Representations of Uncertainty for Tractable Estimation and Control

PhD Oral Examination

Alexandros Tzikas

Stanford Intelligent Systems Lab

April 27, 2026

Motivation: Uncertainty is challenging

- Modern control under uncertainty is **high-dimensional**, **data-driven**, and **sequential**
- Models of uncertainty should be **computationally tractable** and **efficiently learned** from data
- Algorithms for control should **scale** and systematically **incorporate** newly revealed information

A unifying principle

Data samples / distribution → **Representation** simplify & keep what matters → **Control** convex / gradient / MPC

- Different control tasks depend on different aspects of uncertainty
- We can choose a representation of uncertainty that retains only the relevant aspects
- We will present three such representations for use in control

Contributions

1. Enhancement of a *factor model* by adding new statistical factors
   (with applications to financial portfolio construction)

Paper 1 [working paper]: Structural representation
A principled method to extend a provided covariance model

low-rank + diagonal covariance model

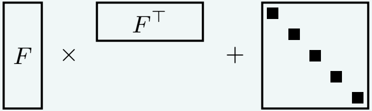

$$\Sigma \approx FF^\top + D$$

Contributions

2. A family of *distances* between probability distributions
   (with applications to distribution-aware control)

Paper 2 [ACC 2026]: Geometric representation
sliced (1D) projections of distributions and gradient-based distribution shaping

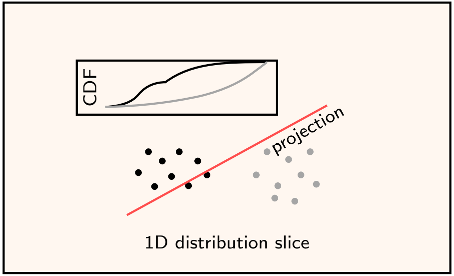

Contributions

3. *Conditional expectation* and convex optimization for sequential allocation under stochastic demands

Paper 3 [ACC 2026]: Informational representation
conditioning + shrinking-horizon re-optimization

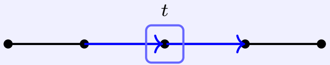

$$\mathbb{P}(\cdot) \mapsto \mathbb{P}(\cdot \mid \mathcal{F}_t)$$

A unifying problem: Financial portfolio construction

Structural.
Enhance the risk model
$\Sigma = \mathbf{cov}(r)$

$$\begin{aligned} \max_{w}.\quad & \boxed{\mathbb{E}[r]^\top w} \\ \text{s.t.}\quad & w^\top \boxed{\Sigma} w \leq \sigma^2. \end{aligned}$$

Geometric.
A more expressive objective measuring the distance to the desired distribution of portfolio returns

Informational. Incorporate newly revealed information, e.g., $\mathbb{E}\left[r \mid \mathcal{F}_t\right]$

Outline

Contribution 1: Problem statement

We are given

1. (approx. zero-mean) returns on n assets $r_1, \ldots, r_T \in \mathbb{R}^n$
2. a *base* factor model

$$r \approx F_{\text{base}} s + \epsilon, \quad s \sim \mathcal{N}(0, \Omega_{\text{base}}) \quad \perp\!\!\!\perp \quad \epsilon \sim \mathcal{N}(0, D_{\text{base}})$$

   D_{base} is diagonal (idiosyncratic risk),
   $F_{\text{base}} \in \mathbb{R}^{n \times k_1}$ are factor exposures ($k_1 \ll n$),
   factor returns s and idiosyncratic returns ϵ

Why use a factor model?

- **Interpretability:** expose shared drivers (factors) behind many variables [Grinold and Kahn, 2000]
- **Regularization:** low-rank structure reduces estimation variance and overfitting [Fan et al., 2008]
- **Fast downstream optimization:** exploit $\Sigma \approx FF^\top + D$ to accelerate risk computations

$$w^\top \Sigma w \leq \sigma^2 \quad \text{cost: } \mathcal{O}(n^3)$$

$$\left\| \begin{bmatrix} F^\top w \\ D^{1/2} w \end{bmatrix} \right\| \leq \sigma \quad \text{cost: } \mathcal{O}(nk^2)$$

[Boyd and Vandenberghe, 2004]

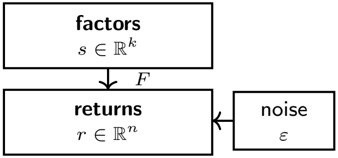

$$r = Fs + \varepsilon$$

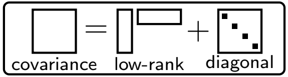

What can the base factor model miss?

- It is typically updated infrequently (e.g., monthly) [MSCI Inc.]
- It may fail to capture transient factors or shifts in market regimes [Lee et al., 2025]
- It may be missing (statistical) factors that explain the asset returns
- These statistical factors seem to be capturing *themes* [Candès et al., 2025]
  - groups of stocks with highly correlated residualized returns

The extended factor model

Extend *base* model to include (transient) themes and/or missing factors

$$r \approx \begin{bmatrix} F_{\text{base}} & F_2 \end{bmatrix} \begin{bmatrix} s^{(1)} \\ s^{(2)} \end{bmatrix} + \epsilon$$

- $F_2 \in \mathbb{R}^{n \times k_2}$ are additional factor exposures/themes $(k_2 \ll k_1)$
- $s^{(2)}$ are new factor variables
- $\begin{bmatrix} s^{(1)} \\ s^{(2)} \end{bmatrix} \sim \mathcal{N}(0, \Sigma_f) \quad \perp\!\!\!\perp \quad \epsilon \sim \mathcal{N}(0, D)$

Equivalent parameterizations of the extended model

- (Σ_f, F_2, D):

$$\Sigma = \begin{bmatrix} F_{\text{base}} & F_2 \end{bmatrix} \Sigma_f \begin{bmatrix} F_{\text{base}} & F_2 \end{bmatrix}^\top + D \tag{full}$$

- (Ω, F_2, D):

$$\Sigma = \begin{bmatrix} F_{\text{base}} & F_2 \end{bmatrix} \begin{bmatrix} \Omega & 0 \\ 0 & I \end{bmatrix} \begin{bmatrix} F_{\text{base}} & F_2 \end{bmatrix}^\top + D \tag{identifiable}$$

- $(\Sigma_f, F_{\text{base},\perp}, D)$:

$$\Sigma = \begin{bmatrix} F_{\text{base}} & F_{\text{base},\perp} \end{bmatrix} \Sigma_f \begin{bmatrix} F_{\text{base}} & F_{\text{base},\perp} \end{bmatrix}^\top + D \tag{orthogonal}$$

where $F_{\text{base},\perp}^\top F_{\text{base}} = 0$

Maximum likelihood estimation

We learn $\theta := (\Omega, F_2, D)$ by maximizing

$$\sum_{\tau \leq T} w_\tau \log p_\theta(r_\tau)$$

using expectation–maximization (EM)

- $p_\theta(\cdot)$ is the density of the Gaussian $\mathcal{N}\left(0, \begin{bmatrix} F_{\text{base}} & F_2 \end{bmatrix} \begin{bmatrix} \Omega & 0 \\ 0 & I \end{bmatrix} \begin{bmatrix} F_{\text{base}} & F_2 \end{bmatrix}^\top + D\right)$
- w_τ are non-negative weights that sum to 1

We can accommodate missing returns

We obtain closed-form expressions for the EM update

Equivalence of objectives

The following are the same

$$\operatorname*{argmax}_{\theta} \sum_{\tau \le T} w_\tau \log p_\theta(r_\tau)$$

$$\operatorname*{argmin}_{\theta} \ \mathrm{KL}\left(\mathcal{N}(0, C) \quad \| \quad \mathcal{N}(0, \hat{\Sigma})\right)$$

- KL is the Kullback–Leibler divergence,

$$C = \sum_{\tau \le T} w_\tau r_\tau r_\tau^\top, \qquad \hat{\Sigma} = \begin{bmatrix} F_{\text{base}} & F_2 \end{bmatrix} \begin{bmatrix} \Omega & 0 \\ 0 & I \end{bmatrix} \begin{bmatrix} F_{\text{base}} & F_2 \end{bmatrix}^\top + D$$

Design choices

We use an exponentially weighted moving average

$$C = \gamma \sum_{t \leq T} \beta^{T-t} r_\tau r_\tau^\top$$

- $\beta \in (0, 1)$ is the forgetting factor, often expressed in terms of the half-life

$$H = -\log 2 / \log \beta$$

- $\gamma = \left(\sum_{t \leq T} \beta^{T-t} \right)^{-1} = \dfrac{1-\beta}{1-\beta^T}$ is the normalizing constant

Judging the quality of the factor model

At every date, we extend the provided base model using the observed returns

We look at:

- R^2 metric for predicting out-of-sample returns
- Out-of-sample log-likelihood
- Performance of a portfolio constructed using the risk model

Predicting out-of-sample returns

At date t, assume the risk model is $F_t F_t^\top + D_t$. Then:

1. learn the most likely factor returns given the next-day returns in the train set

$$\hat{s} \leftarrow \underset{s}{\operatorname{argmin}} \left(r_{t+1}^{\text{train}} - F_t^{\text{train}} s\right)^\top \left(D_t^{\text{train}}\right)^{-1} \left(r_{t+1}^{\text{train}} - F_t^{\text{train}} s\right) + s^\top s$$

2. predict the next-day returns in the test set

$$\hat{r}_{t+1}^{\text{test}} = F_t^{\text{test}} \hat{s}$$

3. compute the R^2

$$R_t^2 = 1 - \frac{\|r_{t+1}^{\text{test}} - \hat{r}_{t+1}^{\text{test}}\|^2}{\|r_{t+1}^{\text{test}} - 0\|^2}$$

Predicting out-of-sample returns

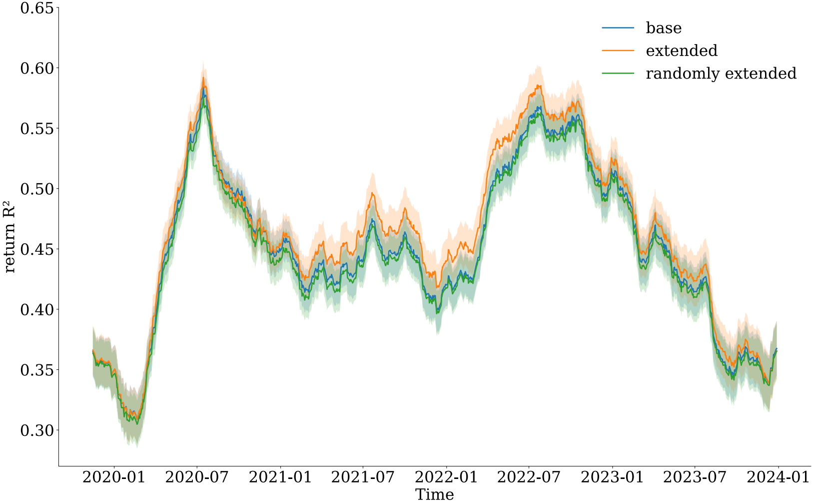

Log-likelihood

Model	Average Log-Likelihood	Std. of Average Log-Likelihood
base	2.679	0.039
extended	**2.726**	0.032

Table: Model comparison based on average out-of-sample (normalized) log-likelihood and its standard deviation.

Portfolio performance

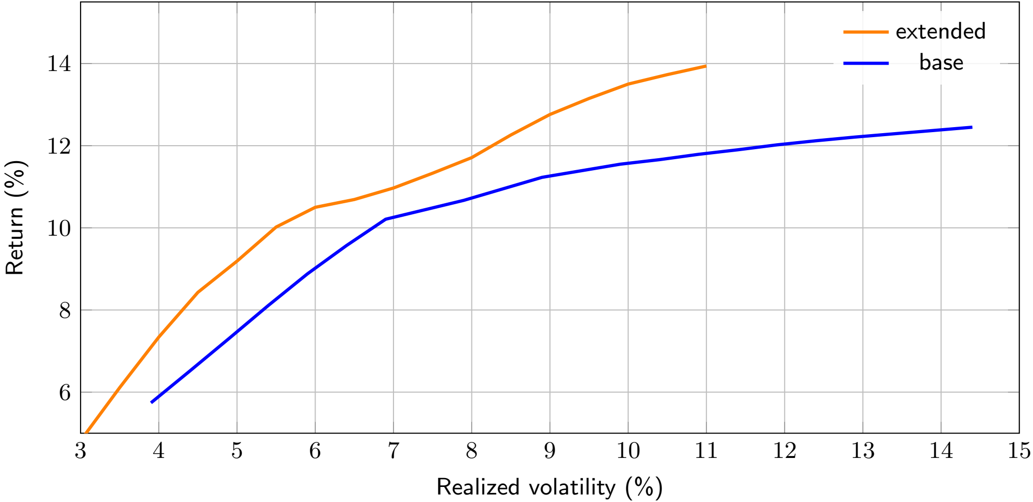

Limitations

- A risk model captures the covariance
- Many control objectives depend on the full distribution (e.g., multimodal uncertainty)
- We need to capture the shape of a distribution or the mismatch between distributions

This is the focus of contribution 2

Outline

Contribution 2: Problem statement

- Two probability distributions μ_x and μ_y over $\mathbb{R}^n$
- Each given by a density or by a collection of samples

What is a **computationally efficient**, **interpretable**, and **control-suitable** notion of similarity between μ_x and μ_y?

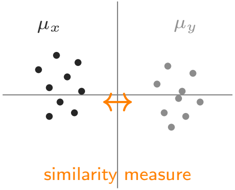

Motivation

- Many control applications: ergodic control [Dressel and Kochenderfer, 2019], and distribution steering [Rapakoulias and Tsiotras, 2024]
- Distances based on 1D projections (i.e., *sliced* distances) exist [Kolouri et al., 2019; Titouan et al., 2019]
- These are not widely used in control
- We unify prior work on sliced distances and propose simple gradient-based controllers for distribution shaping

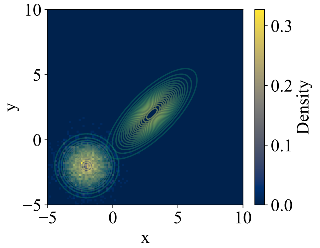

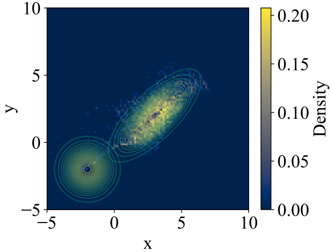

Our family of sliced distances

- Let $\mathcal{C}$ be the set of univariate CDFs
- Let $d : \mathcal{C} \times \mathcal{C} \to \mathbb{R}_+$ be a distance
- For fixed $\tilde{q} \in \mathbb{R}^n$ and random $X \in \mathbb{R}^n$, let $F_X^{\tilde{q}}$ be the CDF of $\tilde{q}^\top X$
    - $\tilde{q}^\top X$ is a 1D projection of X

For random $X, Y, q \in \mathbb{R}^n$ let

$$\Delta(X, Y) = \mathbb{E}_q \left[d\left(F_X^q, F_Y^q \right) \right]$$

be a measure of similarity between the distributions of X and Y

An illustration of the projection

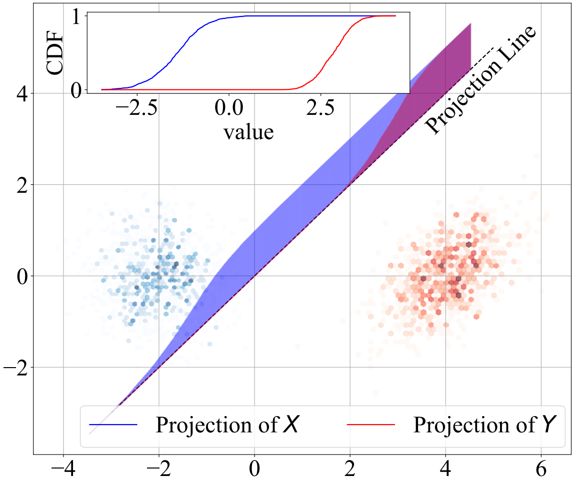

Properties

For random $X, Y, q \in \mathbb{R}^n$ let

$$\Delta(X,Y) = \mathbb{E}_q \left[d \left(F_X^q, F_Y^q \right) \right],$$

be a measure of similarity between the distributions of X and Y

Theorem: $\Delta(X,Y)$ is a distance in the space of distributions

This implies

$$\boxed{\Delta(X,Y) = 0 \Leftrightarrow X \overset{D}{=} Y}$$

Design choices

- For fixed q, $F_X^q(t)$ is easily obtained if, e.g., X is given by samples or as a GMM
- For example
$$d(F_X^q, F_Y^q) = \left(\int c(t) |F_X^q(t) - F_Y^q(t)|^p dt \right)^{1/p}, \quad p \geq 1$$
where $c(\cdot)$ is positive and integrates to 1
- $\Delta(X, Y) \in [0, 1]$
  - but it depends on the choice of $c(\cdot)$ and the distribution of q

An estimator of $\Delta(X, Y)$

- Let q_i be iid vectors on the unit sphere
- We use samples of X and Y to compute the empirical CDFs

$$\hat{F}_X^{q_i}(t), \qquad \hat{F}_Y^{q_i}(t)$$

The sample-based estimator of $\Delta(X, Y)$ is

$$\hat{\Delta}(X, Y) = \frac{1}{H} \sum_{i=1}^{H} d(\hat{F}_X^{q_i}, \hat{F}_Y^{q_i})$$

Replacing $\mathbf{1}(\cdot)$ with a steep sigmoid in the empirical CDFs, we get a differentiable estimator

Distinguishing $X \stackrel{D}{=} Y$ from $X \stackrel{D}{\neq} Y$

$$X_i \stackrel{\text{iid}}{\sim} \mathcal{N}(0,1)$$

$$Y_i \stackrel{\text{iid}}{\sim} \mathcal{N}(0,1) \qquad \text{or} \qquad Y_i \stackrel{\text{iid}}{\sim} \text{Laplace}\left(0, \frac{1}{\sqrt{2}}\right)$$

$$\mathbb{E}X = \mathbb{E}Y \qquad \text{and} \qquad \mathbf{cov}(X) = \mathbf{cov}(Y)$$

At each trial:

- generate $1,000$ samples from each distribution
- evaluate $\hat{\Delta}(X, Y)$ using 600 random q_i

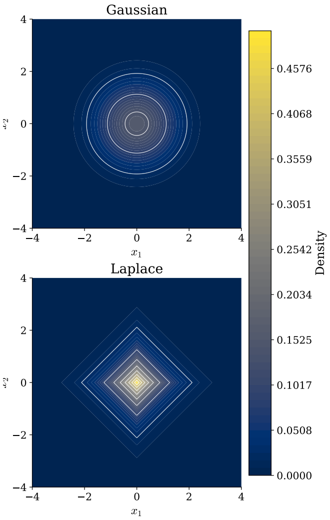

Distinguishing $X \stackrel{D}{=} Y$ from $X \stackrel{D}{\neq} Y$

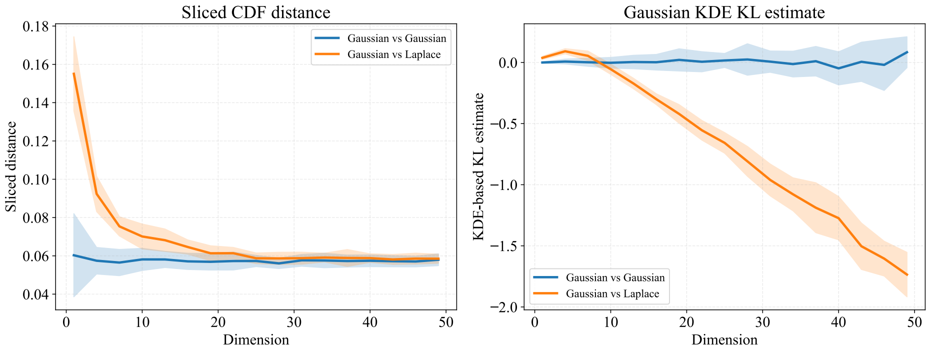

Demonstration: Ergodic control

Unicycle dynamics $x_{t+1} = f(x_t, u_t)$

Goal: Compute the input $\{u_t\}_{t=1}^{T-1}$ such that the empirical distribution of $\{x_t\}_{t=1}^{T}$ is close to the distribution μ_Y

Proposed approach: Minimize sliced distance between $\{x_t\}_{t=1}^{T}$ and μ_Y

$$u_t \leftarrow u_t - \rho \nabla_{u_t} \hat{\Delta} \left(\{x_\tau\}_{\tau=1}^{T}, \mu_Y \right)$$

Baseline approach: Maximize the log-likelihood of $\{x_t\}_{t=1}^{T}$ under μ_Y

Results

Proposed

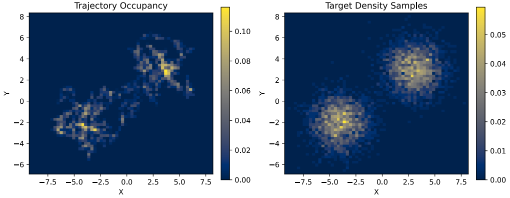

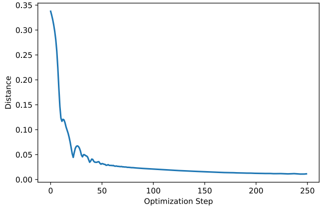

Baseline

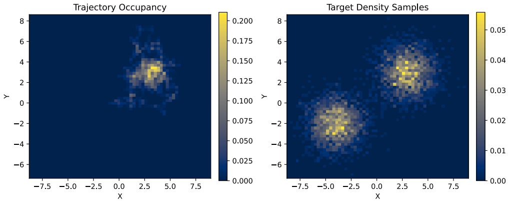

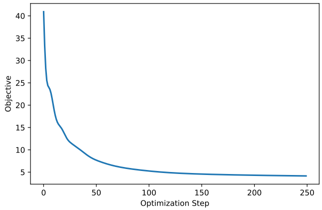

Limitations

- The proposed sliced distance handles expressive objectives

  but

  we need to select the half-spaces and the function $d(\cdot)$,

  performance degrades in high dimensions

- If the control objective depends only on statistical moments, how can we exploit the sequential problem structure to improve performance?

This is the focus of contribution 3

Outline

Contribution 3: Problem statement

Consider a sequential decision-making setting:

- we allocate limited resources a_t over time to satisfy stochastic demands,
- the demands d_t are realized over time

What is a good resource allocation policy?

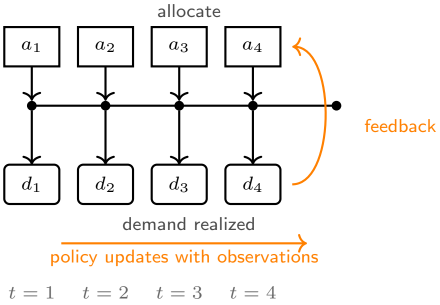

Motivation

- Extensive applications:
  - supply chain [Choi et al., 2008]
  - dynamic pricing [Monahan et al., 2004]
  - dispatch of energy resources [Marchi et al., 2019]
- We propose a simple and computationally efficient framework using conditional moments and convex optimization

Our approach

Unknown demands $d_1, \ldots, d_T \in \mathbb{R}_+$ and known prices $p_1, \ldots, p_T \in \mathbb{R}_+$ per unit of demand
At time 1, we solve

$$
\begin{array}{ll}
\underset{a_1,\ldots,a_T}{\max.} & \mathbb{E}_{d_1,\ldots,d_T}\left[\sum_{t=1}^{T} p_t \min\left(d_t, a_t\right)\right] \\
\text{s.t.} & \sum_{t=1}^{T} a_t = L \\
& a_1, \ldots, a_T \geq 0
\end{array}
$$

- the variables $a_1, \ldots, a_T \in \mathbb{R}$ are the allocated units
- the problem is convex
- we allocate $a_1^\star$ and observe the realized demand d_1

Solution using the KKT conditions

The objective can be re-written as

$$\sum_{t=1}^{T} \int_0^\infty \mathbb{P}\left(\min(d_t, a_t) \geq x\right) dx = \sum_{t=1}^{T} \int_0^{a_t} \mathbb{P}(d_t \geq x) dx$$

The primal–dual (KKT) optimality conditions are

$$\begin{aligned} p_t \mathbb{P}(d_t \geq a_t^\star) + \lambda_t^\star &= \nu^\star \\ \sum_{t=1}^{T} a_t^\star &= L \\ \lambda_t^\star a_t^\star &= 0 \\ \lambda_t^\star \geq 0, \qquad a_t^\star &\geq 0 \end{aligned}$$

Solution using the KKT conditions

$$a_t^\star(\nu^\star) = \begin{cases} F_t^{-1}\left(1 - \frac{\nu^\star}{p_t}\right), & \frac{\nu^\star}{p_t} \leq 1 \\ 0, & \text{otherwise} \end{cases}$$

- $F_t^{-1}(\cdot)$ is the inverse CDF of d_t
- $\sum_{t=1}^{T} a_t^\star(\nu)$ is decreasing in ν
- We compute $\nu^\star$ via bisection on

$$\sum_{t=1}^{T} a_t^\star(\nu) = L$$

Model for demands

Log-normal model for demands

$$\begin{bmatrix} \log d_1 \\ \vdots \\ \log d_T \end{bmatrix} \sim \mathcal{N}(\mu, \Sigma)$$

- demands are non-negative
- demands empirically exhibit right-skewness
- correlations across time are easily modeled

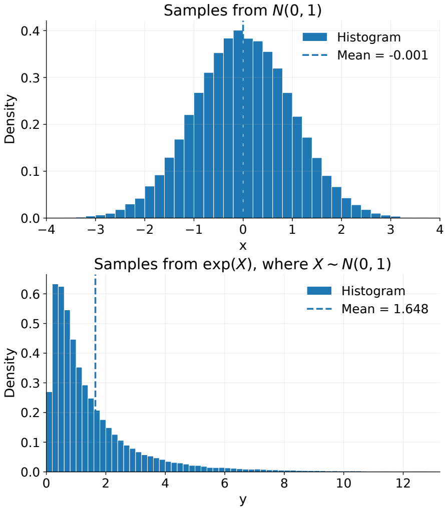

Our approach II

At time 2, we solve

$$\begin{aligned} \max_{a_2,\dots,a_T} \quad & \mathbb{E}_{d_2,\dots,d_T|d_1}\left[\sum_{t=2}^{T} p_t \min\left(d_t, a_t\right)\right] \\ \text{s.t.} \quad & \sum_{t=2}^{T} a_t = L - a_1^\star \\ & a_2,\dots,a_T \geq 0 \end{aligned}$$

If the demands are not independent, conditioning on the observed d_1 adds information

We can apply the same algorithm, after replacing F_t^{-1} with the conditional inverse CDF of d_t

Conditioning

Suppose

$$\log d_1, \ldots, \log d_T \sim \mathcal{N}(\mu, \Sigma)$$

Then

$$\mathbb{P}(d_t \leq x \mid d_1 = x_1, \ldots, d_{t-1} = x_{t-1}) =$$
$$\mathbb{P}(\log d_t \leq \log x \mid \log d_1 = \log x_1, \ldots, \log d_{t-1} = \log x_{t-1}) = \Phi\left(\frac{\log x - \mu_{t|t-1}}{\sigma_{t|t-1}}\right)$$

Φ is the CDF of the standard Gaussian and
$\mu_{t|t-1}, \sigma^2_{t|t-1}$ are the conditional moments of $\log d_t$ (given by standard formulas)

Our algorithm only requires evaluating these conditional CDFs

Results on synthetic data

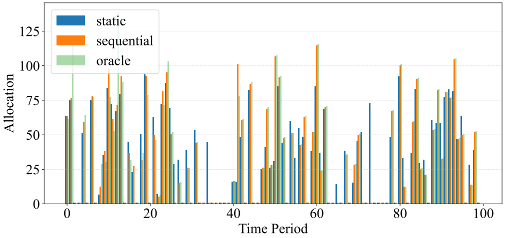

(a) Allocations

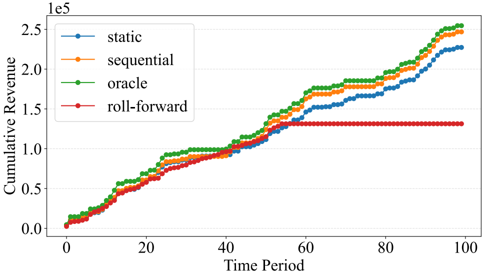

(b) Revenue

The proposed approach (*sequential*) achieves a revenue close to the regret bound

Limitations

- Our algorithm is a heuristic and does not guarantee global optimality
- We could replace the expected value with another point estimate (e.g., MAP)
- Uncertainty in prices is not (yet) incorporated

Outline

The big picture

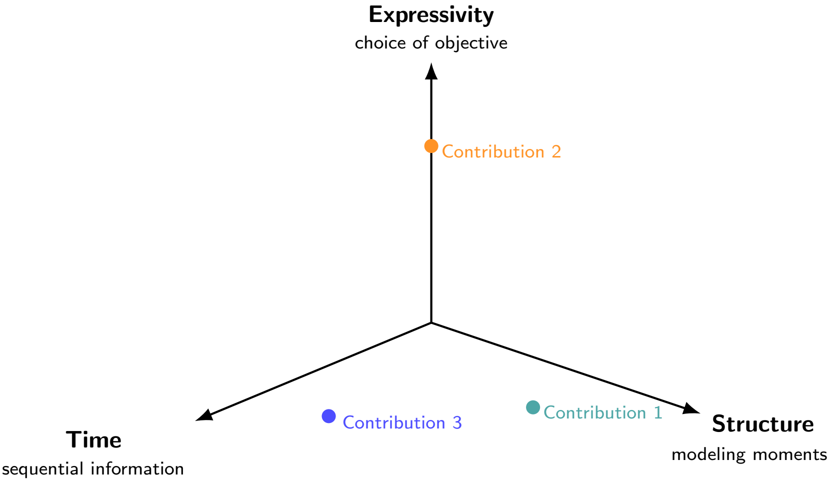

Next steps

Contribution 1

Search for an F_2 with a specified range

Contribution 2

Explore applications in generative modeling

Contribution 3

Evaluate in real-world settings

Outline

Acknowledgments

- Prof. Mykel Kochenderfer
- Prof. Stephen Boyd, Prof. Emmanuel Candès, Prof. Trevor Hastie
- Prof. Marco Pavone, Prof. Mac Schwager
- My collaborators: SISL, CVXGRP, BlackRock AI Lab
- My friends
- My family

Representations of Uncertainty for Tractable Estimation and Control

PhD Oral Examination

Alexandros Tzikas

Stanford Intelligent Systems Lab

April 27, 2026

Why KL instead of Frobenius norm?

$$r = Fs + \epsilon, \quad s \sim \mathcal{N}(0, I), \quad \epsilon \sim \mathcal{N}(0, 0.2^2 I)$$

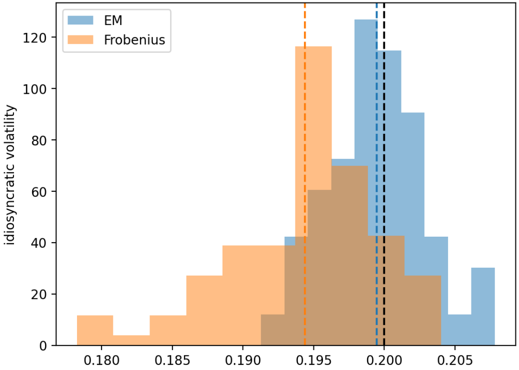

Figure: The Frobenius norm tends to underestimate the idiosyncratic volatility.

Why not used in control?

- Classical control rarely optimizes over full distributions
- The control variable affects the distribution only through the dynamics
- Not Bellman friendly
- Variance of the distance estimate
- Random projections make objective less interpretable
- Safety guarantees are harder

Distinguishing $X \overset{D}{=} Y$ from $X \overset{D}{\neq} Y$

$$X \sim \mathcal{N}(0, I_d),$$

$$Y \sim 0.9\mathcal{N}(0, I_d) + 0.1\mathcal{N}(0, 36I_d)$$

At each trial:

- generate 400 samples from each distribution
- evaluate $\hat{\Delta}(X, Y)$ using 128 random q_i

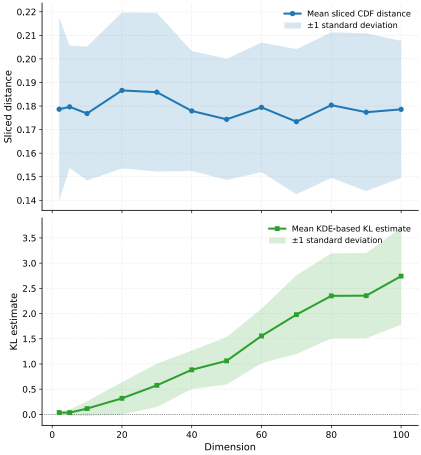

Control application: Distribution steering

Unicycle dynamics $x_{t+1} = f(x_t, u_t)$
$x_0 \sim \mu_{\text{start}}$ (given as set of samples $\{x_0^i\}_i$)

Goal: As $t \to \infty$, $x_t \sim \mu_Y$, where μ_Y is a given density

Approach: At any time t, given the state samples $\{x_t^i\}_i$, design the controller $\pi_t = (K_t, b_t)$, where $u_t = K_t x_t + b_t$,
by gradient descent

$$\pi_t \leftarrow \pi_t - \rho \nabla_{(K_t, b_t)} \hat{\Delta} \left(\{x_i^{(t+\tau)}(K_t, b_t)\}_i, \mu_Y \right),$$

(τ is the horizon and we use the sigmoid approximation in $\hat{\Delta}$)

Results: Distribution steering

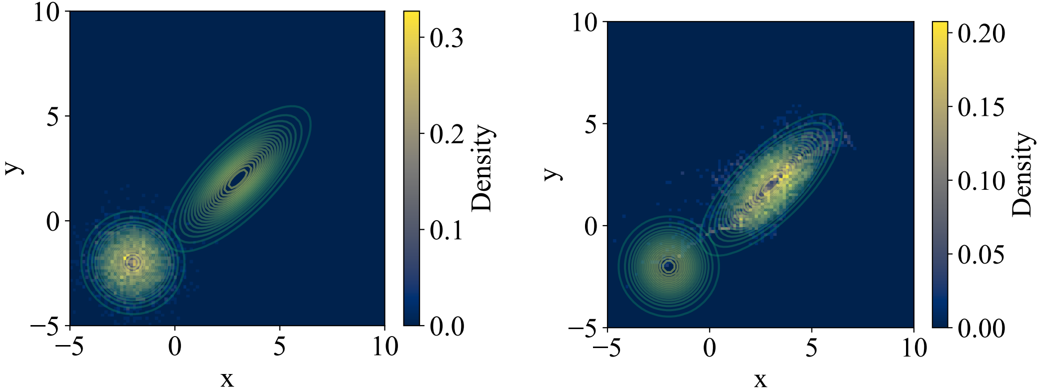

No model misspecification

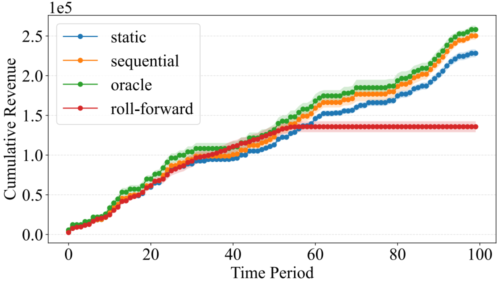

Figure: 20 trials.

Model misspecification: Marginal means

$$\mathbb{E}_{\text{model}}\left[d_t\right] = \mathbb{E}_{\text{real}}\left[d_t\right] + U([-5,5]), \quad \mathbb{E}_{\text{real}}\left[d_t\right] \in [20, 100]$$

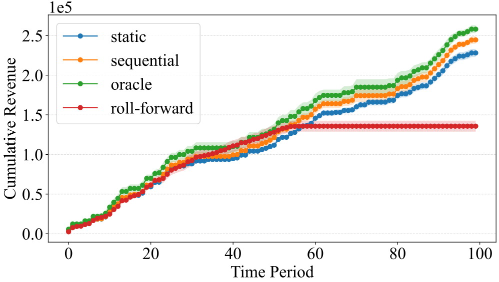

Figure: 20 trials.

Model misspecification: Marginal variances

$$\mathbf{std}_{\text{model}}(d_t) = \mathbf{std}_{\text{real}}(d_t) \times U([0.8, 1.2]), \quad \mathbf{std}_{\text{real}}(d_t) \in [10, 30]$$

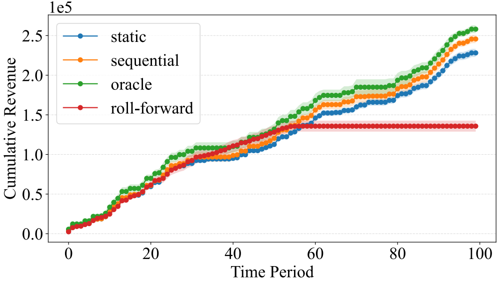

Figure: 20 trials.

Model misspecification: Marginal variances

$$\mathbf{std}_{\text{model}}(d_t) = \mathbf{std}_{\text{real}}(d_t) \times U([0.5, 1.5]), \quad \mathbf{std}_{\text{real}}(d_t) \in [10, 30]$$

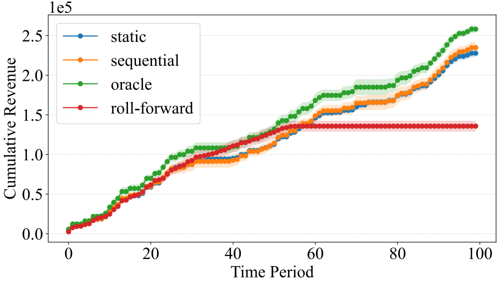

Figure: 20 trials.

Model misspecification: Log demand correlations

$$\mathbf{corr}_{\text{model}}(\log d_\tau, \log d_t) = \mathbf{corr}_{\text{real}}(\log d_\tau, \log d_t) \times U([0.5, 1.5])$$

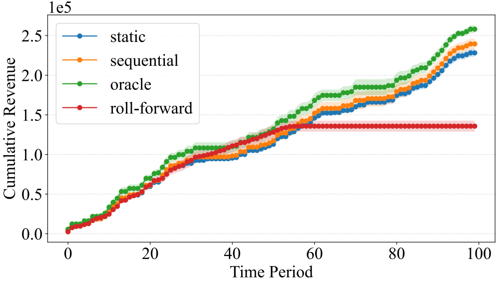

Figure: 20 trials.