

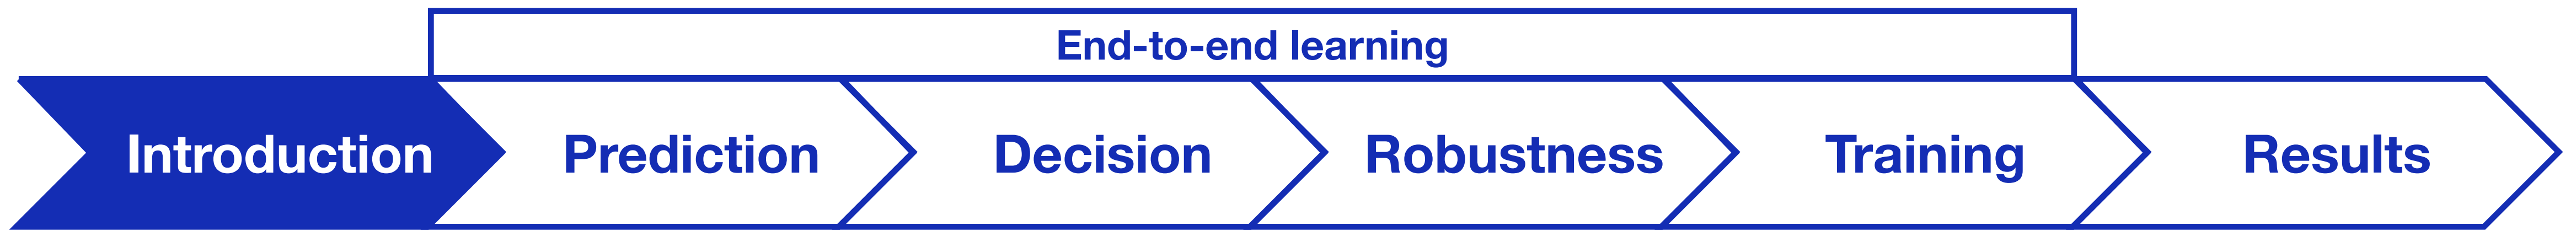
Distributionally Robust End-to-End Portfolio Construction

8th Annual Bloomberg-Columbia
Machine Learning in Finance Workshop

September 2022

Giorgio Costa and **Garud N. Iyengar**





Distributional robustness in end-to-end learning systems




Distributional robustness in end-to-end learning systems

- ▶ **End-to-end:** Integrate the prediction and optimization steps.



Distributional robustness in end-to-end learning systems

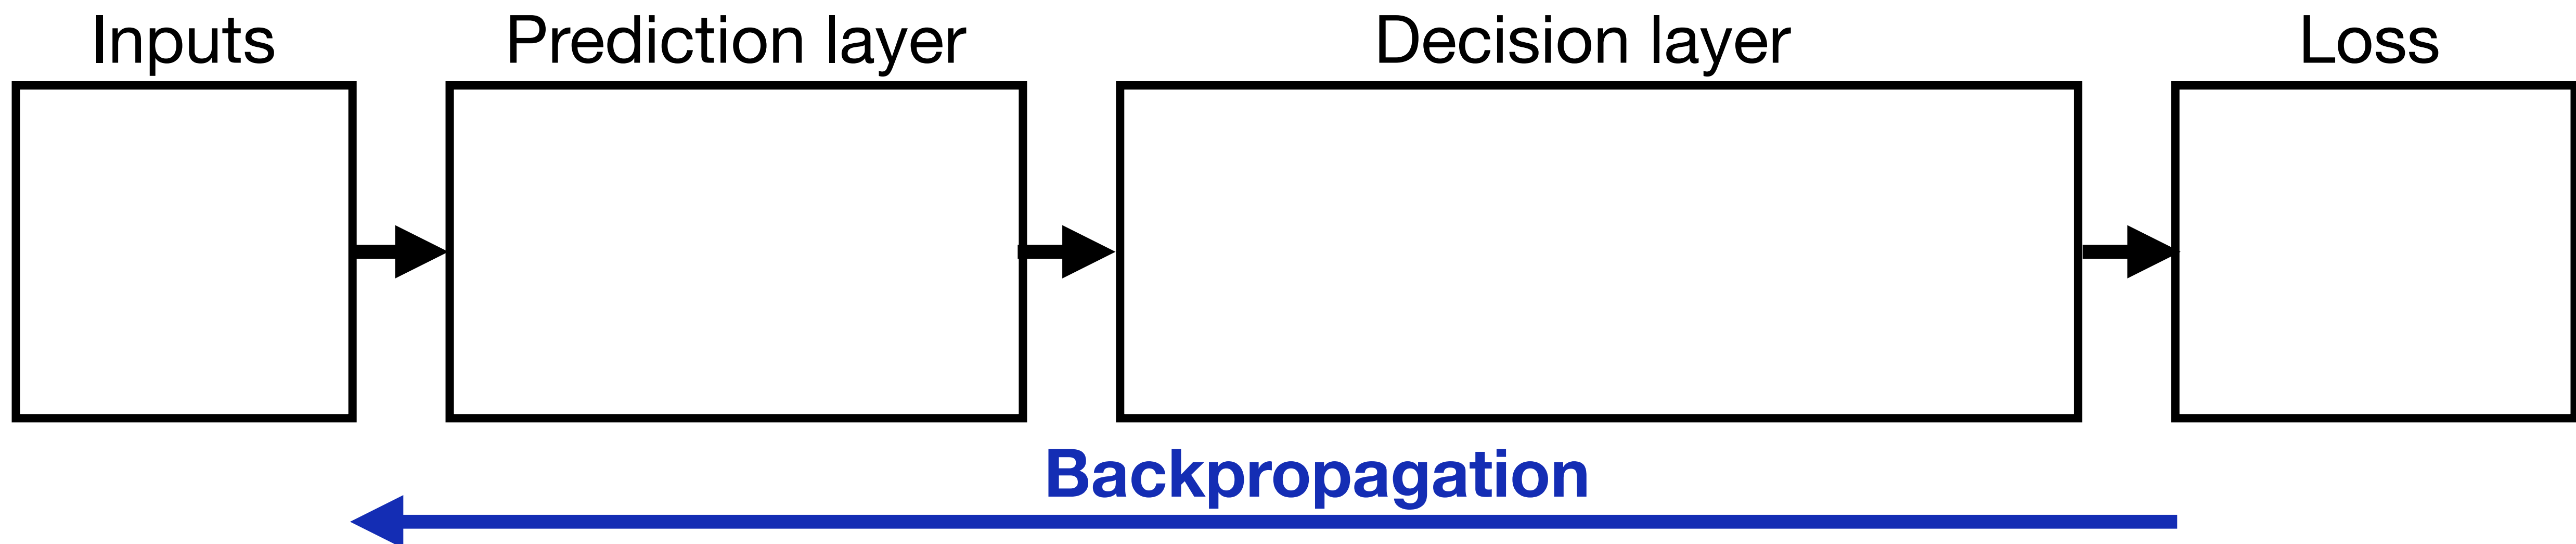
- ▶ **End-to-end:** Integrate the prediction and optimization steps.





Distributional robustness in end-to-end learning systems


- ▶ **End-to-end:** Integrate the prediction and optimization steps.
 - Information is passed between prediction and decision layers during training.







Distributional robustness in end-to-end learning systems

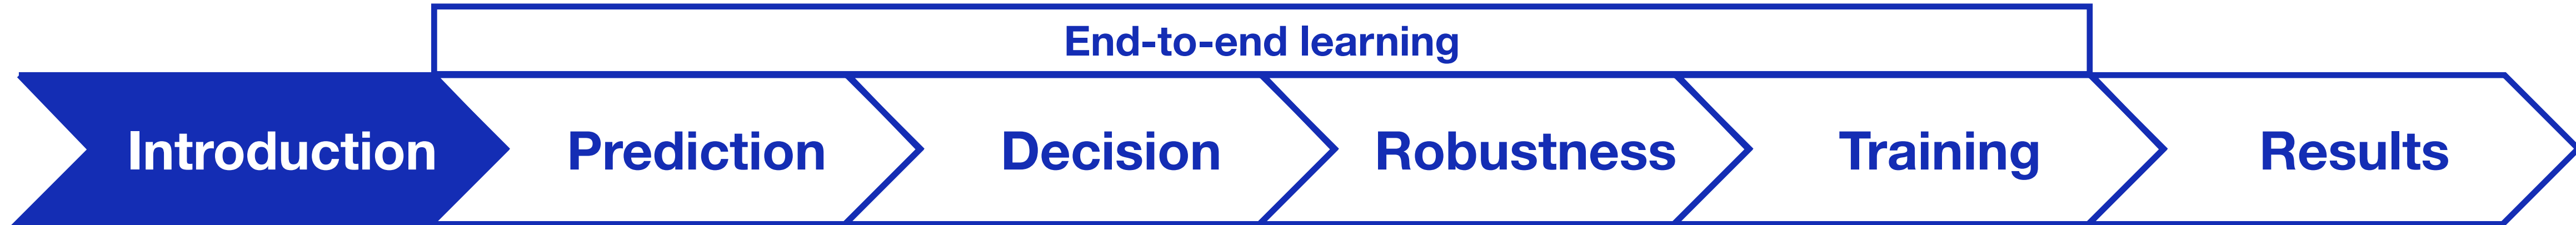
- ▶ **End-to-end:** Integrate the prediction and optimization steps.
 - Information is passed between prediction and decision layers during training.
 - Training is based on the final task rather than predicted performance.






Distributional robustness in end-to-end learning systems





Distributional robustness in end-to-end learning systems


- ▶ The decision layer is a portfolio optimization problem.





Distributional robustness in end-to-end learning systems


- ▶ The decision layer is a portfolio optimization problem.
 - Suffers from sensitivity to prediction errors.






Distributional robustness in end-to-end learning systems


- ▶ The decision layer is a portfolio optimization problem.
 - Suffers from sensitivity to prediction errors.
 - We will use robustness to mitigate model error




End-to-end learning



Data: features and realizations




End-to-end learning




Data: features and realizations

- ▶ $\{\mathbf{x}\}_{j=t-T}^t$: Time series of financial factors
- ▶ $\{\mathbf{y}\}_{j=t-T}^{t+v}$: Time series of asset returns




End-to-end learning




Data: features and realizations

- ▶ $\{\mathbf{x}\}_{j=t-T}^t$: Time series of financial factors
- ▶ $\{\mathbf{y}\}_{j=t-T}^{t+v}$: Time series of asset returns






Prediction layer





End-to-end learning



Prediction layer

► $g_{\theta} : \mathbb{R}^m \rightarrow \mathbb{R}^n$: Prediction model that maps features x_j to predictions \hat{y}_j







Prediction layer

► $g_{\theta} : \mathbb{R}^m \rightarrow \mathbb{R}^n$

: Prediction model that maps features x_j to predictions \hat{y}_j


The prediction model can have any form that allows for **gradient-based learning**






Prediction layer


- ▶ $g_{\theta} : \mathbb{R}^m \rightarrow \mathbb{R}^n$: Prediction model that maps features x_j to predictions \hat{y}_j
- ▶ $\hat{\mathbf{y}} = \{g_{\theta}(\mathbf{x}_j)\}_{j=t-T}^t$: Predicted asset returns






Prediction layer


- ▶ $g_{\theta} : \mathbb{R}^m \rightarrow \mathbb{R}^n$: Prediction model that maps features x_j to predictions \hat{y}_j
- ▶ $\hat{\mathbf{y}} = \{g_{\theta}(\mathbf{x}_j)\}_{j=t-T}^t$: Predicted asset returns
- ▶ $\epsilon = \{\mathbf{y}_j - \hat{\mathbf{y}}_j\}_{j=t-T}^{t-1}$: Prediction errors





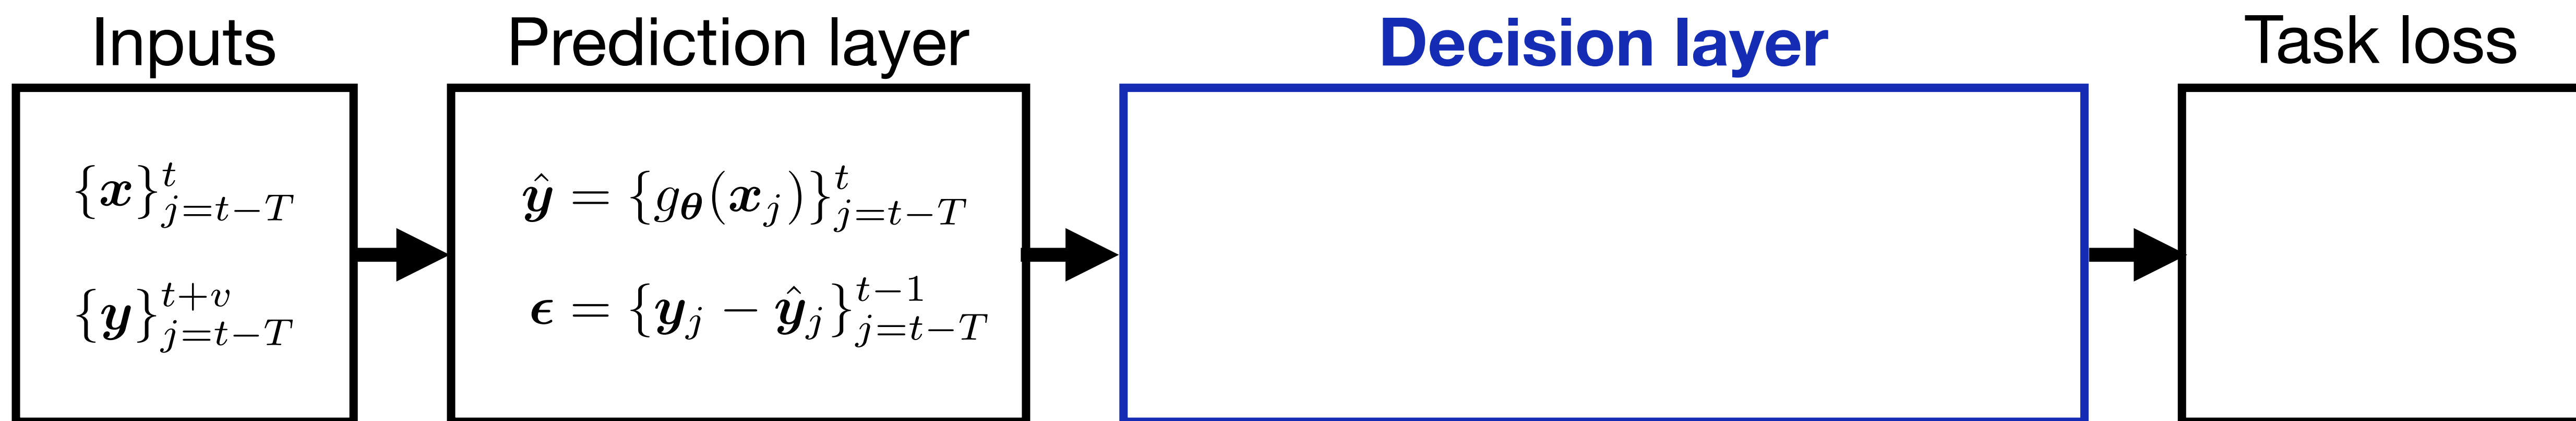
Prediction layer

- ▶ $g_{\theta} : \mathbb{R}^m \rightarrow \mathbb{R}^n$: Prediction model that maps features x_j to predictions \hat{y}_j
- ▶ $\hat{\mathbf{y}} = \{g_{\theta}(\mathbf{x}_j)\}_{j=t-T}^t$: Predicted asset returns
- ▶ $\epsilon = \{\mathbf{y}_j - \hat{\mathbf{y}}_j\}_{j=t-T}^{t-1}$: Prediction errors






Decision layer



End-to-end learning

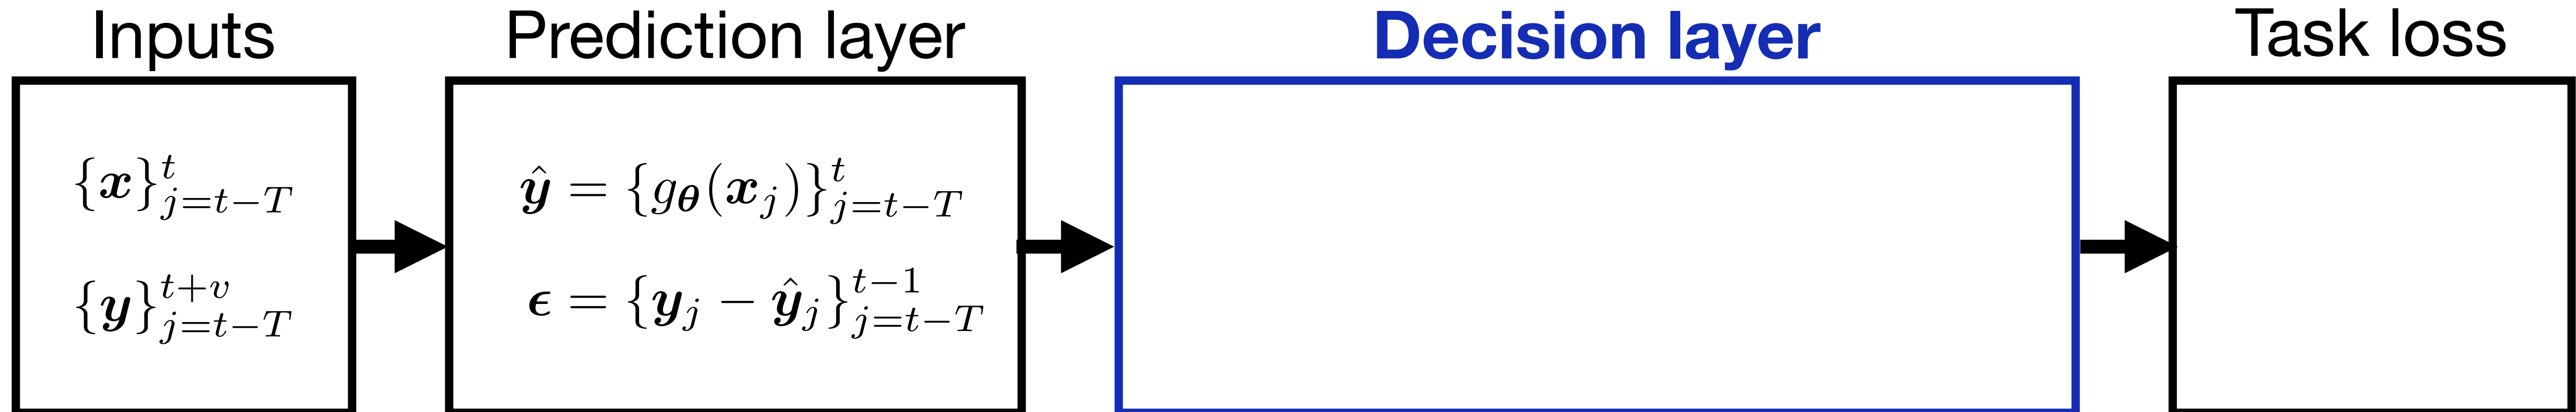


Decision layer


Optimization problem

$$\underbrace{z_t^*}_{\text{Optimal portfolio}} = \operatorname{argmin}_{z \in \mathcal{Z}} f_\epsilon(z) - \gamma \cdot \hat{\mathbf{y}}_t^\top z$$

Optimal portfolio



End-to-end learning




Decision layer


Optimization problem

$$z_t^* = \operatorname{argmin}_{z \in \mathcal{Z}} \underbrace{f_\epsilon(z)}_{\text{Deviation risk measure}} - \gamma \cdot \hat{\mathbf{y}}_t^\top z$$

Deviation risk measure



End-to-end learning

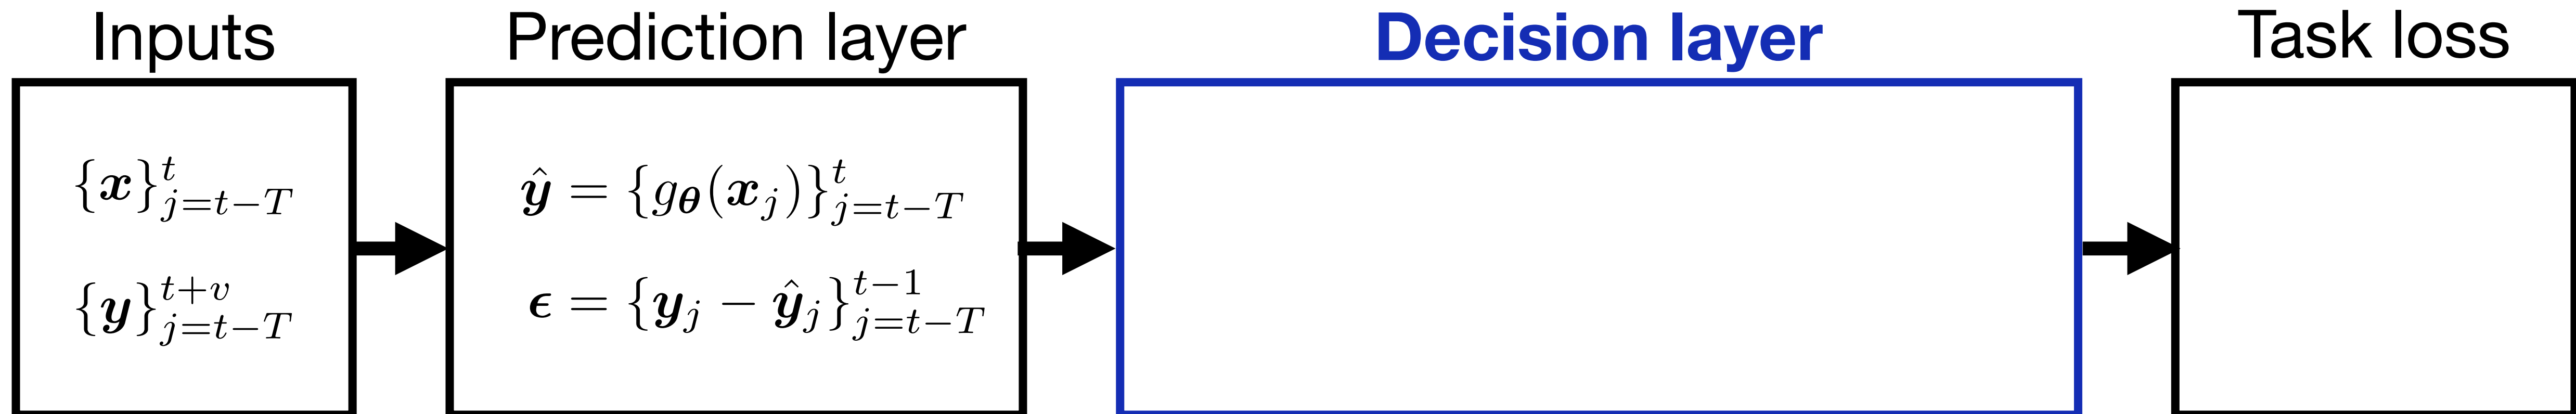


Decision layer


Optimization problem

$$z_t^* = \operatorname{argmin}_{z \in \mathcal{Z}} f_\epsilon(z) - \underbrace{\gamma \cdot \hat{\mathbf{y}}_t^\top z}_{\text{Predicted portfolio return}}$$

Predicted portfolio return



End-to-end learning





Decision layer

Optimization problem

$$z_t^* = \operatorname{argmin}_{z \in \mathcal{Z}} f_\epsilon(z) - \underbrace{\gamma}_{\text{Risk aversion parameter}} \cdot \hat{\mathbf{y}}_t^\top z$$

Risk aversion parameter





Decision layer

Optimization problem


$$z_t^* = \operatorname{argmin}_{z \in \mathcal{Z}} \underbrace{f_\epsilon(z)}_{\text{Deviation risk measure}} - \gamma \cdot \hat{\mathbf{y}}_t^\top z$$


Deviation risk measure

Let $R : \mathbb{R} \rightarrow \mathbb{R}_+ \cup +\infty$ is a closed convex function where $R(0) = 0$ and $R(X) = R(-X)$

Then, $f_\epsilon(z) = \min_c \frac{1}{T} \sum_{j=1}^T R(\epsilon_j^\top z - c)$

is a deviation risk measure





Decision layer

Optimization problem


$$z_t^* = \operatorname{argmin}_{z \in \mathcal{Z}} \underbrace{f_\epsilon(z)}_{\text{Deviation risk measure}} - \gamma \cdot \hat{y}_t^\top z$$


Deviation risk measure

Example: $R(X) = X^2$

$$\text{Then, } f_\epsilon(z) = \min_c \frac{1}{T} \sum_{j=1}^T (\epsilon_j^\top z - c)^2 = \operatorname{var}(\epsilon^\top z)$$

is the portfolio error variance



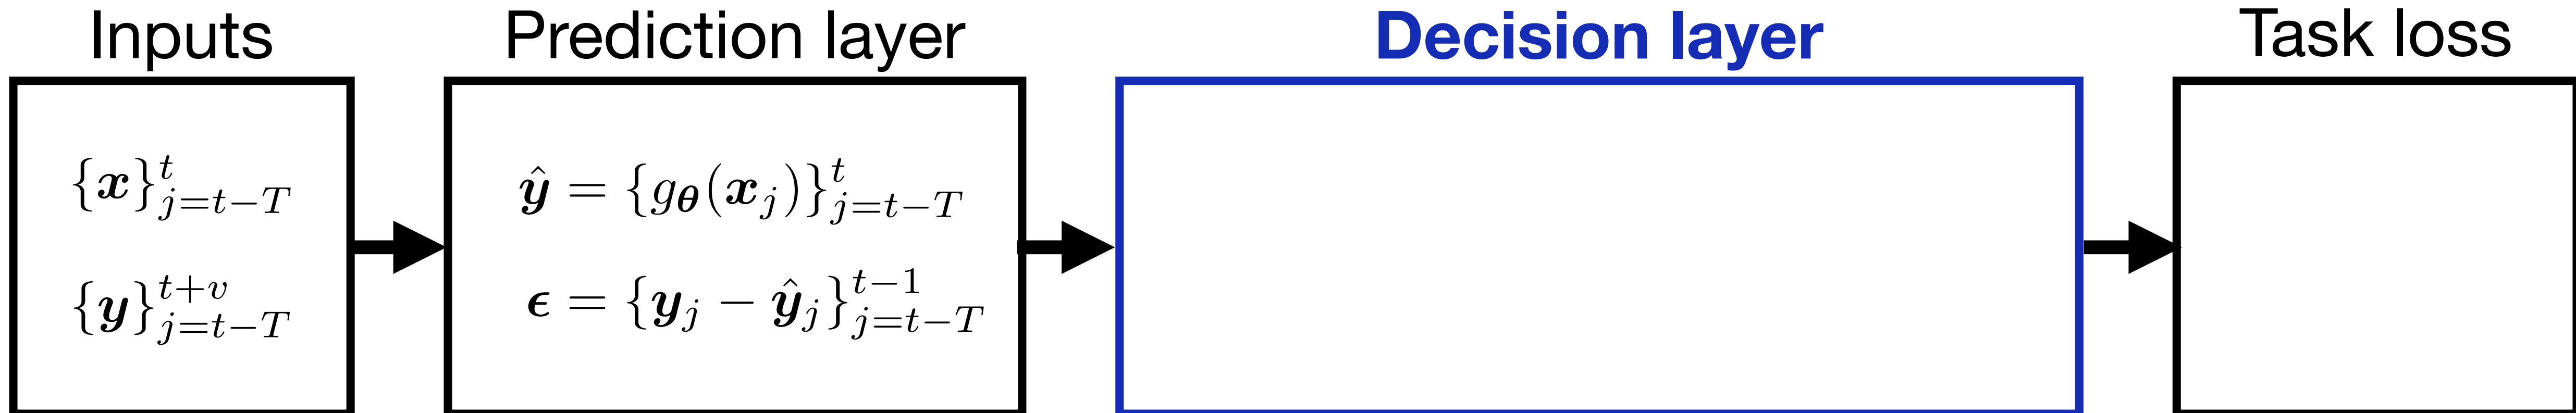



Decision layer

Optimization problem

$$z_t^* = \operatorname{argmin}_{z \in \mathcal{Z}} f_\epsilon(z) - \gamma \cdot \hat{y}_t^\top z$$

- ▶ z_t^* : Optimal portfolio
- ▶ $f_\epsilon(z)$: Convex deviation risk measure
- ▶ $\hat{y}_t^\top z$: Predicted portfolio return
- ▶ γ : Risk aversion parameter







Decision layer

Optimization problem

$$z_t^* = \operatorname{argmin}_{z \in \mathcal{Z}} f_\epsilon(z) - \gamma \cdot \hat{y}_t^\top z$$

- ▶ z_t^* : Optimal portfolio
- ▶ $f_\epsilon(z)$: Convex deviation risk measure
- ▶ $\hat{y}_t^\top z$: Predicted portfolio return
- ▶ γ : Risk appetite - *Learnable parameter*



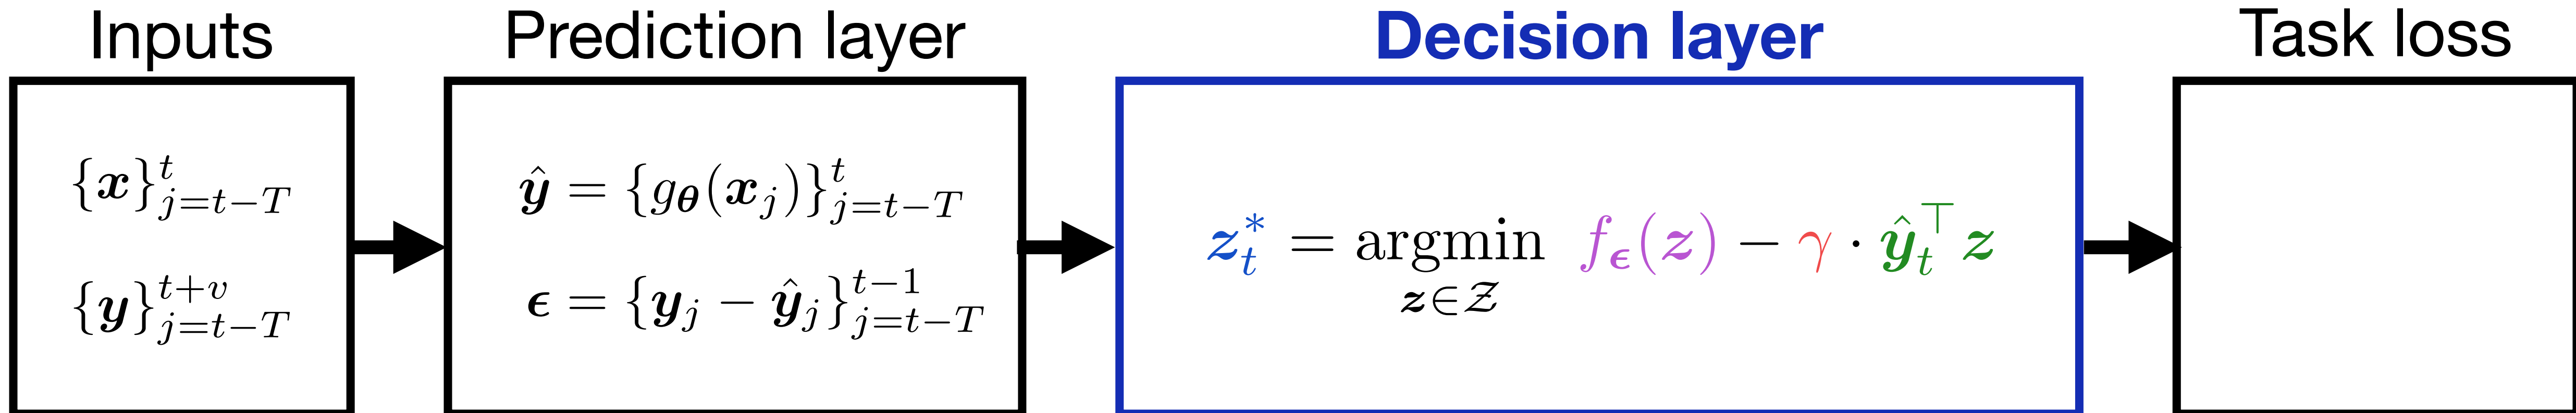


Decision layer

Optimization problem

$$z_t^* = \operatorname{argmin}_{z \in \mathcal{Z}} f_\epsilon(z) - \gamma \cdot \hat{y}_t^\top z$$

- ▶ z_t^* : Optimal portfolio
- ▶ $f_\epsilon(z)$: Convex deviation risk measure
- ▶ $\hat{y}_t^\top z$: Predicted portfolio return
- ▶ γ : Risk appetite - *Learnable parameter*



End-to-end learning

Introduction

Prediction

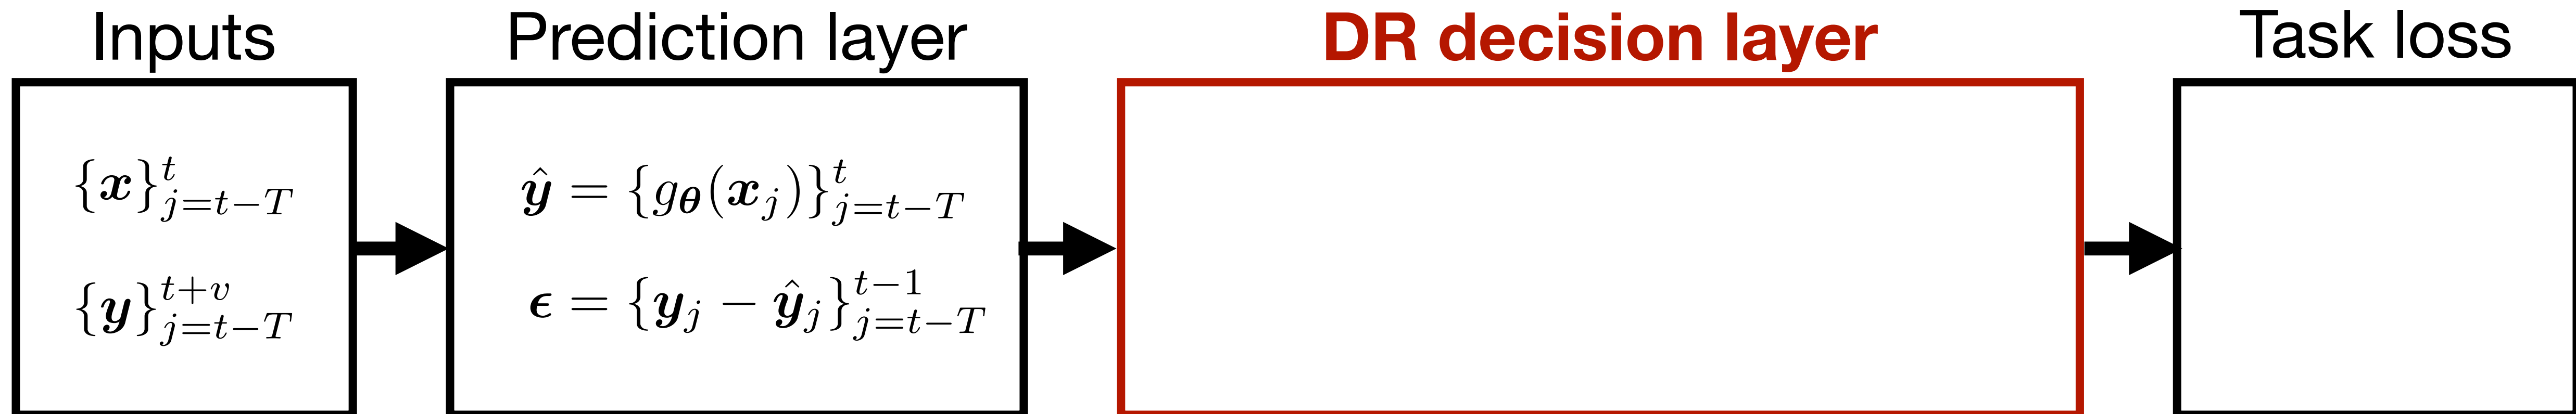
Decision

Robustness

Training

Results

Distributionally robust decision layer



End-to-end learning

Introduction

Prediction

Decision

Robustness

Training

Results

Distributionally robust decision layer

Nominal decision layer

$$\mathbf{z}_t^* = \operatorname{argmin}_{\mathbf{z} \in \tilde{\mathcal{Z}}} f_\epsilon(\mathbf{z}) - \gamma \cdot \hat{\mathbf{y}}_t^\top \mathbf{z}$$

Inputs

$$\{\mathbf{x}\}_{j=t-T}^t$$

$$\{\mathbf{y}\}_{j=t-T}^{t+v}$$


Prediction layer

$$\hat{\mathbf{y}} = \{g_\theta(\mathbf{x}_j)\}_{j=t-T}^t$$

$$\epsilon = \{\mathbf{y}_j - \hat{\mathbf{y}}_j\}_{j=t-T}^{t-1}$$

DR decision layer

Task loss



End-to-end learning

Introduction

Prediction

Decision

Robustness

Training


Results

Distributionally robust decision layer

Nominal decision layer

$$z_t^* = \operatorname{argmin}_{z \in \mathcal{Z}} f_\epsilon(z) - \gamma \cdot \hat{y}_t^\top z$$

Derived from a set of past prediction errors



Distributionally robust decision layer

Nominal decision layer

$$z_t^* = \operatorname{argmin}_{z \in \mathcal{Z}} f_\epsilon(z, \mathbf{q}) - \gamma \cdot \hat{\mathbf{y}}_t^\top z$$

Deviation risk measure

$$f_\epsilon(z, \mathbf{q}) = \min_c \sum_{j=1}^T q_j \cdot R(\boldsymbol{\epsilon}_j^\top z - c)$$

Derived from a set of past prediction errors

- ▶ Nominal assumption: All scenarios are equally likely, $q_j = 1/T$ for $j = 1, \dots, T$



Distributionally robust decision layer

Nominal decision layer

$$z_t^* = \operatorname{argmin}_{z \in \mathcal{Z}} f_\epsilon(z, \mathbf{q}) - \gamma \cdot \hat{\mathbf{y}}_t^\top z$$

- ▶ Nominal assumption: All scenarios are equally likely, $q_j = 1/T$ for $j = 1, \dots, T$
- ▶ Can we protect against scenario probabilities changing in the future?

Introduction

Prediction

Decision

Robustness


Training

Results

Distributionally robust decision layer


Nominal decision layer


$$z_t^* = \operatorname{argmin}_{z \in \mathcal{Z}} f_\epsilon(z, \mathbf{q}) - \gamma \cdot \hat{\mathbf{y}}_t^\top z$$



DR decision layer

$$z_t^* = \operatorname{argmin}_{z \in \mathcal{Z}} \max_{p \in \mathcal{P}(\delta)} f_\epsilon(z, \mathbf{p}) - \gamma \cdot \hat{\mathbf{y}}_t^\top z$$






Distributionally robust decision layer

Nominal decision layer

$$z_t^* = \operatorname{argmin}_{z \in \mathcal{Z}} f_\epsilon(z, \mathbf{q}) - \gamma \cdot \hat{\mathbf{y}}_t^\top z$$




DR decision layer

$$z_t^* = \operatorname{argmin}_{z \in \mathcal{Z}} \max_{\mathbf{p} \in \mathcal{P}(\delta)} f_\epsilon(z, \mathbf{p}) - \gamma \cdot \hat{\mathbf{y}}_t^\top z$$

$$\underbrace{\mathcal{P}(\delta)}_{\text{Ambiguity set}} \triangleq \{ \mathbf{p} \in \mathbb{R}^T : \mathbf{p} \geq \mathbf{0}, \mathbf{1}^\top \mathbf{p} = 1, I_\phi(\mathbf{p}, \mathbf{q}) \leq \delta \}$$

Ambiguity set


- ▶ \mathbf{p} : probability mass function



Distributionally robust decision layer

Nominal decision layer

$$z_t^* = \operatorname{argmin}_{z \in \mathcal{Z}} f_\epsilon(z, \mathbf{q}) - \gamma \cdot \hat{\mathbf{y}}_t^\top z$$




DR decision layer

$$z_t^* = \operatorname{argmin}_{z \in \mathcal{Z}} \max_{\mathbf{p} \in \mathcal{P}(\delta)} f_\epsilon(z, \mathbf{p}) - \gamma \cdot \hat{\mathbf{y}}_t^\top z$$

$$\mathcal{P}(\delta) \triangleq \{ \mathbf{p} \in \mathbb{R}^T : \mathbf{p} \geq \mathbf{0}, \mathbf{1}^\top \mathbf{p} = 1, I_\phi(\mathbf{p}, \mathbf{q}) \leq \delta \}$$

Probability simplex


- ▶ \mathbf{p} : probability mass function



Distributionally robust decision layer

Nominal decision layer

$$z_t^* = \operatorname{argmin}_{z \in \mathcal{Z}} f_\epsilon(z, \mathbf{q}) - \gamma \cdot \hat{\mathbf{y}}_t^\top z$$



DR decision layer

$$z_t^* = \operatorname{argmin}_{z \in \mathcal{Z}} \max_{\mathbf{p} \in \mathcal{P}(\delta)} f_\epsilon(z, \mathbf{p}) - \gamma \cdot \hat{\mathbf{y}}_t^\top z$$

$$\mathcal{P}(\delta) \triangleq \{ \mathbf{p} \in \mathbb{R}^T : \mathbf{p} \geq \mathbf{0}, \mathbf{1}^\top \mathbf{p} = 1, \underbrace{I_\phi(\mathbf{p}, \mathbf{q})}_{\delta\text{-constrained distance measure}} \leq \delta \}$$

δ -constrained distance measure

- ▶ \mathbf{p} : probability mass function
- ▶ Distance measure: ϕ -divergence (e.g., Kullback-Leibler, Hellinger)

End-to-end learning

Introduction

Prediction

Decision

Robustness


Training

Results

Distributionally robust decision layer

Nominal decision layer

$$z_t^* = \operatorname{argmin}_{z \in \mathcal{Z}} f_\epsilon(z, \mathbf{q}) - \gamma \cdot \hat{\mathbf{y}}_t^\top z$$



DR decision layer

$$z_t^* = \operatorname{argmin}_{z \in \mathcal{Z}} \max_{\mathbf{p} \in \mathcal{P}(\delta)} f_\epsilon(z, \mathbf{p}) - \gamma \cdot \hat{\mathbf{y}}_t^\top z$$

Inputs

$$\{\mathbf{x}\}_{j=t-T}^t$$
$$\{\mathbf{y}\}_{j=t-T}^{t+v}$$

Prediction layer

$$\hat{\mathbf{y}} = \{g_\theta(\mathbf{x}_j)\}_{j=t-T}^t$$
$$\epsilon = \{\mathbf{y}_j - \hat{\mathbf{y}}_j\}_{j=t-T}^{t-1}$$

DR decision layer

$$z_t^* = \operatorname{argmin}_{z \in \mathcal{Z}} \max_{\mathbf{p} \in \mathcal{P}(\delta)} f_\epsilon(z, \mathbf{p}) - \gamma \cdot \hat{\mathbf{y}}_t^\top z$$

Task loss

End-to-end learning

Introduction

Prediction

Decision

Robustness

Training

Results

Distributionally robust decision layer

Minimax problem

$$\mathbf{z}_t^* = \operatorname{argmin}_{\mathbf{z} \in \mathcal{Z}} \max_{\mathbf{p} \in \mathcal{P}(\delta)} f_\epsilon(\mathbf{z}, \mathbf{p}) - \gamma \cdot \hat{\mathbf{y}}_t^\top \mathbf{z}$$

Inputs

$$\{\mathbf{x}\}_{j=t-T}^t$$

$$\{\mathbf{y}\}_{j=t-T}^{t+v}$$


Prediction layer

$$\hat{\mathbf{y}} = \{g_\theta(\mathbf{x}_j)\}_{j=t-T}^t$$

$$\epsilon = \{\mathbf{y}_j - \hat{\mathbf{y}}_j\}_{j=t-T}^{t-1}$$

DR decision layer

Task loss



End-to-end learning



Distributionally robust decision layer

Minimax problem

$$z_t^* = \operatorname{argmin}_{z \in \mathcal{Z}} \max_{p \in \mathcal{P}(\delta)} f_\epsilon(z, p) - \gamma \cdot \hat{y}_t^\top z$$


Duality



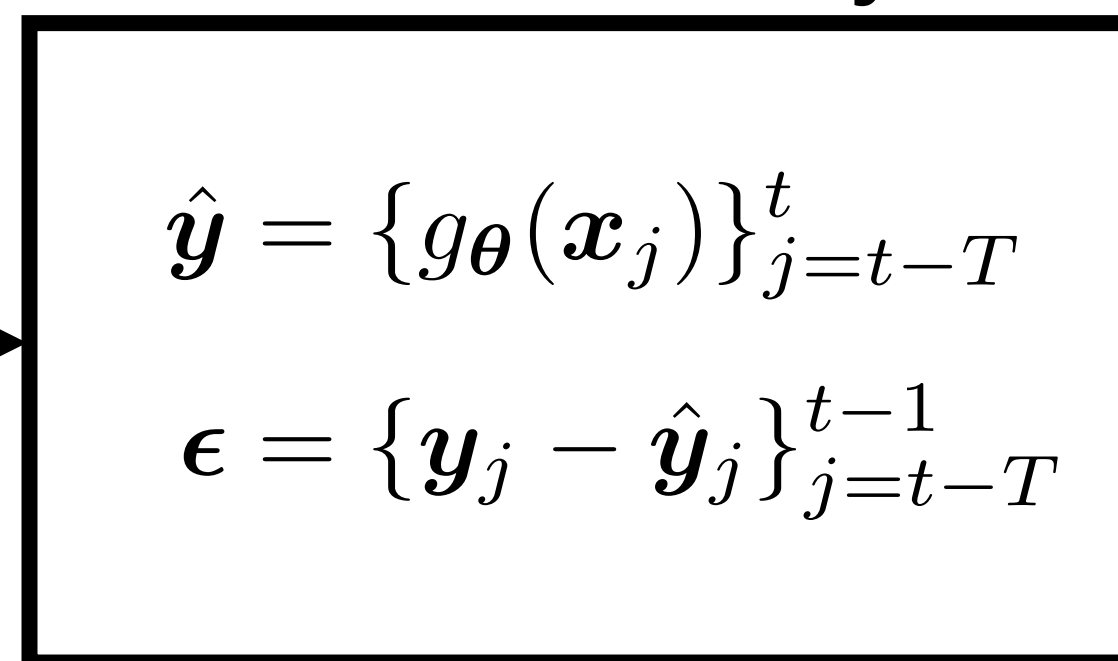
Convex minimization problem

$$z_t^* = \operatorname{argmin}_{z \in \mathcal{Z}, \lambda \geq 0, \xi, c} f_\epsilon^\delta(z, c, \lambda, \xi) - \gamma \cdot \hat{y}_t^\top z$$

Inputs




Prediction layer




DR decision layer



Task loss






Distributionally robust decision layer

Minimax problem

$$z_t^* = \operatorname{argmin}_{z \in \mathcal{Z}} \max_{p \in \mathcal{P}(\delta)} f_\epsilon(z, p) - \gamma \cdot \hat{y}_t^\top z$$

Duality




Convex minimization problem

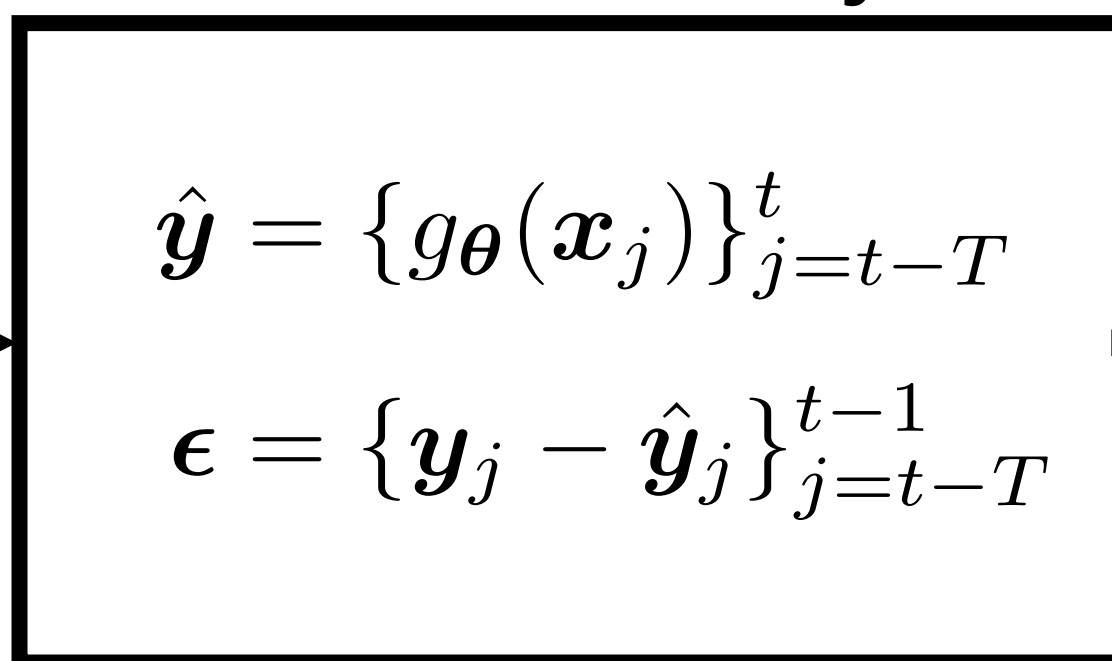
$$z_t^* = \operatorname{argmin}_{z \in \mathcal{Z}, \lambda \geq 0, \xi, c} f_\epsilon^\delta(z, c, \lambda, \xi) - \gamma \cdot \hat{y}_t^\top z$$

δ is now learnable

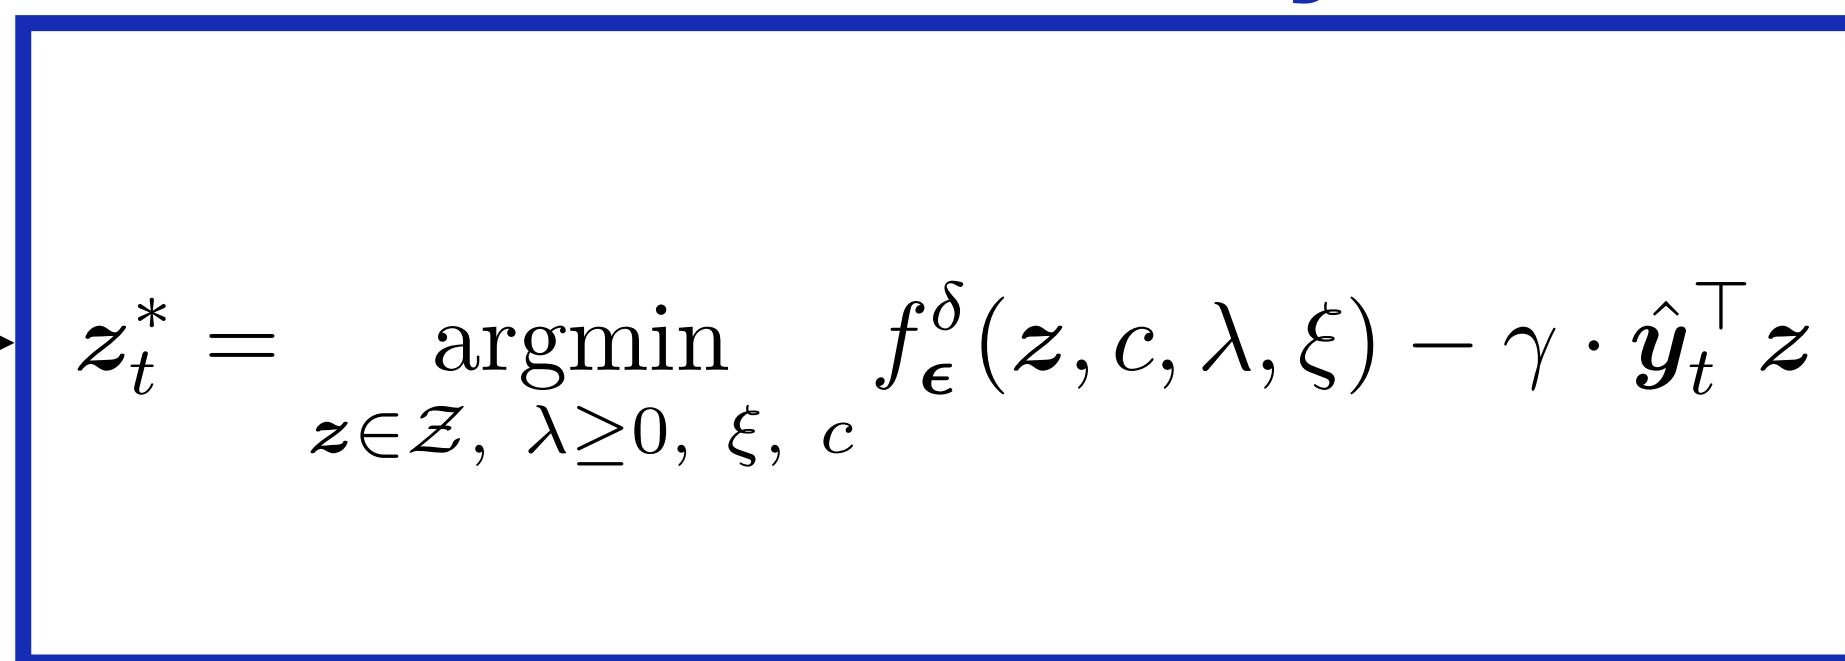
Inputs




Prediction layer



DR decision layer



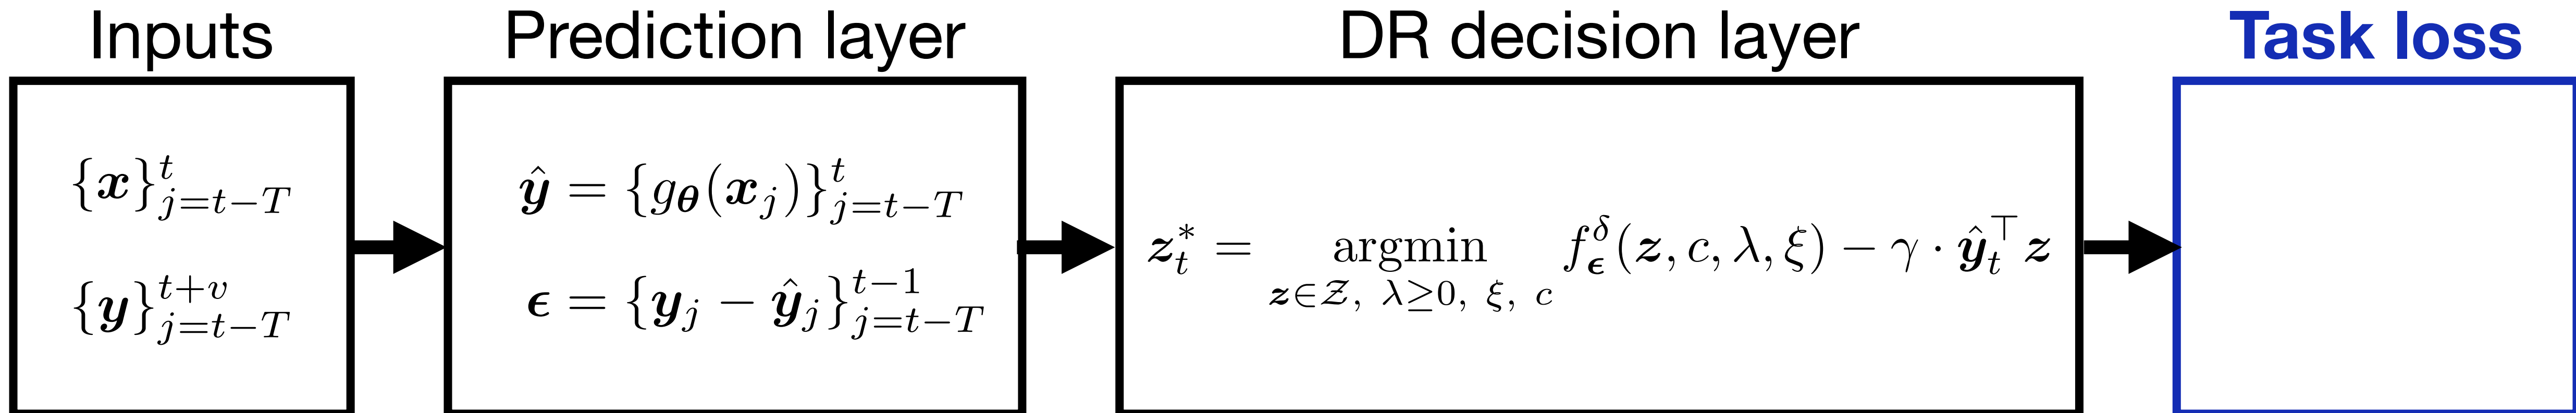
Task loss




End-to-end learning



Task loss

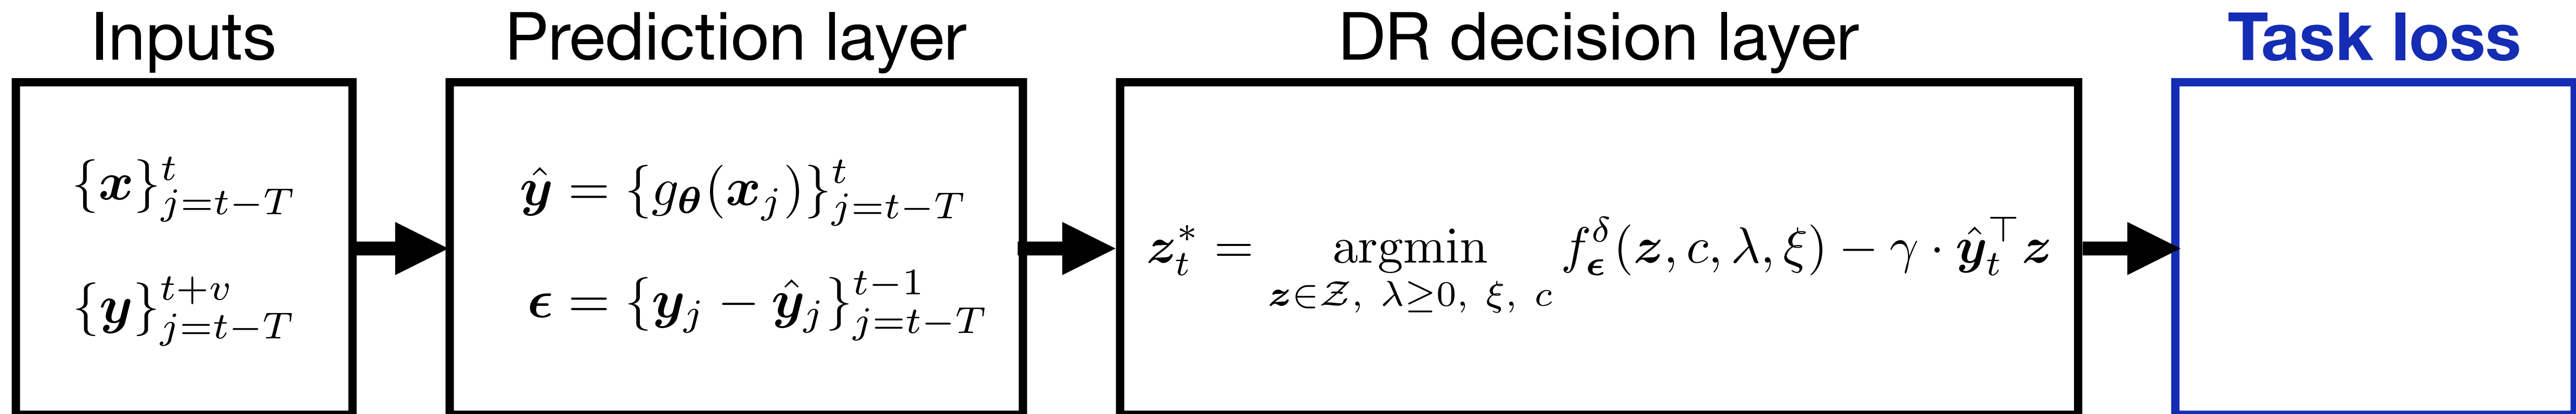


End-to-end learning




Task loss

- ▶ Standard supervised learning: loss function = prediction error.




End-to-end learning




Task loss

- ▶ Standard supervised learning: loss function = prediction error.
- ▶ End-to-end system: Task loss = out-of-sample performance of the decision.




End-to-end learning




Task loss

- ▶ Standard supervised learning: loss function = prediction error.
- ▶ End-to-end system: Task loss = out-of-sample performance of the decision.
- ▶ Task loss function \neq objective function of the decision layer.




End-to-end learning



Task loss

- ▶ Define the task loss as the financial performance over the next v time steps.

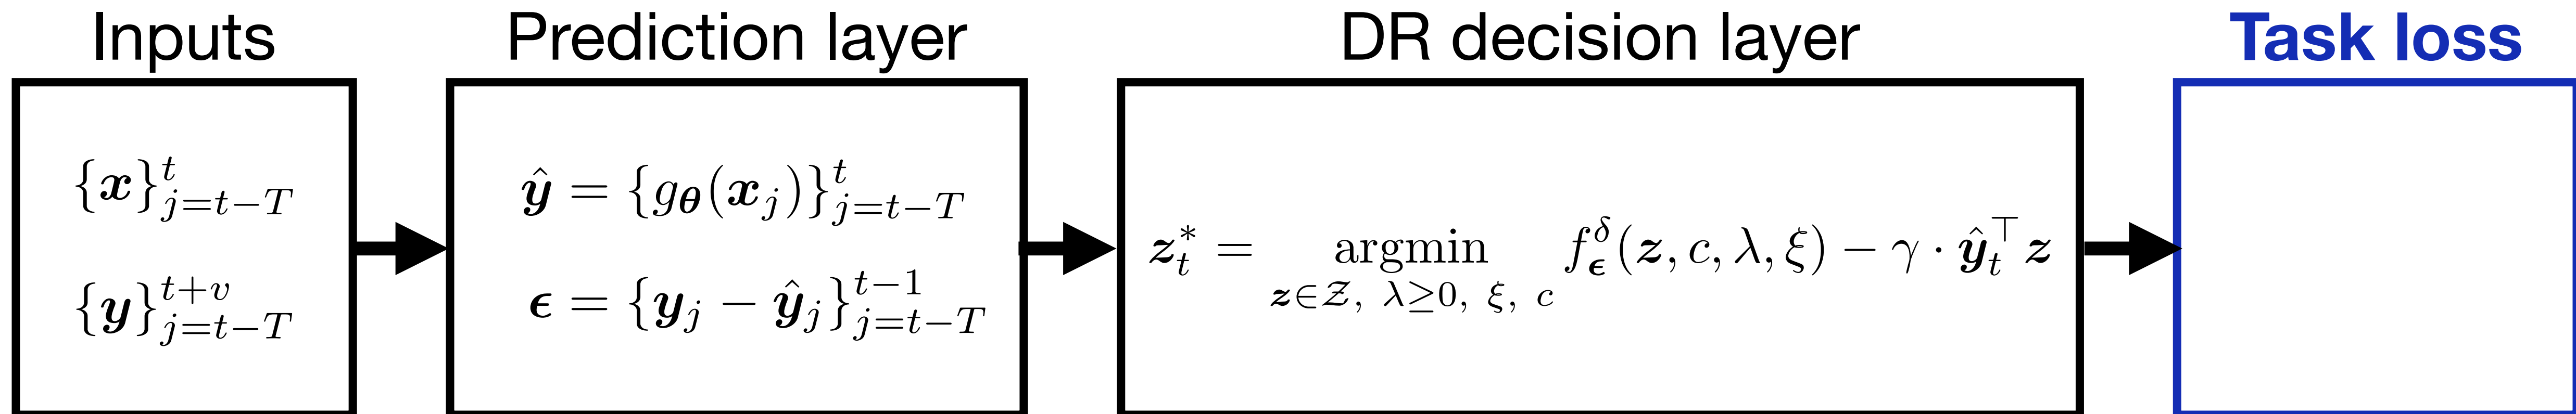





Task loss

- ▶ Define the task loss as the financial performance over the next v time steps.
- ▶ For example, the task loss may be defined as the Sharpe ratio:

$$l(\mathbf{z}_t^*, \{\mathbf{y}_j\}_{j=t}^{t+v}) = -\frac{\text{mean}(\{\mathbf{y}_j^\top \mathbf{z}_t^*\}_{j=t}^{t+v})}{\text{std}(\{\mathbf{y}_j^\top \mathbf{z}_t^*\}_{j=t}^{t+v})}$$

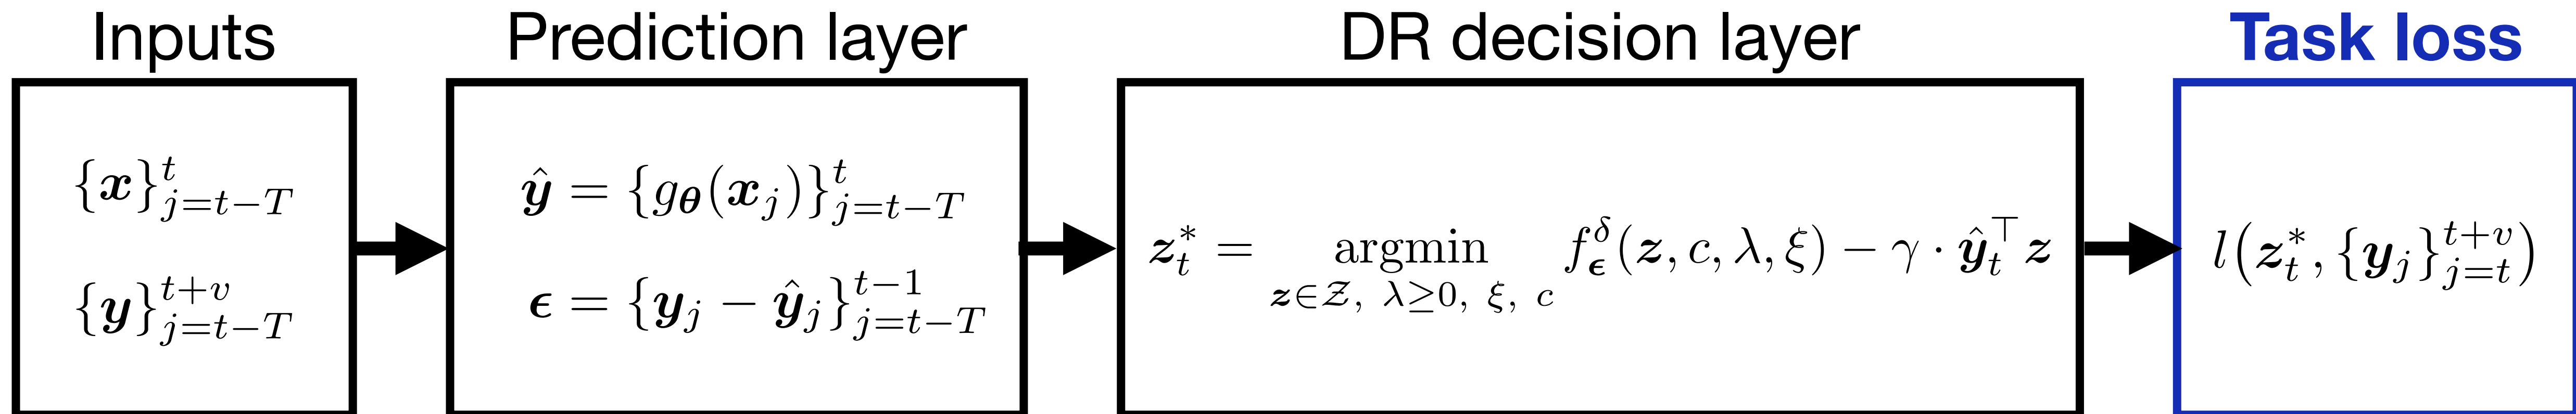




Task loss

- ▶ Define the task loss as the financial performance over the next v time steps.
- ▶ For example, the task loss may be defined as the Sharpe ratio:

$$l(\mathbf{z}_t^*, \{\mathbf{y}_j\}_{j=t}^{t+v}) = -\frac{\text{mean}(\{\mathbf{y}_j^\top \mathbf{z}_t^*\}_{j=t}^{t+v})}{\text{std}(\{\mathbf{y}_j^\top \mathbf{z}_t^*\}_{j=t}^{t+v})}$$



End-to-end learning

Introduction

Prediction

Decision

Robustness

Training

Results

Training

End-to-end learning

Introduction

Prediction

Decision

Robustness


Training

Results

Training

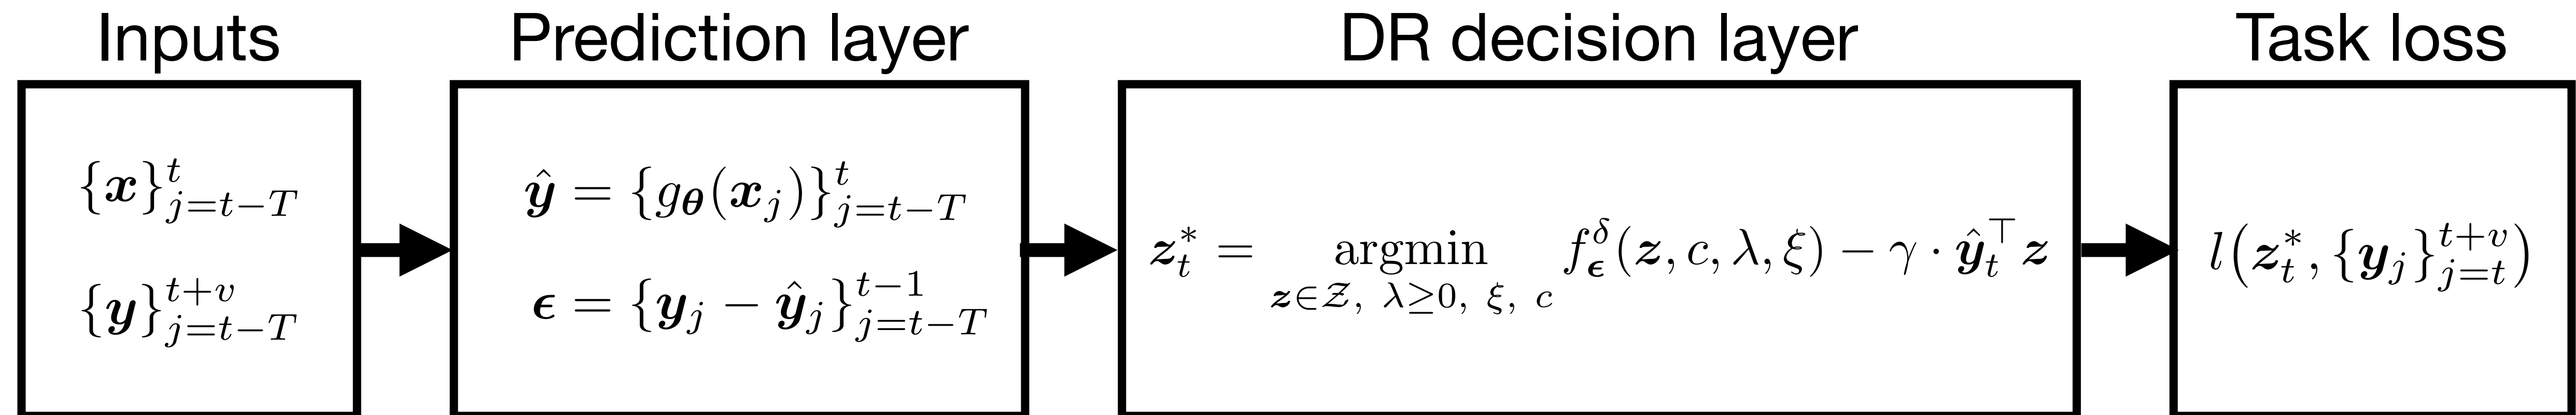
- ▶ The model is trained through gradient descent.

End-to-end learning



Training

Forward Pass:

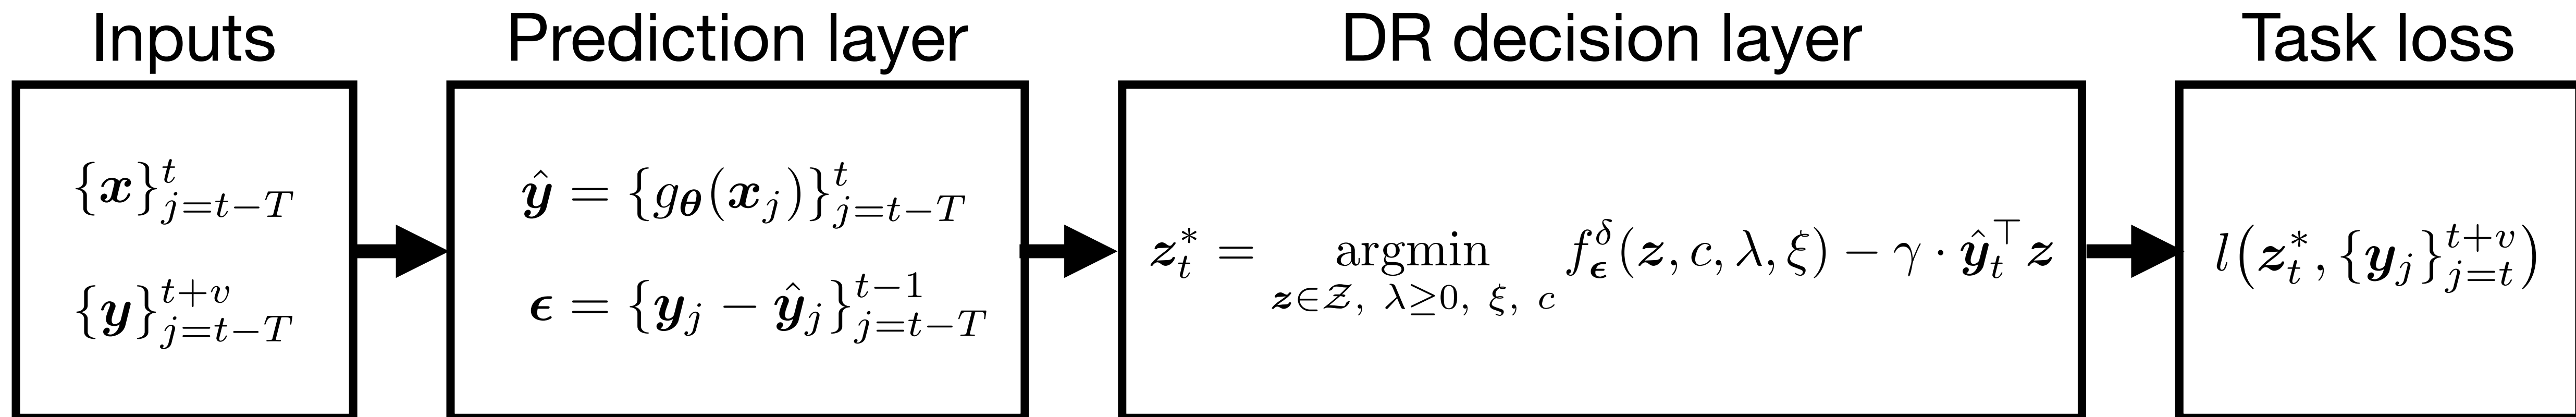


End-to-end learning




Training


Forward Pass:



Backward Pass:





End-to-end learning



Training


- ▶ Backpropagation through an optimization problem is possible by **differentiating** the system of equations arising from the **KKT optimality conditions**.





Training

- ▶ Backpropagation through an optimization problem is possible by differentiating the system of equations arising from the KKT optimality conditions.
- ▶ Both γ and δ are **learned parameters** to enhance out-of-sample performance.



End-to-end learning

Introduction

Prediction

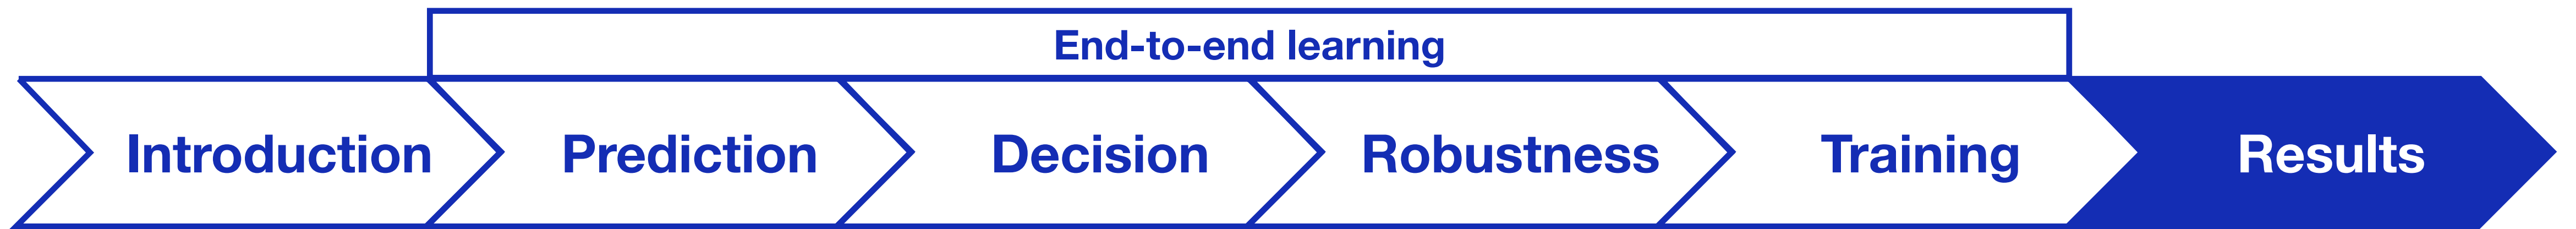
Decision

Robustness

Training

Results

Numerical experiment




Numerical experiment

- ▶ **Data:** Weekly, 07-Jan-2000 to 01-Oct-2021
 - *Assets:* 20 stocks from the S&P500
 - *Features:* 8 Fama-French factors
- ▶ **System:**
 - ▶ *Prediction layer:* Linear
 - ▶ *Task loss:* Sharpe ratio + prediction MSE




Numerical experiment

- ▶ **Data:** Weekly, 07–Jan–2000 to 01-Oct-2021
 - *Assets:* 20 stocks from the S&P500
 - *Features:* 8 Fama-French factors
- ▶ **System:**
 - ▶ *Prediction layer:* Linear
 - ▶ *Task loss:* Sharpe ratio + prediction MSE
- ▶ **Training:** 07–Jan–2000 to 18–Jan–2013
 - Prediction layer initialized to OLS weights
 - For each point prediction, prediction errors computed from $T = 104$ observations
 - The Sharpe ratio is computed over the subsequent $v = 13$ weeks
 - Time series split cross-validation is used to calibrate the learning rate and number of epochs
- ▶ **Testing:** 25–Jan–2013 to 01–Oct–2021
 - Systems are retrained every 2 years



Numerical experiment: Competing models

- ▶ **Equal weight:** $z_t = 1/n$




Numerical experiment: Competing models

► **Equal weight**

► **Predict-then-optimize:** Prediction layer → Fixed OLS weights

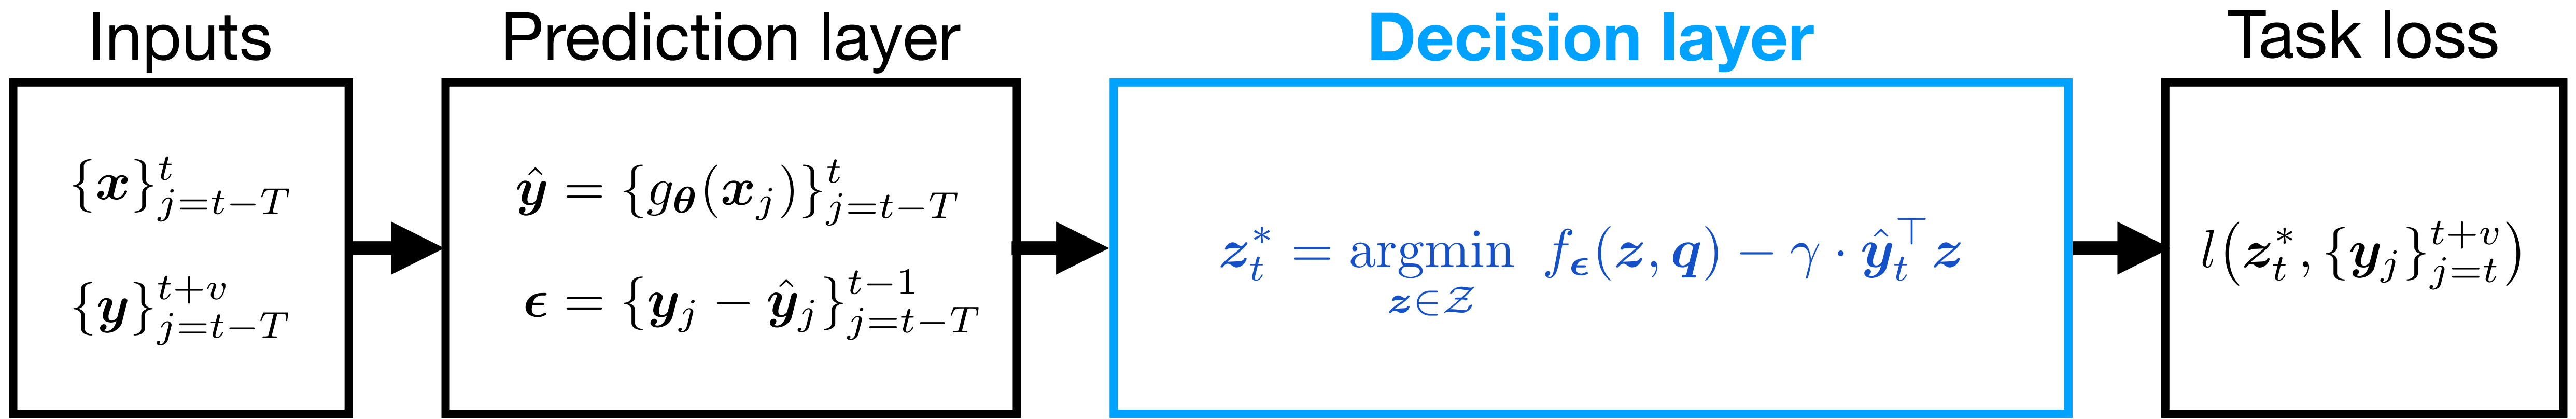
○ Decision layer → $z_t^* = \operatorname{argmin}_{z \in \mathcal{Z}} f_\epsilon(z, q) - \gamma \cdot \hat{y}_t^\top z$

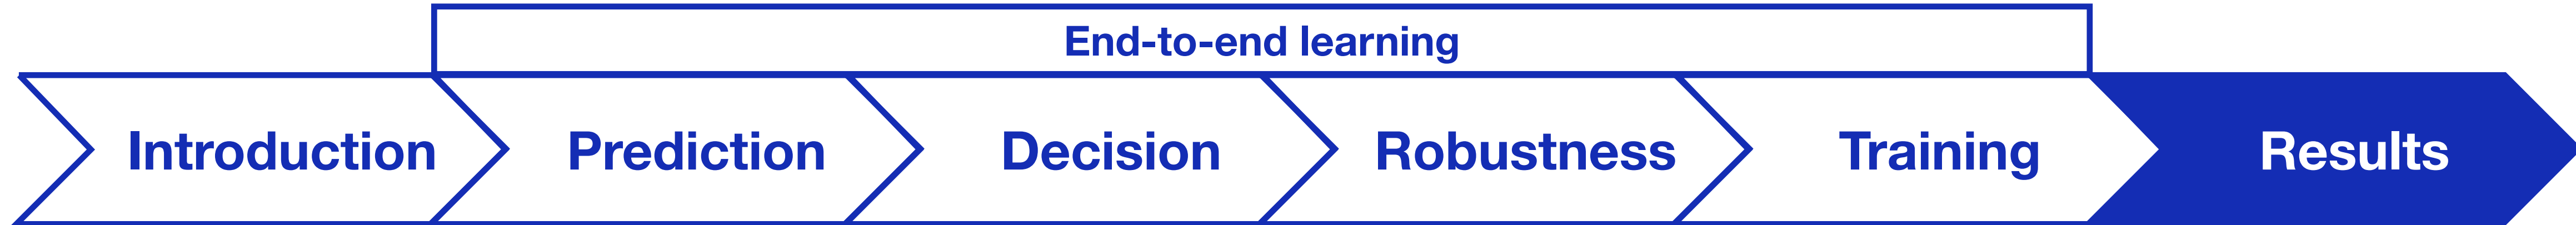
↑ ↑
Constants



Numerical experiment: Competing models

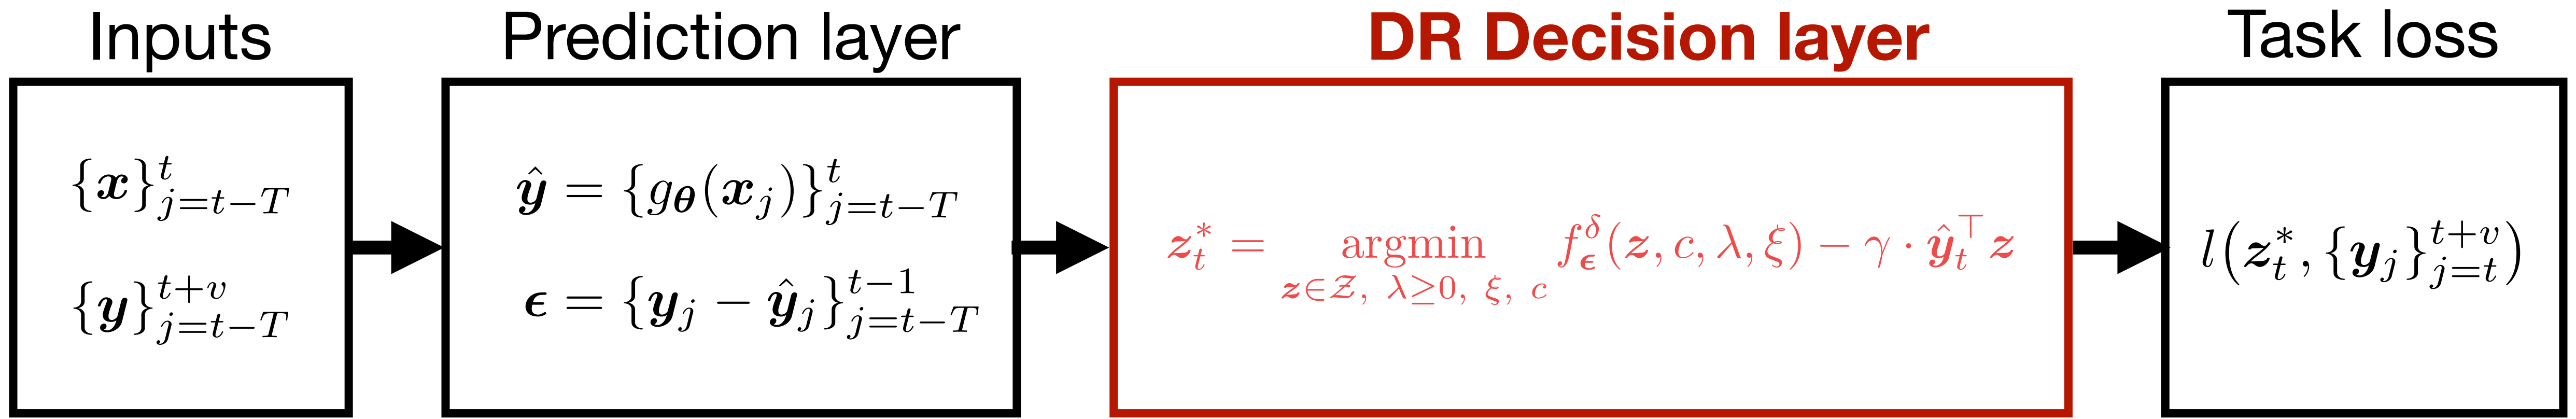
- ▶ Equal weight
- ▶ Predict-then-optimize
- ▶ Nominal:





Numerical experiment: Competing models

- ▶ **Equal weight**
- ▶ **Predict-then-optimize**
- ▶ **Nominal**
- ▶ **DR (Hellinger-based):**



End-to-end learning

Introduction

Prediction


Decision

Robustness

Training

Results

Out-of-sample financial performance



Sharpe ratio

Equal weight: 1.05

Pred-then-Opt: 0.88

Nominal: 1.24

DR: 1.30

End-to-end learning

Introduction

Prediction

Decision

Robustness

Training

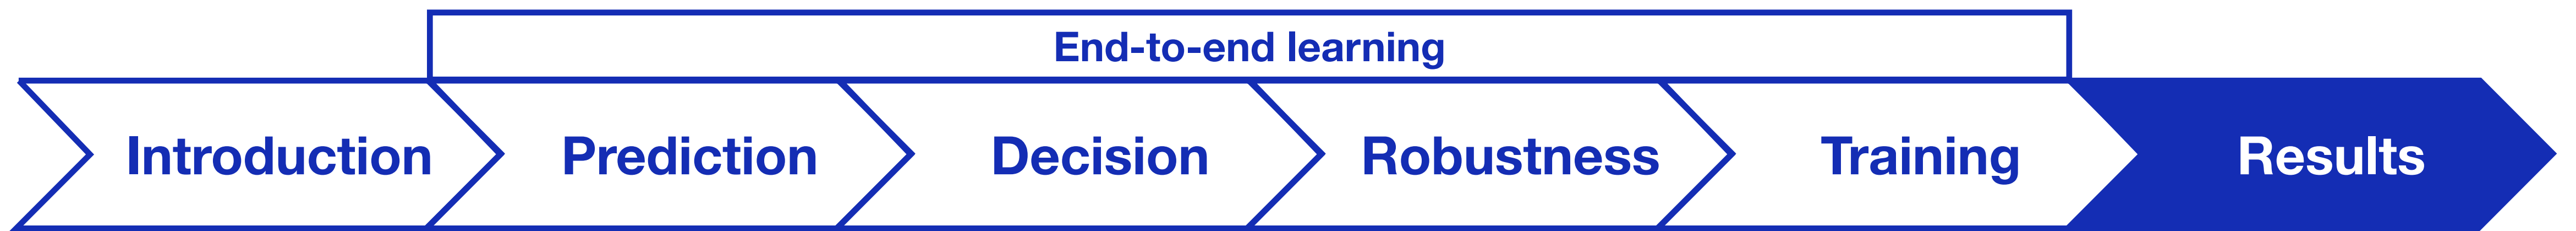
Results

Conclusion




Conclusion

- ▶ End-to-end system with a **robust** decision layer that explicitly incorporates prediction model risk.




Conclusion

- ▶ End-to-end system with a **robust** decision layer that explicitly incorporates prediction model risk.
 - By design, we pass both the prediction and a set of prediction errors to the decision layer.
 - Furthermore, we introduce robustness by taking the worst-case risk over a set of probability measures.



Conclusion

- ▶ End-to-end system with a **robust** decision layer that explicitly incorporates prediction model risk.
 - By design, pass both the prediction and a set of prediction errors to the decision layer.
 - Furthermore, introduce robustness by taking the worst-case risk over a set of probability measures.
- ▶ Use convex duality to show that the DR decision layer is computationally tractable.



Conclusion

- ▶ End-to-end system with a **robust** decision layer that explicitly incorporates prediction model risk.
 - By design, pass both the prediction and a set of prediction errors to the decision layer.
 - Furthermore, introduce robustness by taking the worst-case risk over a set of probability measures.
- ▶ Use convex duality to show that the DR decision layer is computationally tractable.
- ▶ Risk and robustness parameters are *learned* directly from data.



Conclusion

- ▶ End-to-end system with a **robust** decision layer that explicitly incorporates prediction model risk.
 - By design, we pass both the prediction and a set of prediction errors to the decision layer.
 - Furthermore, we introduce robustness by taking the worst-case risk over a set of probability measures.
- ▶ Use convex duality to show that the DR decision layer is computationally tractable.
- ▶ Risk and robustness parameters are **learned** directly from data.
 - Being able to learn these parameters relieves practitioners from the challenge of determining them a priori.

References

- ▶ Amos, B. and Kolter, J. Z. Optnet: Differentiable optimization as a layer in neural networks. In *International Conference on Machine Learning*, pp. 136–145. PMLR, 2017.
- ▶ Ben-Tal, A., Den Hertog, D., De Waegenaere, A., Melenberg, B., and Rennen, G. Robust solutions of optimization problems affected by uncertain probabilities. *Management Science*, 59(2):341–357, 2013.
- ▶ Calafiore, G. C. Ambiguous risk measures and optimal robust portfolios. *SIAM Journal on Optimization*, 18(3): 853–877, 2007.
- ▶ Donti, P. L., Amos, B., and Kolter, J. Z. Task-based end-to-end model learning in stochastic optimization. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pp. 5490–5500, 2017.

A draft version our paper is now available

- ▶ Costa, G., & Iyengar, G. N. (2022). Distributionally Robust End-to-End Portfolio Construction. arXiv preprint arXiv:2206.05134.

Thank you!