c_s)^\top + \Sigma_n(c) - \Sigma_n(c_s)).$$

This estimator uses *all past samples* while enjoying the desired expectation

$$\begin{aligned} \mathbb{E} \left[ \hat{\Sigma}_{\hat{x}}(c) \right] &= \mathbb{E} \left[ \sum_{s=1}^t (\hat{x}_s(c_s) \hat{x}_s(c_s)^\top + \Sigma_n(c) - \Sigma_n(c_s)) \right] \\ &= \mathbb{E} \left[ \sum_{s=1}^t ((\Sigma_x + \Sigma_n(c_s)) + \Sigma_n(c) - \Sigma_n(c_s)) \right] \\ &= \mathbb{E} \left[ \sum_{s=1}^t (\Sigma_x + \Sigma_n(c)) \right] = t \Sigma_{\hat{x}}(c). \end{aligned}$$

Inspired by this, we add a correction term of the form  $\nu^\top (\Sigma_n(c) - \Sigma_n(c_s)) \nu$  to the natural empirical loss and define the following loss estimator:

$$\hat{L}_t^{kc}(c, \nu) = \sum_{s=1}^t \left( (\hat{x}_s^\top \nu - y_s)^2 + \nu^\top (\Sigma_n(c) - \Sigma_n(c_s)) \nu + \lambda c \right).$$

We indeed show that this empirical loss uniformly converges to the expected loss  $\ell(c, \nu)$  and is almost convex:

**Proposition 2.** With probability at least  $1 - 3\delta$ , for all  $t \geq 1$ ,  $c \in [0, 1]$  and  $\nu \in \mathbb{R}^d$  s.t.  $\|\nu\| \leq S$ , it holds that

$$\left| \hat{L}_t^{kc}(c, \nu) - t\ell(c, \nu) \right| \leq 9S^2 \sqrt{8t\beta_t^2 \ln \frac{3dt(t+1)}{\delta}}$$

where

$$\begin{aligned} \beta_t &= R^2 \left( d + 2\sqrt{d \ln \frac{3t(t+1)}{\delta}} + 2 \ln \frac{3t(t+1)}{\delta} \right) \\ &\quad + S^2 \left( 1 + 2\sqrt{\frac{\ln \frac{3t(t+1)}{\delta}}{d}} \right) = \tilde{O}(R^2 d + S^2). \end{aligned}$$

Furthermore, under the same event, the regularized loss

$$\hat{L}_t^{kc,R}(c, \nu) = \hat{L}_t^{kc}(c, \nu) + \gamma_t \|\nu\|^2 \quad (1)$$

for  $\gamma_t = 2\sqrt{8t\beta_t^2 \ln \frac{3dt(t+1)}{\delta}}$  is quadratic and strictly convex in  $\nu$  for all  $t \geq 1$ .

*Proof Sketch.* As a first step, we show how to decompose the loss difference into the following quadratic terms:

$$\begin{aligned} \hat{L}_t^{kc}(c, \nu) - t\ell(c, \nu) &= \nu^\top \left( \sum_{s=1}^t (\hat{x}_s \hat{x}_s^\top - \Sigma_x - \Sigma_n(c_s)) \right) \nu \\ &\quad + \theta^{*\top} \left( \sum_{s=1}^t x_s x_s^\top - \Sigma_x \right) \theta^* \\ &\quad - 2\nu^\top \left( \sum_{s=1}^t x_s x_s^\top - \Sigma_x \right) \theta^* \\ &\quad - 2\nu^\top \left( \sum_{s=1}^t n_s x_s^\top \right) \theta^*. \end{aligned}$$

In particular, the first term already incorporates the covariance transfer from  $\Sigma_n(c_s)$  to  $\Sigma_n(c)$ , leaving a well-behaved error term that depends on all the past payments, even when different costs were paid. The first three terms are quadratic forms w.r.t. the error between the empirical covariance of subgaussian vectors and their real covariance. The last term can also be written as such, noting that

$$\nu^\top n_s x_s^\top \theta^* = \begin{pmatrix} 0 \\ \nu \end{pmatrix}^\top \left( \begin{pmatrix} x_s \\ n_s \end{pmatrix} \begin{pmatrix} x_s \\ n_s \end{pmatrix}^\top - \begin{pmatrix} \Sigma_x & 0 \\ 0 & \Sigma_n(c_s) \end{pmatrix} \right) \begin{pmatrix} \theta^* \\ 0 \end{pmatrix}.$$

Using the inequality  $x^\top A y \leq \|x\| \|y\| \|A\|_{op}$ , and given that the norms of  $\nu, \theta^*$  are bounded, we can uniformly bound the loss difference by upper bounding the operator norms of the covariance error terms. To bound these norms, we adapt martingale matrix concentration results (Tropp 2012) and apply them to the empirical covariance of subgaussian

---

**Algorithm 1: Online Regression with Paid Features – Known Noise Covariances**


---

**Require:**  $\delta \in (0, 1)$ ,  $K \in \mathbb{N}$   
**Initialize:**  $\hat{L}_0^{kc,R}(c, \nu) = 0$   
**for**  $t = 1, \dots, T$  **do**  
  **for**  $k = 1, \dots, K$  **do**  
    Calculate  $\hat{\nu}_{t-1}(k) \in \arg \min_{\nu: \|\nu\| \leq S} \hat{L}_{t-1}^{kc,R}(k/K, \nu)$   
  **end for**  
  Set  $k_{t-1} \in \arg \min_{k \in [K]} \hat{L}_{t-1}^{kc,R}(k/K, \hat{\nu}_{t-1}(k))$   
  Pay  $c_t = k_{t-1}/K$  and observe  $\hat{x}_t = \hat{x}_t(c_t)$   
  Predict  $\hat{y}_t = \hat{x}_t^\top \hat{\nu}_{t-1}(k_{t-1})$  and observe  $y_t$   
  Update  $\hat{L}_t^{kc,R}(k/K, \nu)$  for all  $k \in [K]$   
**end for**

---

vectors. In particular, if  $Z_t$  is a conditionally subgaussian vector martingale of conditional covariance  $\Sigma_t$ , it holds with high probability that

$$\left\| \sum_{s=1}^t Z_s Z_s^\top - \Sigma_s \right\|_{op} < \gamma_t/2$$

(see Lemma 6 in Appendix D for the precise statement and proof). This leads to the uniform concentration of the loss. One direct implication is the strict convexity of the regularized loss – noting that the loss is quadratic in  $\nu$  and that  $\hat{x}_t$  is  $2R^2$ -subgaussian, a regularization term of  $\gamma_t \|\nu\|^2$  is enough to turn the quadratic loss term into strictly convex.  $\square$

**Remark 3.** *The concentration bound of Proposition 2 can be easily extended to the case where the regression output is noisy, namely,  $y_t = x_t^\top \theta^* + \eta_t$  for conditionally zero-mean subgaussian noise  $\eta_t$ . This extra noise does not affect the algorithm or the concentration rates; see further discussion in Appendix B.2.*

A natural algorithmic approach is to minimize the (regularized) empirical loss – which we indeed do in Algorithm 1. Specifically, the algorithm minimizes this loss (w.r.t. the linear predictor  $\nu$ ) for every cost on a dense discretized grid  $c \in \{k/K\}_{k \in [K]}$  and pays the cost

$$c_t \in \arg \min_{c \in \{k/K\}_{k \in [K]}} \min_{\nu: \|\nu\| \leq S} \hat{L}_t^{kc,R}(c, \nu)$$

that achieves the minimal empirical loss on this grid. We note that the grid is only used to facilitate the optimal cost calculation; as long as it is sufficiently fine, it does not affect performance. Then, upon observing  $\hat{x}_t(c_t)$ , a prediction  $\hat{y}_t(\hat{x}_t) = \hat{x}_t^\top \nu_t$  is calculated using the best empirical linear predictor for  $c_t$ :

$$\nu_t = \hat{\nu}_{t-1}(c_t) \in \arg \min_{\nu: \|\nu\| \leq S} \hat{L}_t^{kc,R}(c_t, \nu).$$

This algorithm enjoys the following performance bound (see Appendix B for the proof):

**Theorem 1.** *For any  $T \geq 1$ , set  $\delta = 1/T$  and  $K = \lceil \lambda T \rceil$ . Then, the regret of Algorithm 1 is bounded by*

$$\text{Reg}(T) = \tilde{\mathcal{O}}\left(S^2(R^2 d + S^2)\sqrt{T} + \lambda\right).$$

*Proof Sketch.* For any  $c$ , the uniform concentration in Proposition 2 guarantees that the empirical loss  $\hat{L}_{t-1}^{kc}$  is similar to the expected one  $\ell$  both for  $\hat{\nu}_{t-1}(c)$  and for  $\nu^*(c)$ . Since  $\hat{\nu}_{t-1}(c)$  outperforms  $\nu^*(c)$  on the empirical loss  $\hat{L}_{t-1}^{kc}$ , the loss similarity implies that it cannot be much worse than  $\nu^*(c)$  on the expected loss  $\ell$  for any payment  $c \in [0, 1]$ . The suboptimality of  $\hat{\nu}_{t-1}(c)$  compared to  $\nu^*(c)$  will be proportional to the maximal error between the average empirical loss  $\frac{1}{t} \hat{L}_t^{kc}$  and the true loss  $\ell$ : by Proposition 2, it is of order  $\tilde{\mathcal{O}}(S^2 \beta_t / \sqrt{t})$ .

As for the choice of payment, the information sharing enables us to collect data on all costs simultaneously, without the need for explicit exploration. Thus, we can choose the best payment greedily on a payment grid, and the uniform concentration will again imply that the best empirical payment on the grid will perform similarly to the best cost on the grid (up to a similar error of  $\tilde{\mathcal{O}}(S^2 \beta_t / \sqrt{t})$ ). All that remains is to relate the optimal cost on the chosen grid to the best continuous payment: we do so leveraging the one-sided Lipschitzness of the loss, as stated in Claim 2. In particular, we choose a sufficiently fine grid to ensure that the discretization error is negligible. Thus, the instantaneous error compared to the optimal loss  $\ell^*$  is of order  $\tilde{\mathcal{O}}(S^2 \beta_t / \sqrt{t})$ , and accumulating it across all rounds leads to the desired regret bound.  $\square$

We now discuss two important aspects of our algorithm: its computational tractability and statistical efficiency (optimality of the regret bound).

**Tractability of the loss minimization.** One caveat in the covariance correction term  $\Sigma_n(c) - \Sigma_n(c_s)$  is that it might render the loss nonconvex. To mitigate this, we add a regularization term  $\gamma_t \|\nu\|^2$  to our loss. Then, we show in Proposition 2 that with high probability, calculating  $\hat{\nu}_{t-1}(k)$  requires minimizing a strictly convex loss over a ball, which can be done efficiently. In particular, cases where the objective is nonconvex on some rounds are of low probability, so the learner can just choose arbitrary costs and predictions without affecting the performance. Nonetheless, it is important to note that the regularization is not strictly necessary for computational efficiency; the minimization of the unregularized empirical loss  $\hat{L}_t^{kc}$  is in fact an instance of the trust region subproblem that can be efficiently solved even if the quadratic problem is nonconvex (see, e.g., Section 8.2.7 in Beck 2014). We also note that one could further increase the regularization term to make the objective  $\Omega(R^2 d + S^2)$ -strongly convex, while only deteriorating the regret by constant factors, thereby greatly accelerating the loss minimization step.

**Statistical efficiency.** At a glance, the regret rate that we obtained does not seem tight: the information sharing makes the problem similar to online linear regression, for which an  $\mathcal{O}(d \ln T)$  regret is achievable (Vovk 1997). Somewhat surprisingly, we show that this is not the case in our setting: even in one-dimensional problems, an  $\Omega(\sqrt{T})$  is unavoidable (see proof in Appendix B.1):

**Theorem 2.** *For any algorithm and  $T \geq 1$ , there exists a one-dimensional instance with a known noise variance profile*

---

**Algorithm 2: Online Regression with Paid Features – Unknown Noise Covariances**


---

**Require:**  $\delta \in (0, 1)$ ,  $K \in \mathbb{N}$   
**for**  $k = 1, \dots, K$  **do**  
  Pay  $c_k = k/K$  and  $\hat{y}_k = 0$ ; observe  $\hat{x}_k = \hat{x}_k(c_k)$  and  $y_k$   
  Set  $\hat{L}_K^{uc}(k/K, \nu) = (\hat{x}_k^\top \nu - y_k)^2 + \lambda k/K$   
**end for**  
**for**  $t = K + 1, \dots, T$  **do**  
  **for**  $k = 1, \dots, K$  **do**  
    Calculate  $\hat{\nu}_{t-1}(k) \in \arg \min_{\nu: \|\nu\| \leq S} \hat{L}_{t-1}^{uc}(k/K, \nu)$   
  **end for**  
  choose  $k_{t-1}$  according to eq. (2)  
  Pay  $c_t = k_{t-1}/K$  and observe  $\hat{x}_t = \hat{x}_t(c_t)$   
  Predict  $\hat{y}_t = \hat{x}_t^\top \hat{\nu}_{t-1}(k_{t-1})$  and observe  $y_t$   
  Update  $\hat{L}_t^{uc}(k_{t-1}/K, \nu)$   
**end for**

---

such that  $\mathbb{E}[R(T)] \geq \frac{\sqrt{T}}{240}$ .

The construction fixes the feature noise to be  $\mathcal{N}(0, 1)$  for all payments  $c < 1/2$  and the noise to be deterministically zero when  $c \geq 1/2$ . We then create two instances with known  $\theta^* = 1$  and  $x_t \sim \mathcal{N}(0, 1 \pm \epsilon)$  and show that, depending on the instance, it is optimal to play either  $c = 0$  or  $c = 1/2$ , thus reducing the problem to a two-armed bandit instance. From there, we use techniques from (Lattimore and Szepesvári 2020) to derive the lower bound. We hypothesize that the  $\sqrt{T}$  rate results from the discontinuity in the noise variance, and that for sufficiently smooth noise, a rate of  $\mathcal{O}(\ln T)$  might still be achievable, but leave this study for future work.

As a final remark, we note that the lower bound can be immediately generalized to  $d$ -dimensional features. In particular, we can choose a feature distribution that first samples a single coordinate (uniformly at random) to be non-zero, and then assigns its value according to the lower bound construction. This effectively creates  $d$  independent instances, each running for approximately  $T/d$  rounds. The regret of each instance will be lower bounded by  $\Omega(\sqrt{T/d})$ , so the overall regret from the interaction will be bounded by  $\Omega(d\sqrt{T/d}) = \Omega(\sqrt{Td})$ .

## 5 Extension to Unknown Noise Covariances

We now tackle the more challenging case in which the noise covariances are not known in advance. In this situation, we unfortunately cannot share information between different costs as we did in the previous section, and resort to estimating the loss at a cost  $c$  only based on samples collected with this specific cost. Formally, we define the loss

$$\hat{L}_t^{uc}(c, \nu) = \sum_{s=1}^t \left( (\hat{x}_s^\top \nu - y_s)^2 + \lambda c \right) \mathbb{1}\{c_s = c\},$$

which fully overlaps with  $\hat{L}_t^{kc}(c, \nu)$  when all samples are collected with  $c_s = c$ . Leveraging this insight, we utilize Proposition 2 to derive concentration results on  $\hat{L}_t^{uc}(c, \nu)$  (see proof in Appendix C).

**Corollary 3.** Fix  $c \in [0, 1]$  and let  $N_t(c) = \sum_{s=1}^t \mathbb{1}\{c_s = c\}$ . Then, with probability at least  $1 - 3\delta$ , for all  $t \geq 1$  and  $\nu \in \mathbb{R}^d$  s.t.  $\|\nu\| \leq S$ , it holds that

$$\left| \hat{L}_t^{uc}(c, \nu) - N_t(c) \ell(c, \nu) \right| \leq 9S^2 \sqrt{8N_t(c) \beta_t^2 \ln \frac{3dt(t+1)}{\delta}},$$

where  $\beta_t = \tilde{\mathcal{O}}(R^2 d + S^2)$  is specified at Proposition 2.

We remark that without the information sharing term  $\nu^\top (\Sigma_n(c) - \Sigma_n(c_s)) \nu$ , the loss  $\hat{L}_t^{uc}(c, \nu)$  is always convex; therefore, the regularization is no longer needed and is omitted from the corollary.<sup>3</sup> A byproduct of separately estimating the loss for each cost is the need for exploration: as playing one cost does not yield sufficient information on all potential payments, we must test a diverse set of payments to identify the optimal cost  $c^*$ . We encourage our algorithm to explore by adding *optimism* to  $\hat{L}_t^{uc}$  (Auer 2002), that is, replacing the empirical loss by the smallest plausible loss allowed by the confidence interval. Intuitively, the optimism introduces penalty terms that reduce the loss at payment levels that were not sufficiently observed, biasing toward their exploration.

The full algorithm is depicted in Algorithm 2. We first limit the learner to a fixed grid of potential payments  $c \in \{k/K\}_{k \in [K]}$ ; exploration is costly, so only these costs will be explored. We play each of these costs once as a loss initialization step. Then, at each round, the algorithm chooses the cost index that minimizes the optimistic average loss:

$$k_{t-1} \in \arg \min_{k \in [K]} \left\{ \frac{\hat{L}_{t-1}^{uc}(k/K, \hat{\nu}_{t-1}(k))}{N_{t-1}(k/K)} - 9S^2 \sqrt{\frac{8\beta_t^2 \ln \frac{3dt(t+1)}{\delta}}{N_{t-1}(k/K)}} \right\}, \quad (2)$$

where

$$\hat{\nu}_{t-1}(k) \in \arg \min_{\nu: \|\nu\| \leq S} \hat{L}_{t-1}^{uc}(k/K, \nu)$$

is a linear predictor that minimizes the empirical loss. As before, after paying  $c_t = k_{t-1}/K$ , the learner observes  $\hat{x}_t(c_t)$  and uses  $\hat{\nu}_{t-1}(k_{t-1})$  to predict  $\hat{y}_t = \hat{x}_t(c_t)^\top \hat{\nu}_{t-1}(k_{t-1})$ . Afterwards, upon observing  $y_t$ , the learner only updates the loss at  $c_t$  and continues to the next round.

The algorithm enjoys the following regret bound (see Appendix C for the proof):

**Theorem 3.** For any  $T \geq 1$ , set  $K = \left\lceil \frac{T^{1/3} \lambda^{2/3}}{(S^2(R^2 d + S^2))^{2/3}} \right\rceil$  and  $\delta = \frac{1}{KT}$ . Then, the regret of Algorithm 2 is bounded by

$$\text{Reg}(T) = \tilde{\mathcal{O}}\left( (S^2(R^2 d + S^2))^{2/3} \lambda^{1/3} T^{2/3} \right).$$

*Proof Sketch.* The proof again relies on the loss concentration (Corollary 3), now combined with bandit techniques. In particular, we show that with high probability, the instantaneous regret of playing a cost  $c_t$  compared to the best cost on the grid  $\{k/K\}_{k \in [K]}$  is inversely proportional to

<sup>3</sup>The regularization could still be added to ensure strong convexity without affecting the results while requiring only minor algorithmic modifications.

the number of times it was previously played: of the order  $\tilde{\mathcal{O}}\left(S^2\beta_{t-1}/\sqrt{N_{t-1}(c_t)}\right)$ . On the other hand, working with a grid of  $K$  costs leads to an instantaneous discretization error of  $\lambda/K$  (due to the Lipschitzness proved in Claim 2). Therefore, the cumulative regret is bounded (w.h.p.) by

$$\text{Reg}(T) \lesssim \sum_{t=K+1}^T \left( \frac{S^2\beta_T}{\sqrt{N_{t-1}(c_t)}} + \frac{\lambda}{K} \right) \approx S^2\beta_T\sqrt{KT} + \frac{\lambda T}{K}.$$

The last relation uses standard bandit arguments: upon choosing a cost  $c_t$ , its corresponding count must increase, decreasing the denominator in future rounds in which this cost is played. Thus, the sum of the inverse counts cannot be too big. Fixing a grid size that optimally balances the two terms leads to the stated regret bound.  $\square$

**Relation to Lipschitz Bandits.** The discretization approach on the costs and local loss estimation resembles existing approaches in the Lipschitz bandit literature (also called  $\mathcal{X}$ -armed bandits, Bubeck et al. 2011); there, a regret bound of  $\tilde{\mathcal{O}}(T^{2/3})$  is obtained for one-dimensional problems. Indeed, although our problem is not Lipschitz, the one-sided Lipschitzness (Claim 2) might suffice for such a reduction. However, the resulting Lipschitz instance will be  $d+1$ -dimensional (representing both costs and predictions), thus leading to a regret bound of  $\tilde{\mathcal{O}}(T^{\frac{d+2}{d+3}})$ . In contrast, we show how to separately optimize over the linear predictors, so that we can perform the discretization only in one dimension (somewhat similar to Tullii et al. 2024). This allows us to achieve an improved rate of  $\tilde{\mathcal{O}}(T^{2/3})$ .

Next, we discuss the tightness of Theorem 3. Our discretization-based bandit approach yields a worse regret compared to the  $\tilde{\mathcal{O}}(\sqrt{T})$  rate with known noise covariances. This degradation is unavoidable, as we prove in the following one-dimensional lower bound:

**Theorem 4.** *For any algorithm and  $T \geq 1$ , there exists a one-dimensional instance with an unknown noise variance profile such that  $\mathbb{E}[R(T)] \geq \frac{T^{2/3}}{256}$ .*

The full proof can be found at Appendix C.1.

*Proof Sketch.* To derive this bound, we first show that for one dimensional problems with the parameters  $\Sigma_x \triangleq \sigma_x^2 = 1$ ,  $\theta^* = 1$  and  $\lambda = 1/2$ , the variance profile  $\Sigma_n(c) \triangleq \sigma_n^2(c) = f(c) = \frac{1-c}{1+c}$  achieves a constant loss  $\ell^*(c) = 1/2$  for all  $c \in [0, 1]$ . Then, we carefully devise  $K$  noise variance profiles (‘instances’) that smoothly deviate from  $f(c)$  only at a small interval (‘modified interval’) of size  $\approx 1/K$ , with no interval overlaps between instances. In particular, the minimal loss inside these intervals is slightly lower than  $1/2$  (by  $\approx 1/K$ ). We remark that although our problem is very different from Lipschitz bandits, our construction draws inspiration from characteristics of hard Lipschitz bandit instances, and we indeed obtain the same bound as Kleinberg (2004).

We then compare the behavior of any algorithm in these instances to that on a nominal instance with  $\sigma_n^2(c) = f(c)$ . By the pigeonhole principle, there must be an instance  $k^*$  such that its modified interval was sampled less than  $T/K$

times on average in the nominal instance. Since the nominal instance and  $k^*$  are identical outside the interval (no information on the instance is gained), we show that for  $K \approx T^{1/3}$ , the algorithm cannot behave too differently on  $k^*$ ; using information-theoretic tools (Garivier, Ménard, and Stoltz 2019), we show that it chooses costs in the modified intervals no more than  $3T/4$  times. Since the regret for not choosing a cost in the modified interval is roughly  $1/K$  (loss of  $1/2$  instead of  $1/2 - 1/K$ ), this implies a lower bound of

$$\text{Reg}(T) \gtrsim \left(T - \frac{3T}{4}\right) \frac{1}{K} = \Omega(T^{2/3}).$$

$\square$

## 6 Summary and Future Work

In this work, we studied online linear regression problems with stochastic features that are corrupted by noise, and the learner can pay to reduce this noise. This setting is applicable, for example, in settings where features are measured in experiments, problems with privacy-protected information, and more. We focused on two variants of this problem: in the first, the learner has full knowledge of the noise covariance across all payment levels (e.g., privacy noise with known characteristics); in the second, the learner must estimate the effect of the payment on the quality of the observed features. In both cases, we devised learner algorithms that achieve order-optimal regret bounds as a function of the interaction length (up to logarithmic factors).

While the algorithms presented in this work achieve order-optimal horizon-dependence in their regret bounds, it is unclear whether these algorithms are optimal with respect to other problem parameters. In particular, all our lower bounds are derived on one-dimensional instances, and so the optimal dimension dependence remains to be determined. In addition, while both algorithms are polynomial, each time step requires solving a convex optimization problem over a grid of costs. It would be beneficial to derive more computationally efficient algorithms, including algorithms that build adaptive grids or avoid discretizing the payments.

Our work could also be extended beyond linear regression. Specifically, it would be interesting to extend our results to non-linear regression problems  $y_t = f(x_t)$  (for possibly non-linear  $f \in \mathcal{F}$ ), also potentially deviating from the quadratic loss. In this context, it is especially relevant to study *agnostic* settings: limiting the algorithm to a smaller class of predictors that does not include the mapping that minimizes the loss between  $\hat{y}(x_t)$  and  $y_t$ . This might be necessary since the optimal predictors could be extremely complicated, as we discussed in Remark 1. Specifically, we limited our predictors to be linear, similarly to the mapping that created the original outputs  $y_t$ , but for more complex regression models, it might be beneficial to tailor the prediction class  $\mathcal{P}$  to be different than the true regression model class  $\mathcal{F}$ .

Finally, in some situations, we need to pay for each feature separately and/or can obtain the same feature from multiple sources (for different payments and with different noise). It is then interesting to determine the optimal (combinatorial) payment scheme and devise algorithms that learn to optimally aggregate features gathered from different sources.

## Acknowledgments

NM was supported by Israel Science Foundation research grant (ISF's No. 4118/25) and the Maimonides Fund's Future Scientists Center. NCB acknowledges the financial support from the MUR PRIN grant 2022EKNE5K (Learning in Markets and Society), the EU Horizon CL4-2021-HUMAN-01 research and innovation action under grant agreement 101070617, project ELSA, and the FAIR (Future Artificial Intelligence Research) project, funded by the NextGenerationEU program within the PNRR-PE-AI scheme. KJ was supported by the Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) [RS-2021-II211341, Artificial Intelligence Graduate School Program (Chung-Ang University)].

## References

- Agarwal, A.; Harris, K.; Whitehouse, J.; and Wu, S. Z. 2023. Adaptive principal component regression with applications to panel data. *Advances in Neural Information Processing Systems*, 36: 77104–77118.
- Agarwal, A.; Shah, D.; Shen, D.; and Song, D. 2019. On robustness of principal component regression. *Advances in Neural Information Processing Systems*, 32.
- Agarwal, A.; and Singh, R. 2021. Causal Inference with Corrupted Data: Measurement Error, Missing Values, Discretization, and Differential Privacy. *arXiv preprint arXiv:2107.02780*.
- Auer, P. 2002. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3(Nov): 397–422.
- Beck, A. 2014. *Introduction to nonlinear optimization: Theory, algorithms, and applications with MATLAB*. SIAM.
- Belloni, A.; Rosenbaum, M.; and Tsybakov, A. B. 2017. Linear and conic programming estimators in high dimensional errors-in-variables models. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 79(3): 939–956.
- Bubeck, S.; Munos, R.; Stoltz, G.; and Szepesvári, C. 2011. X-Armed Bandits. *Journal of Machine Learning Research*, 12(5).
- Cesa-Bianchi, N.; Shalev-Shwartz, S.; and Shamir, O. 2011. Online learning of noisy data. *IEEE Transactions on Information Theory*, 57(12): 7907–7931.
- Chen, Y.; and Caramanis, C. 2013. Noisy and missing data regression: Distribution-oblivious support recovery. In *International Conference on Machine Learning*, 383–391. PMLR.
- Datta, A.; and Zou, H. 2017. Cocolasso for high-dimensional error-in-variables regression. *Annals of Statistics*, 45(6): 2400–2426.
- Duchi, J. C.; Jordan, M. I.; and Wainwright, M. J. 2013. Local privacy, data processing inequalities, and minimax rates. *arXiv preprint arXiv:1302.3203*.
- Fallah, A.; Makhdomi, A.; Malekian, A.; and Ozdaglar, A. 2024. Optimal and differentially private data acquisition: Central and local mechanisms. *Operations Research*, 72(3): 1105–1123.
- Figuerola-Zúñiga, J. I.; Bayes, C. L.; Leiva, V.; and Liu, S. 2022. Robust beta regression modeling with errors-in-variables: a Bayesian approach and numerical applications. *Statistical Papers*, 63(3): 919–942.
- Fukuchi, K.; Tran, Q. K.; and Sakuma, J. 2017. Differentially private empirical risk minimization with input perturbation. In *International Conference on Discovery Science*, 82–90. Springer.
- Fuller, W. A. 2009. *Measurement error models*. John Wiley & Sons.
- Garivier, A.; Ménard, P.; and Stoltz, G. 2019. Explore first, exploit next: The true shape of regret in bandit problems. *Mathematics of Operations Research*, 44(2): 377–399.
- Ghosh, A.; and Roth, A. 2011. Selling privacy at auction. In *Proceedings of the 12th ACM conference on Electronic commerce*, 199–208.
- Hsu, D.; Kakade, S. M.; and Zhang, T. 2012. A tail inequality for quadratic forms of subgaussian random vectors. *Electronic Communications in Probability*, 17: 1–6.
- Hsu, J.; Gaboardi, M.; Haeblerlen, A.; Khanna, S.; Narayan, A.; Pierce, B. C.; and Roth, A. 2014. Differential privacy: An economic method for choosing epsilon. In *2014 IEEE 27th Computer Security Foundations Symposium*, 398–410. IEEE.
- Kasiviswanathan, S. P.; Lee, H. K.; Nissim, K.; Raskhodnikova, S.; and Smith, A. 2011. What can we learn privately? *SIAM Journal on Computing*, 40(3): 793–826.
- Kaul, A.; and Koul, H. L. 2015. Weighted  $\ell_1$ -penalized corrected quantile regression for high dimensional measurement error models. *Journal of Multivariate Analysis*, 140: 72–91.
- Kim, J.-h.; Yun, S.-Y.; Jeong, M.; Nam, J.; Shin, J.; and Combes, R. 2023. Contextual linear bandits under noisy features: Towards bayesian oracles. In *International Conference on Artificial Intelligence and Statistics*, 1624–1645. PMLR.
- Kleinberg, R. 2004. Nearly tight bounds for the continuum-armed bandit problem. *Advances in Neural Information Processing Systems*, 17.
- Lattimore, T.; and Szepesvári, C. 2020. *Bandit algorithms*. Cambridge University Press.
- Loh, P.-L.; and Wainwright, M. J. 2011. High-dimensional regression with noisy and missing data: Provable guarantees with non-convexity. *Advances in neural information processing systems*, 24.
- Nissim, K.; Orlandi, C.; and Smorodinsky, R. 2012. Privacy-aware mechanism design. In *Proceedings of the 13th ACM conference on electronic commerce*, 774–789.
- Nissim, K.; Vadhan, S.; and Xiao, D. 2014. Redrawing the boundaries on purchasing data from privacy-sensitive individuals. In *Proceedings of the 5th conference on Innovations in theoretical computer science*, 411–422.
- Reilly, P. M.; and Patino-Leal, H. 1981. A Bayesian study of the error-in-variables model. *Technometrics*, 23(3): 221–231.
- Rosenbaum, M.; and Tsybakov, A. B. 2010. Sparse Recovery under Matrix Uncertainty. *The Annals of Statistics*, 38(5): 2620–2651.

- Rosenbaum, M.; and Tsybakov, A. B. 2013. Improved matrix uncertainty selector. In *From Probability to Statistics and Back: High-Dimensional Models and Processes—A Festschrift in Honor of Jon A. Wellner*, volume 9, 276–291. Institute of Mathematical Statistics.
- Smith, A.; Thakurta, A.; and Upadhyay, J. 2017. Is interaction necessary for distributed private learning? In *2017 IEEE Symposium on Security and Privacy (SP)*, 58–77. IEEE.
- Tropp, J. A. 2012. User-friendly tail bounds for sums of random matrices. *Foundations of computational mathematics*, 12: 389–434.
- Tullii, M.; Gaucher, S.; Merlis, N.; and Perchet, V. 2024. Improved algorithms for contextual dynamic pricing. *Advances in Neural Information Processing Systems*, 37: 126088–126117.
- Ungarala, S.; and Bakshi, B. R. 2000. A multiscale, Bayesian and error-in-variables approach for linear dynamic data rectification. *Computers & Chemical Engineering*, 24(2-7): 445–451.
- Van Der Hoeven, D.; Pike-Burke, C.; Qiu, H.; and Cesa-Bianchi, N. 2023. Trading-off payments and accuracy in online classification with paid stochastic experts. In *International Conference on Machine Learning*, 34809–34830. PMLR.
- Vovk, V. 1997. Competitive on-line linear regression. *Advances in Neural Information Processing Systems*, 10.
- Wang, D.; and Xu, J. 2019. On sparse linear regression in the local differential privacy model. In *International Conference on Machine Learning*, 6628–6637. PMLR.
- Wang, W.; Ying, L.; and Zhang, J. 2018. The value of privacy: Strategic data subjects, incentive mechanisms, and fundamental limits. *ACM Transactions on Economics and Computation (TEAC)*, 6(2): 1–26.
- Yang, M.; Guo, T.; Zhu, T.; Tjuawinata, I.; Zhao, J.; and Lam, K.-Y. 2024. Local differential privacy and its applications: A comprehensive survey. *Computer Standards & Interfaces*, 89: 103827.
- Zheng, K.; Cai, T.; Huang, W.; Li, Z.; and Wang, L. 2020. Locally differentially private (contextual) bandits learning. *Advances in Neural Information Processing Systems*, 33: 12300–12310.