Robust Gittins for Stochastic Scheduling

Heather Newman (Carnegie Mellon)

SIGMETRICS 2025

Joint work with: Ben Moseley (Carnegie Mellon), Kirk Pruhs (U. of Pittsburgh), and Rudy Zhou (Microsoft)

Motivation

Motivation

Many stochastic optimization policies assume perfectly accurate distributions

Motivation

Many stochastic optimization policies assume perfectly accurate distributions

↓

rich information

Motivation

Many stochastic optimization policies assume perfectly accurate distributions

↓

rich information

↓

brittle algorithms

Motivation

unrealistic

Many stochastic optimization policies assume perfectly ~~accurate distributions~~

↓

rich information

↓

brittle algorithms

Motivation

unrealistic

Many stochastic optimization policies assume ~~perfectly accurate distributions~~

↓

rich information

↓

brittle algorithms

Motivation: develop stochastic optimization algorithms that are **robust** to **imperfect predicted distributions**

Motivation

unrealistic

Many stochastic optimization policies assume perfectly ~~accurate distributions~~

↓

rich information

↓

brittle algorithms

Motivation: develop stochastic optimization algorithms that are **robust** to **imperfect predicted distributions**

Today:

preemptively schedule stochastic jobs on single machine
to minimize total completion time (no release dates)

Related Work

Related Work

M/G/1 queue

Related Work

M/G/1 queue

- Scully, Grosof, & Mitzenmacher '22: give scheduler **stochastic estimate** z_j of true size s_j where $(s_j, z_j) \sim (S, Z)$ and $z_j \in [\beta \cdot s_j, \alpha \cdot s_j]$ → compare against SRPT

Related Work

M/G/1 queue

- Scully, Grosof, & Mitzenmacher '22: give scheduler **stochastic estimate** z_j of true size s_j where $(s_j, z_j) \sim (S, Z)$ and $z_j \in [\beta \cdot s_j, \alpha \cdot s_j] \rightarrow$ compare against SRPT
- Scully & Harchol-Balter '18: consider error on job **ages**

Related Work

M/G/1 queue

- Scully, Grosof, & Mitzenmacher '22: give scheduler **stochastic estimate** z_j of true size s_j where $(s_j, z_j) \sim (S, Z)$ and $z_j \in [\beta \cdot s_j, \alpha \cdot s_j]$ → compare against SRPT
- Scully & Harchol-Balter '18: consider error on job **ages**

Scheduling with predictions (non-stochastic setting)

Related Work

M/G/1 queue

- Scully, Grosof, & Mitzenmacher '22: give scheduler **stochastic estimate** z_j of true size s_j where $(s_j, z_j) \sim (S, Z)$ and $z_j \in [\beta \cdot s_j, \alpha \cdot s_j]$ → compare against SRPT
- Scully & Harchol-Balter '18: consider error on job **ages**

Scheduling with predictions (non-stochastic setting)

- Purohit et. al. '18; Azar et. al '21, '22; Im et. al. '23

Related Work

M/G/1 queue

- Scully, Grosof, & Mitzenmacher ’22: give scheduler **stochastic estimate** z_j of true size s_j where $(s_j, z_j) \sim (S, Z)$ and $z_j \in [\beta \cdot s_j, \alpha \cdot s_j] \rightarrow$ compare against SRPT
- Scully & Harchol-Balter ’18: consider error on job **ages**

Scheduling with predictions (non-stochastic setting)

- Purohit et. al. ’18; Azar et. al ’21, ’22; Im et. al. ’23

Non-scheduling problems

Related Work

M/G/1 queue

- Scully, Grosof, & Mitzenmacher '22: give scheduler **stochastic estimate** z_j of true size s_j where $(s_j, z_j) \sim (S, Z)$ and $z_j \in [\beta \cdot s_j, \alpha \cdot s_j] \rightarrow$ compare against SRPT
- Scully & Harchol-Balter '18: consider error on job **ages**

Scheduling with predictions (non-stochastic setting)

- Purohit et. al. '18; Azar et. al '21, '22; Im et. al. '23

Non-scheduling problems

- Dütting & Kesselheim '19: Prophet inequalities

Related Work

M/G/1 queue

- Scully, Grosof, & Mitzenmacher '22: give scheduler **stochastic estimate** z_j of true size s_j where $(s_j, z_j) \sim (S, Z)$ and $z_j \in [\beta \cdot s_j, \alpha \cdot s_j]$ → compare against SRPT
- Scully & Harchol-Balter '18: consider error on job **ages**

Scheduling with predictions (non-stochastic setting)

- Purohit et. al. '18; Azar et. al '21, '22; Im et. al. '23

Non-scheduling problems

- Dütting & Kesselheim '19: Prophet inequalities
- Banishashem et. al. '25: Pandora's box

Related Work

M/G/1 queue

- Scully, Grosof, & Mitzenmacher '22: give scheduler **stochastic estimate** z_j of true size s_j where $(s_j, z_j) \sim (S, Z)$ and $z_j \in [\beta \cdot s_j, \alpha \cdot s_j] \rightarrow$ compare against SRPT
- Scully & Harchol-Balter '18: consider error on job **ages**

Scheduling with predictions (non-stochastic setting)

- Purohit et. al. '18; Azar et. al '21, '22; Im et. al. '23

Non-scheduling problems

- Dütting & Kesselheim '19: Prophet inequalities
- Banishashem et. al. '25: Pandora's box
- Kim & Lim '16: Multi-armed bandits

Problem Definition

Problem Definition

$\mathscr{D}_1$

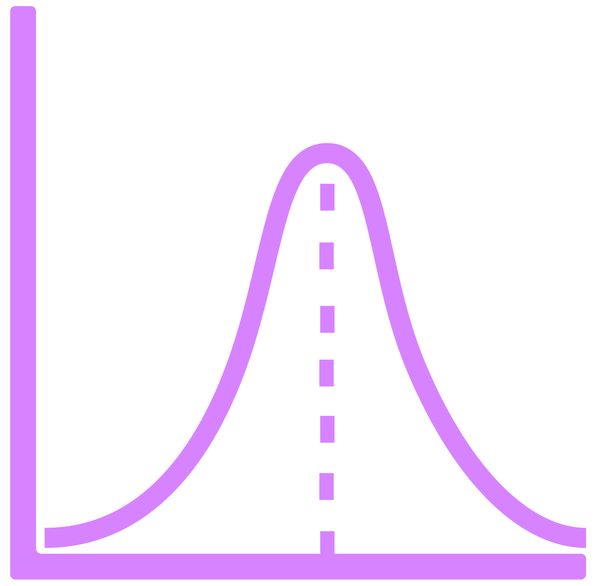

$\mathscr{D}_2$

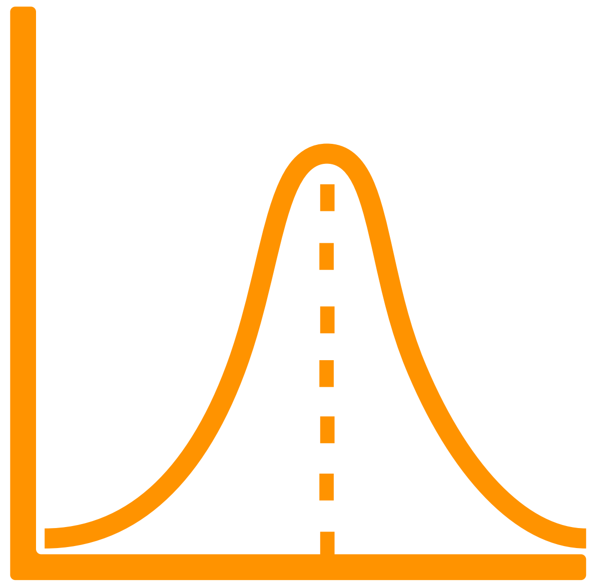

$\mathscr{D}_3$

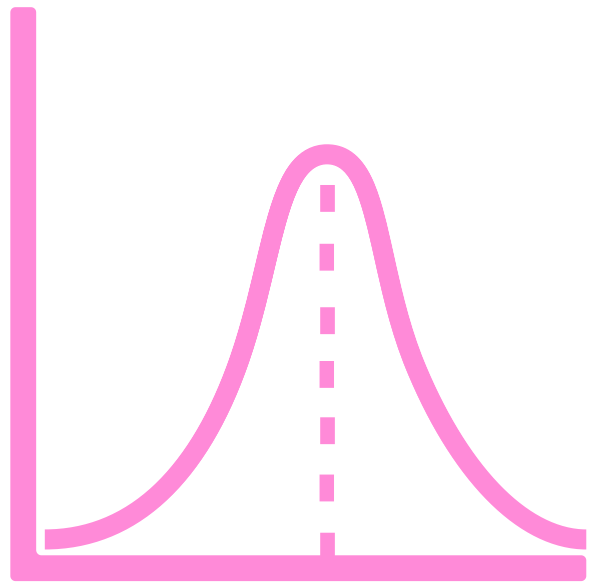

Problem Definition

independent
(not necessarily
identical)

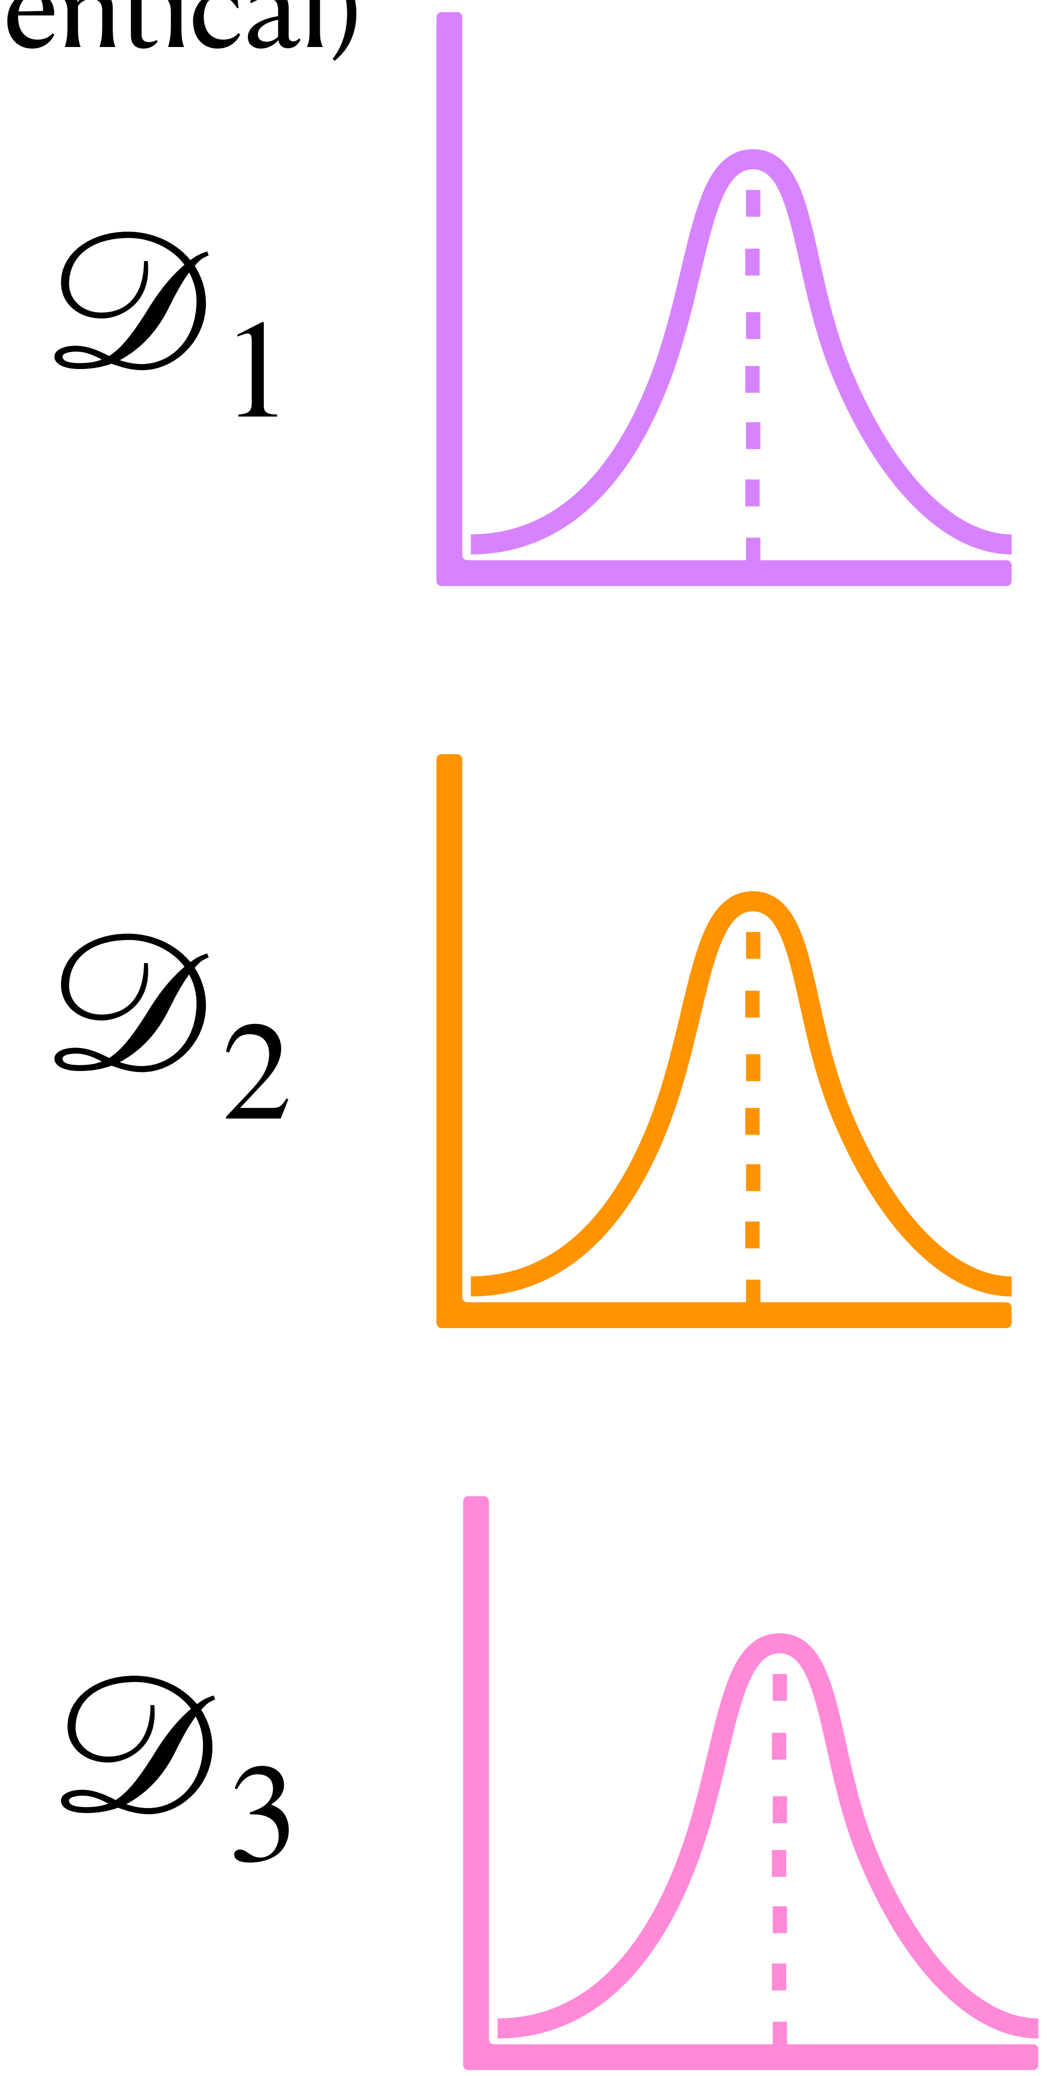

Problem Definition

independent
(not necessarily
identical)

$\mathscr{D}_1$

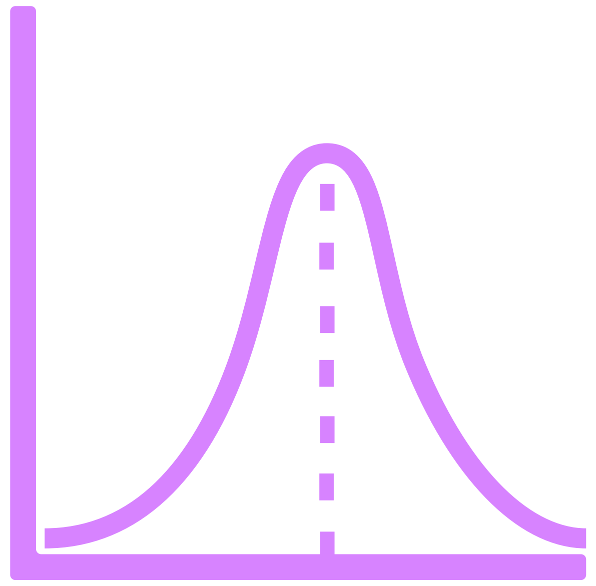

P_1

$\mathscr{D}_2$

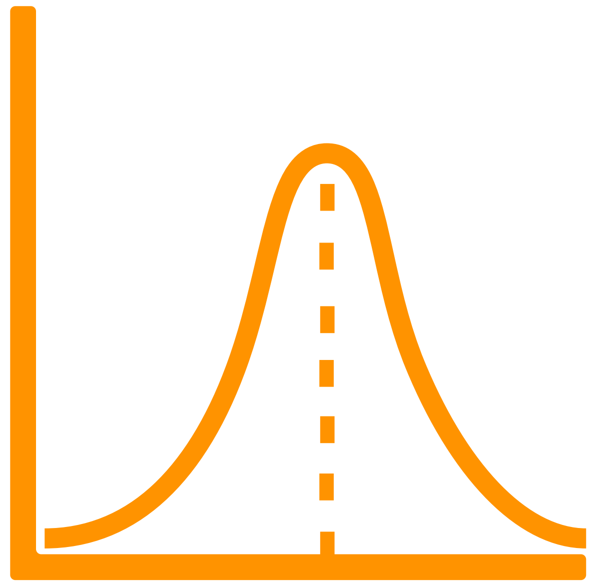

P_2

$\mathscr{D}_3$

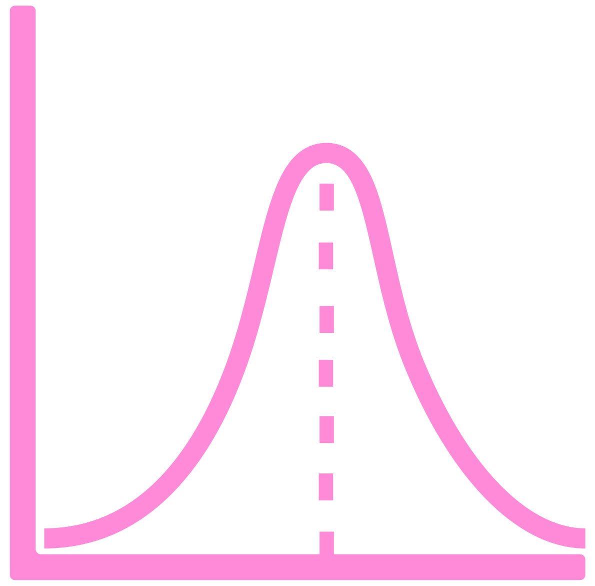

P_3

Problem Definition

independent
(not necessarily
identical)

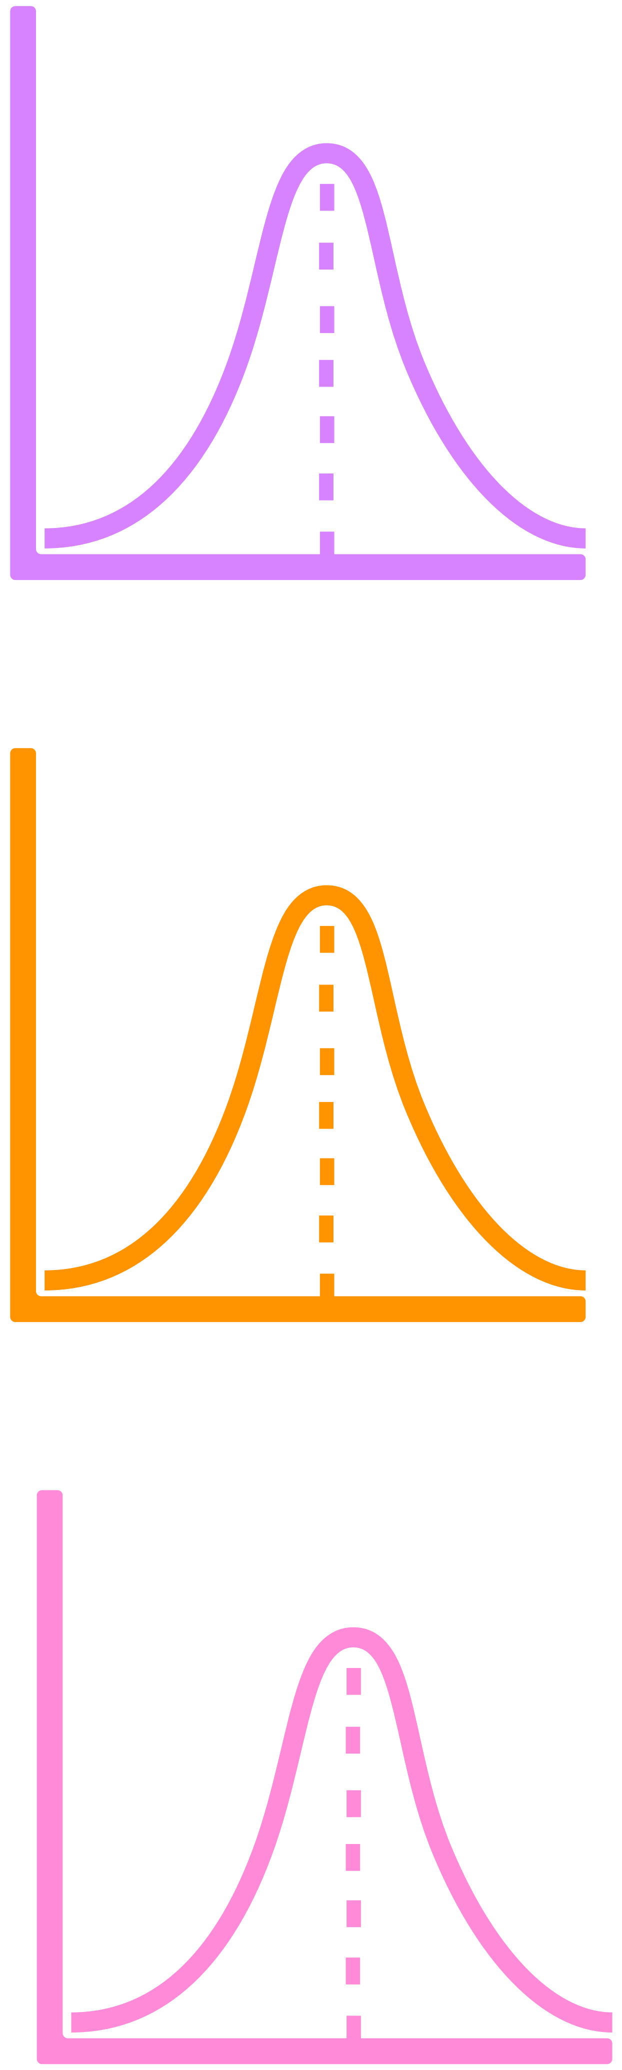

$\mathcal{D}_1$ P_1

$\mathcal{D}_2$ P_2

$\mathcal{D}_3$ P_3

- Nonanticipatory

Problem Definition

independent
(not necessarily
identical)

- Nonanticipatory
- Preemption allowed

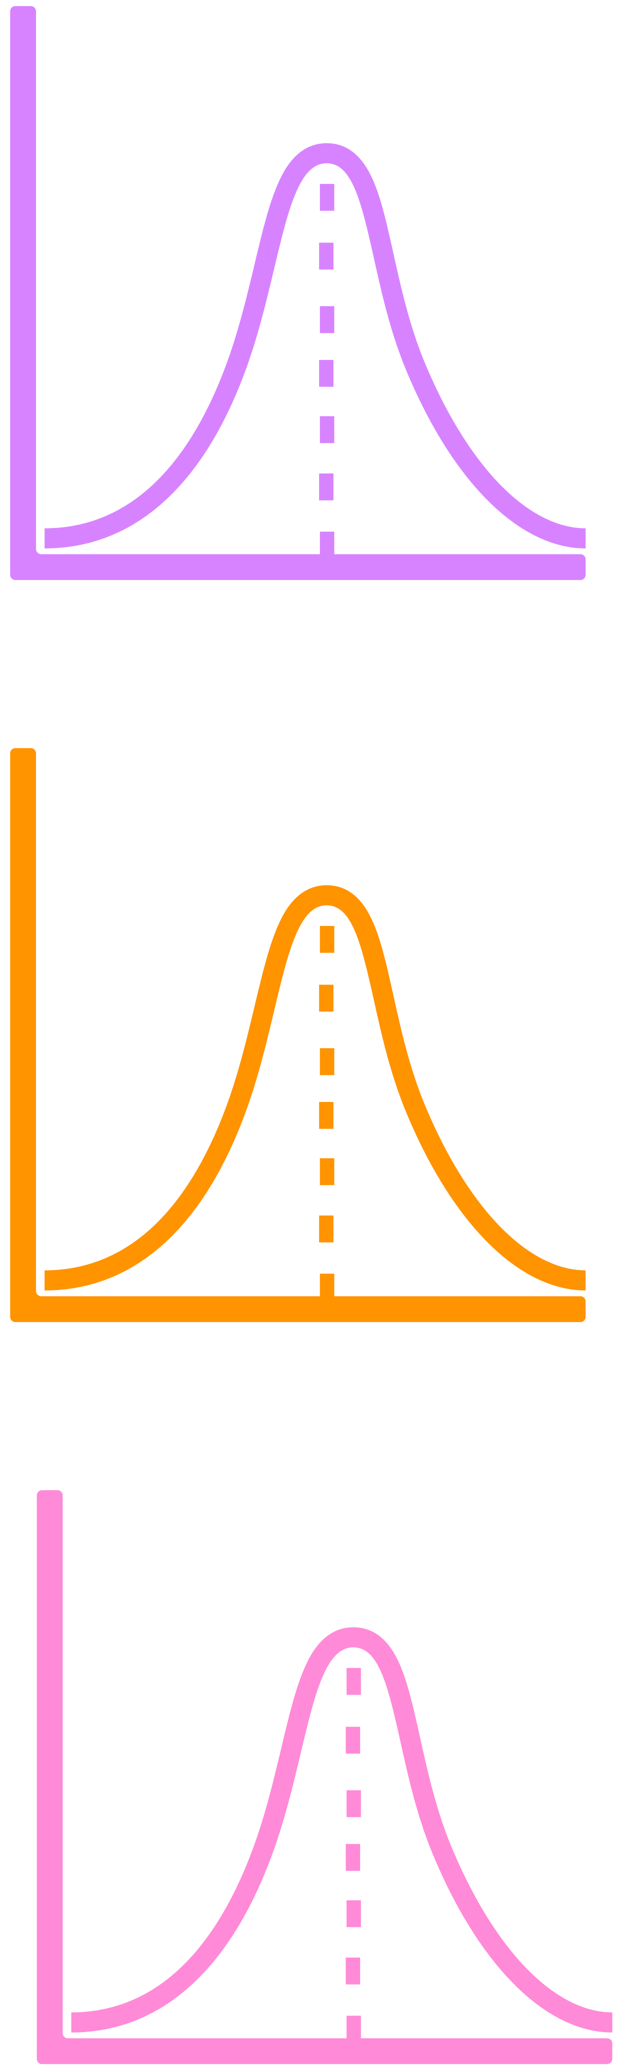

$\mathscr{D}_1$ P_1

$\mathscr{D}_2$ P_2

$\mathscr{D}_3$ P_3

Problem Definition

independent
(not necessarily
identical)

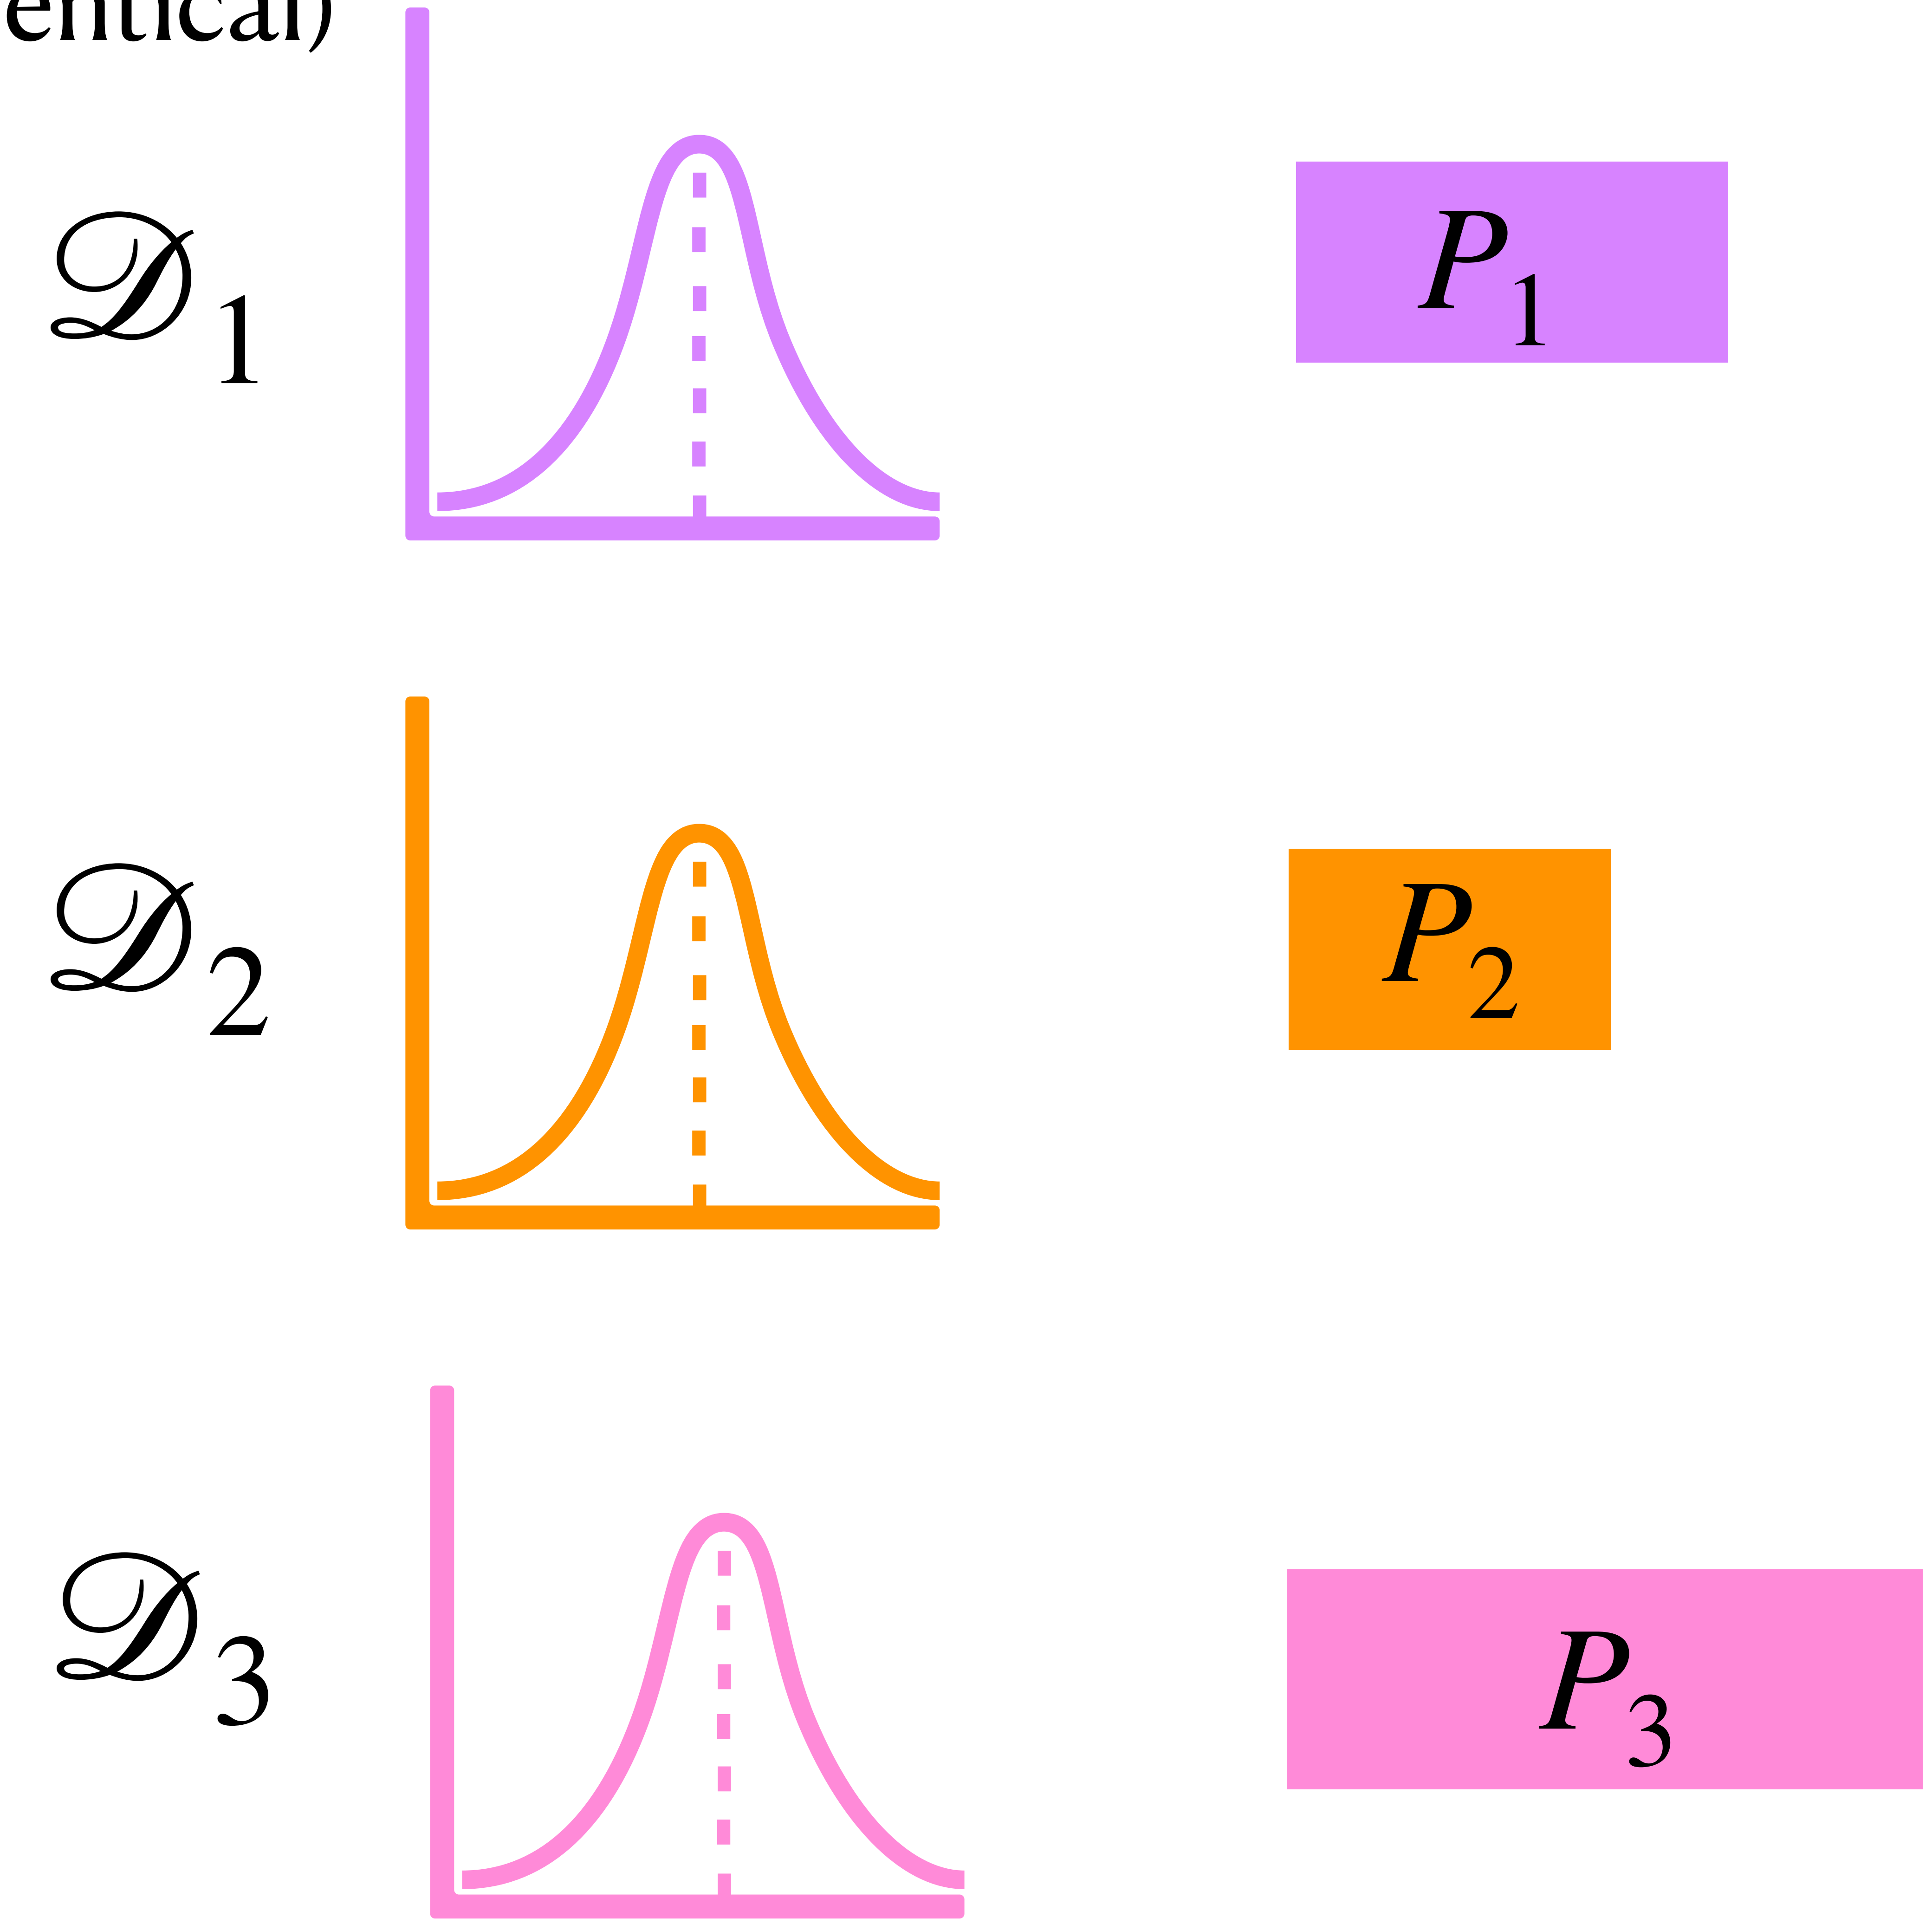

$\mathscr{D}_1$ P_1

$\mathscr{D}_2$ P_2

$\mathscr{D}_3$ P_3

- Nonanticipatory
- Preemption allowed
- No release dates

Problem Definition

independent
(not necessarily
identical)

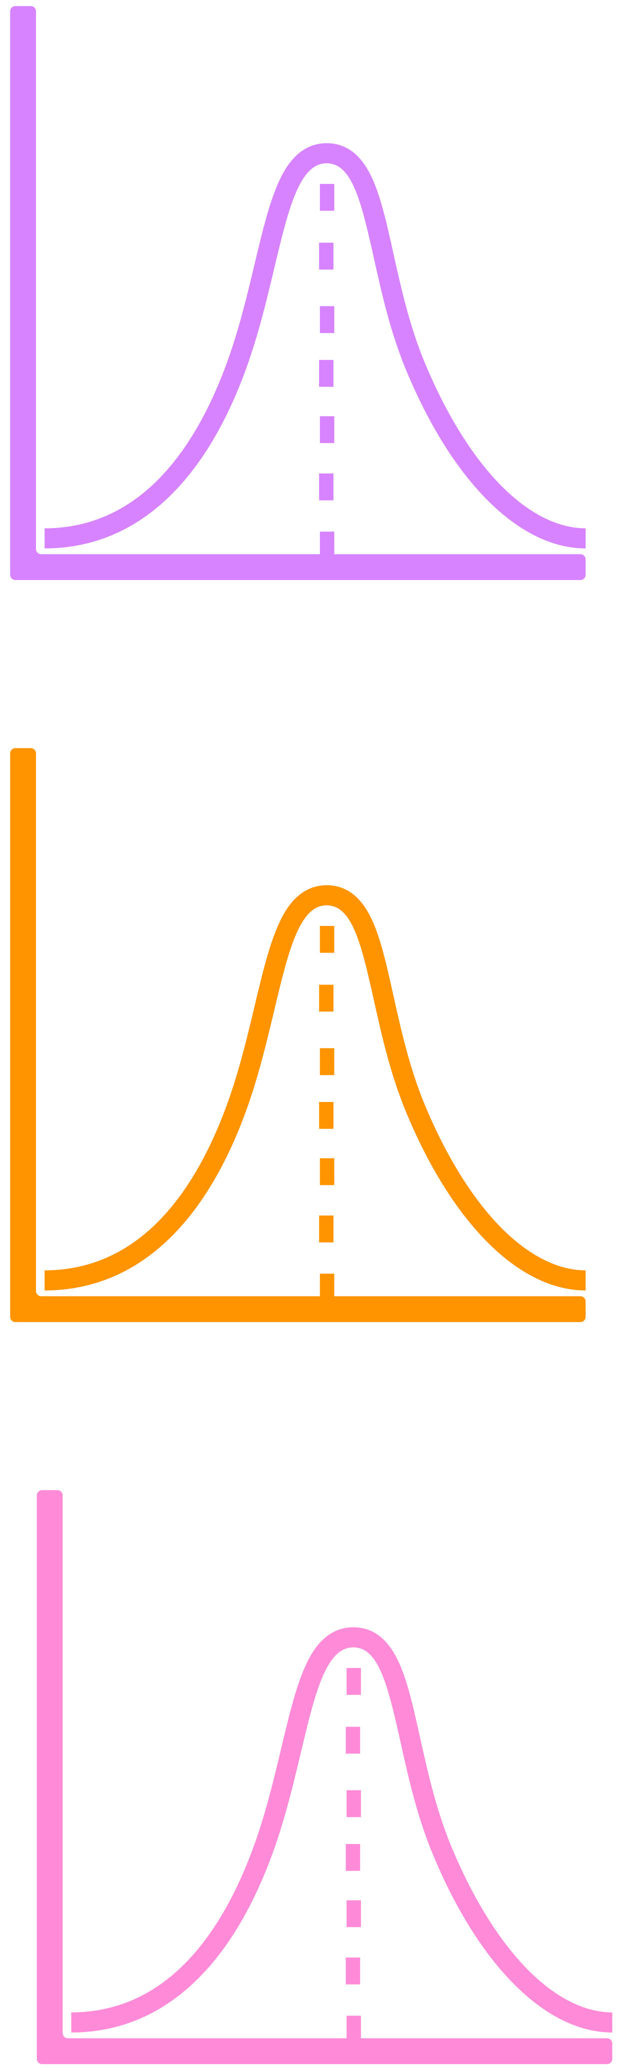

$\mathscr{D}_1$ P_1

$\mathscr{D}_2$ P_2

$\mathscr{D}_3$ P_3

- Nonanticipatory
- Preemption allowed
- No release dates
- Single machine

Problem Definition

independent
(not necessarily
identical)

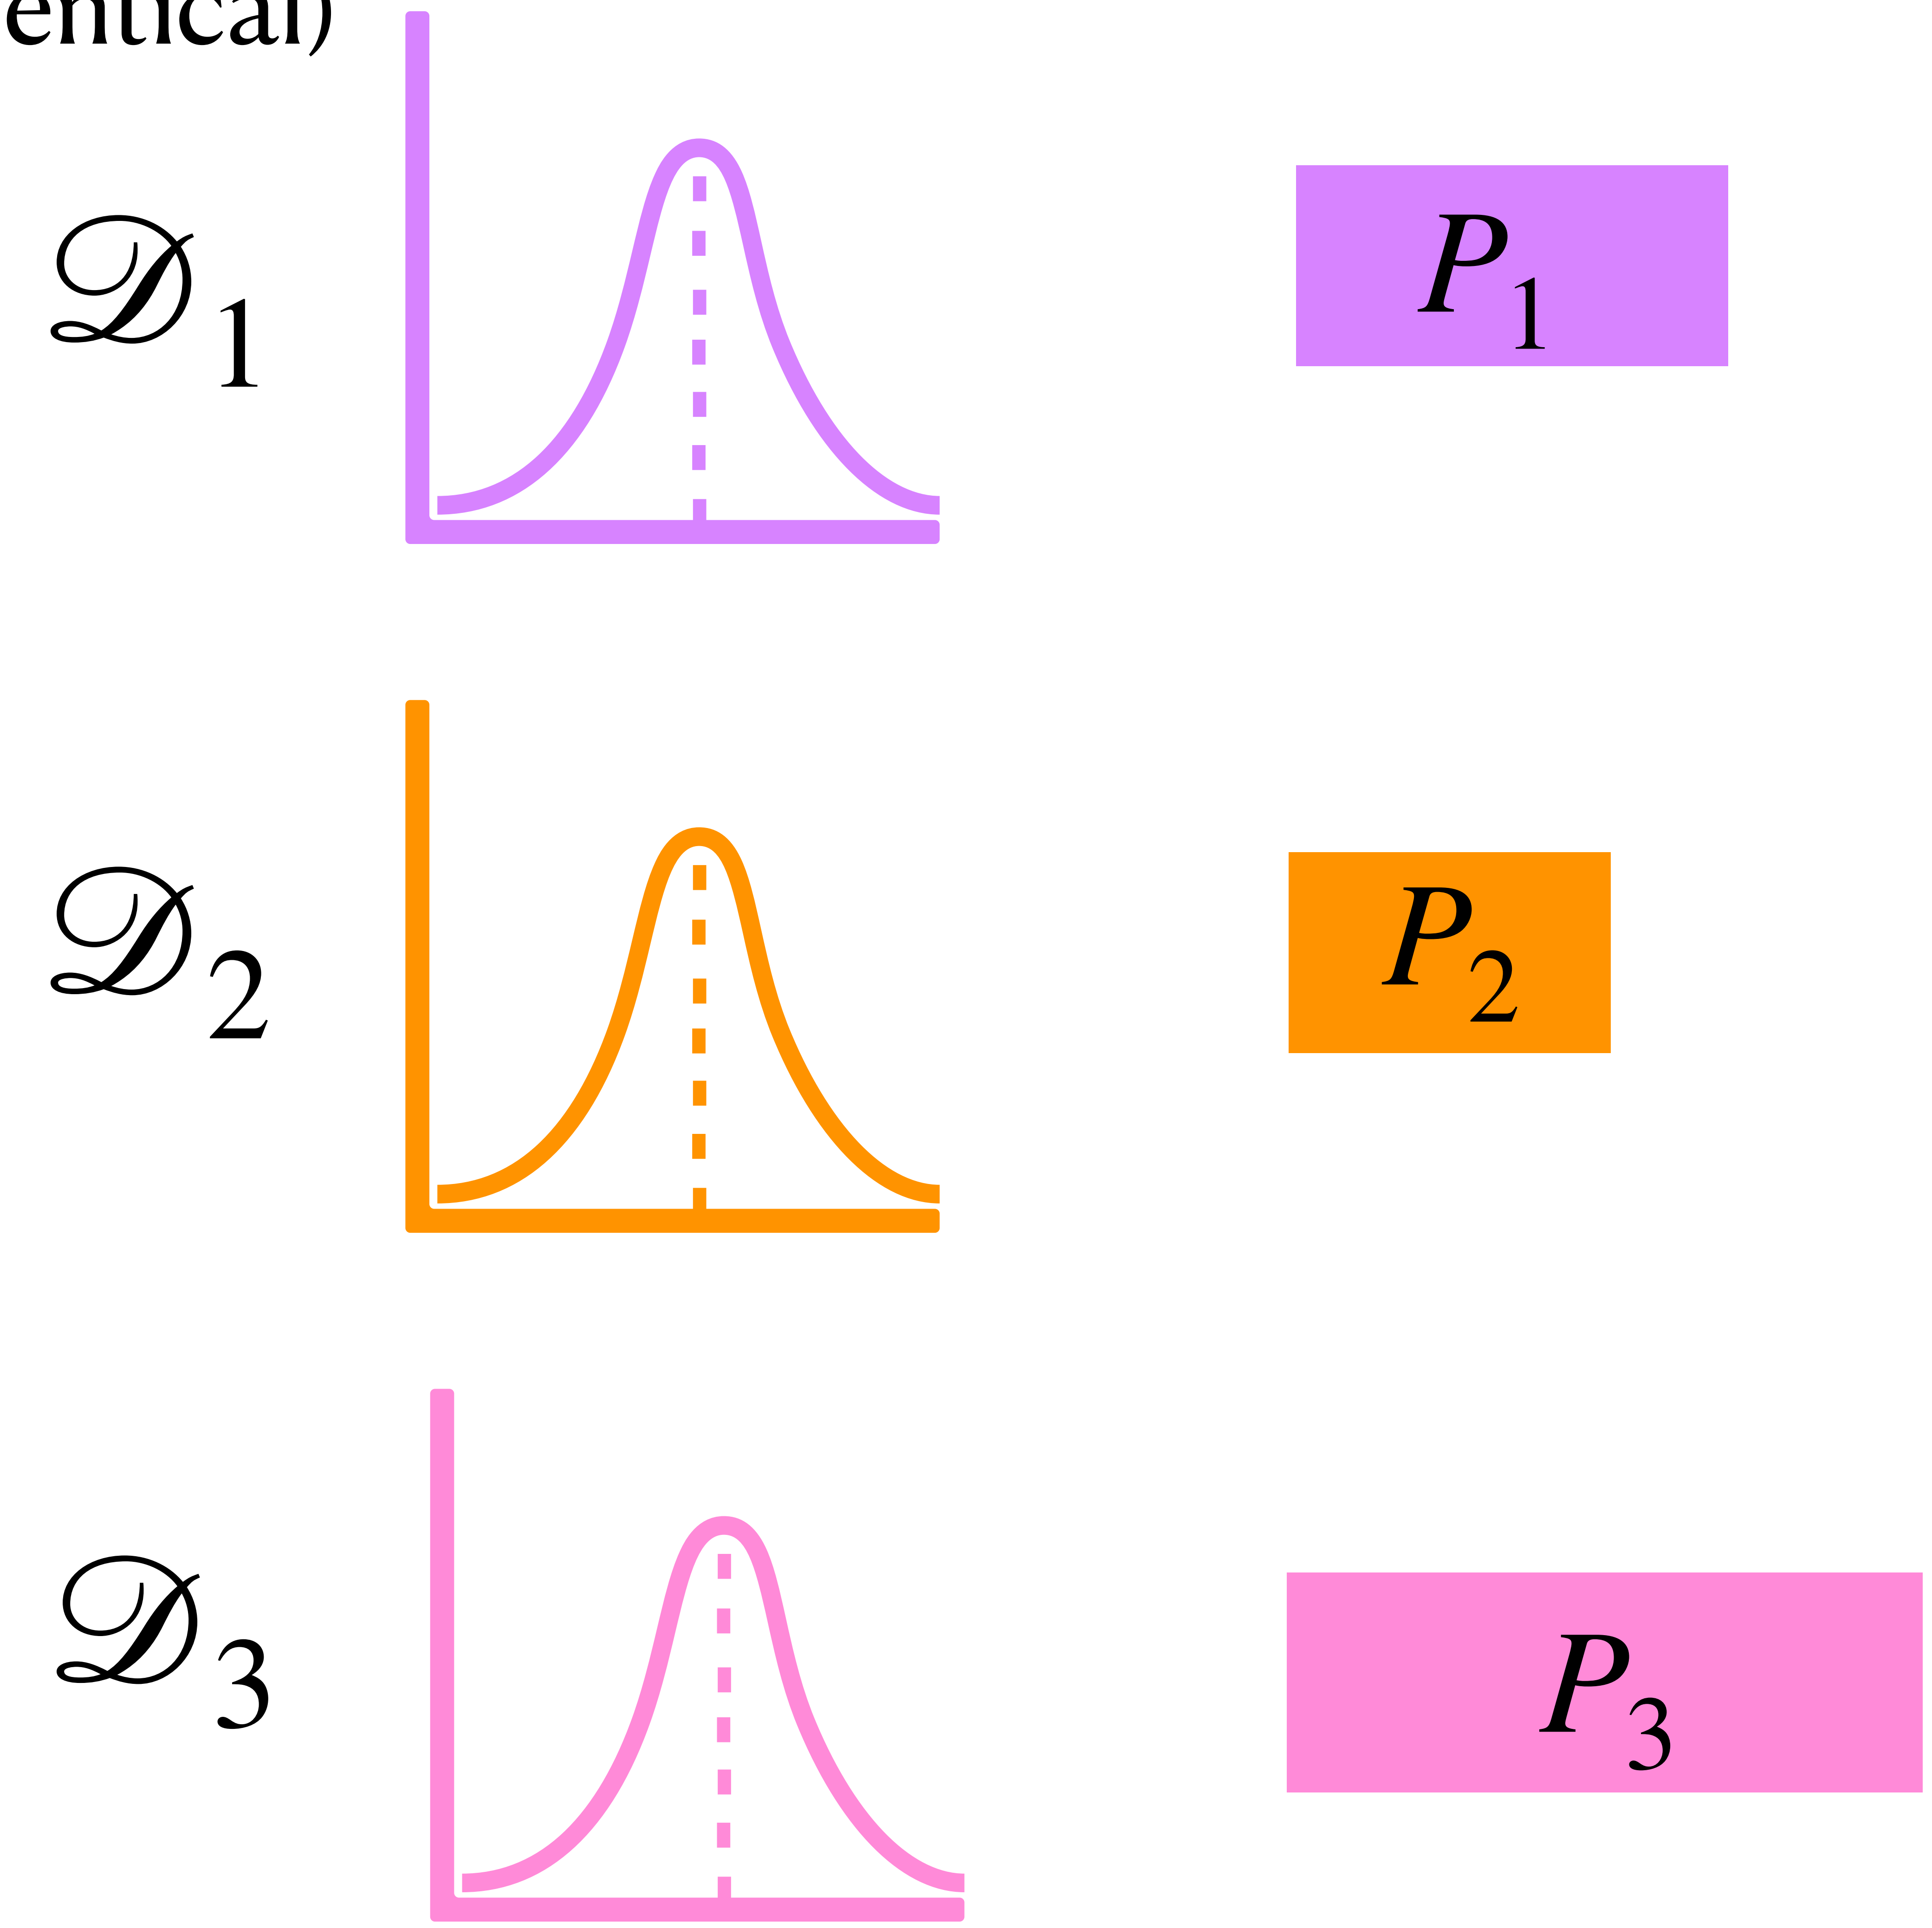

- Nonanticipatory
- Preemption allowed
- No release dates
- Single machine

Problem Definition

independent
(not necessarily
identical)

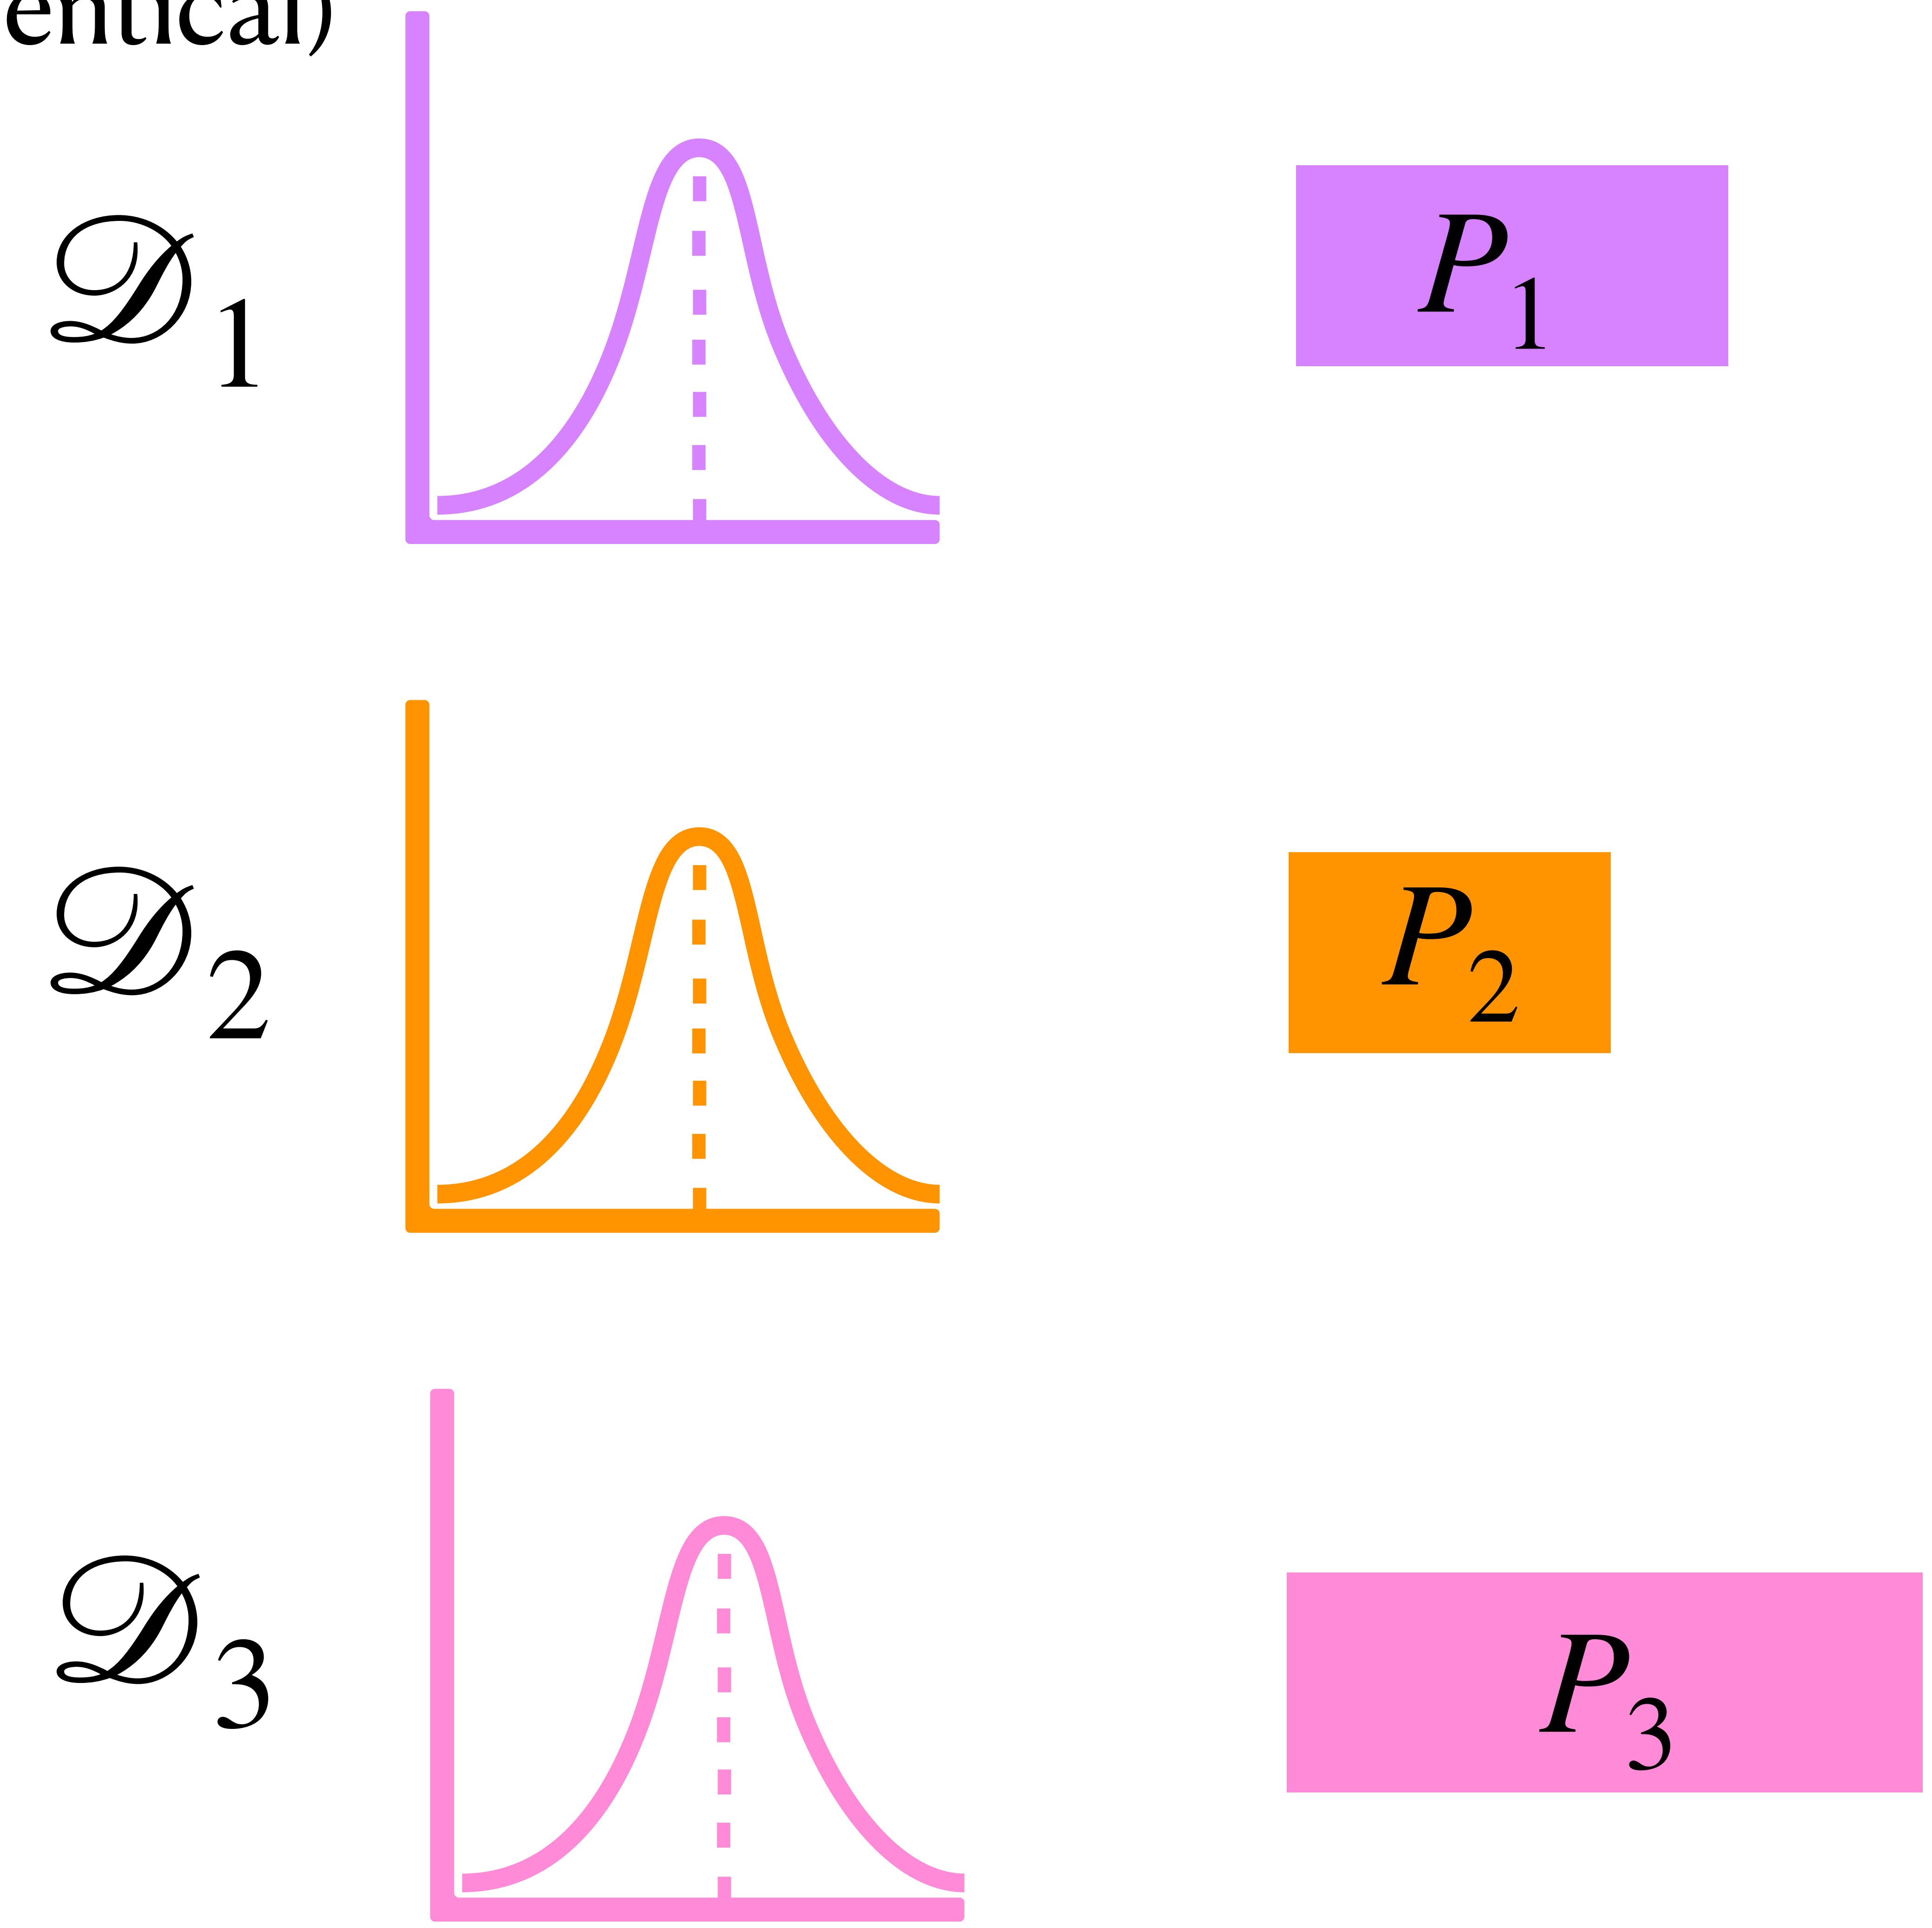

$\mathscr{D}_1$ P_1

$\mathscr{D}_2$ P_2

$\mathscr{D}_3$ P_3

- Nonanticipatory
- Preemption allowed
- No release dates
- Single machine

Problem Definition

independent (not necessarily identical)

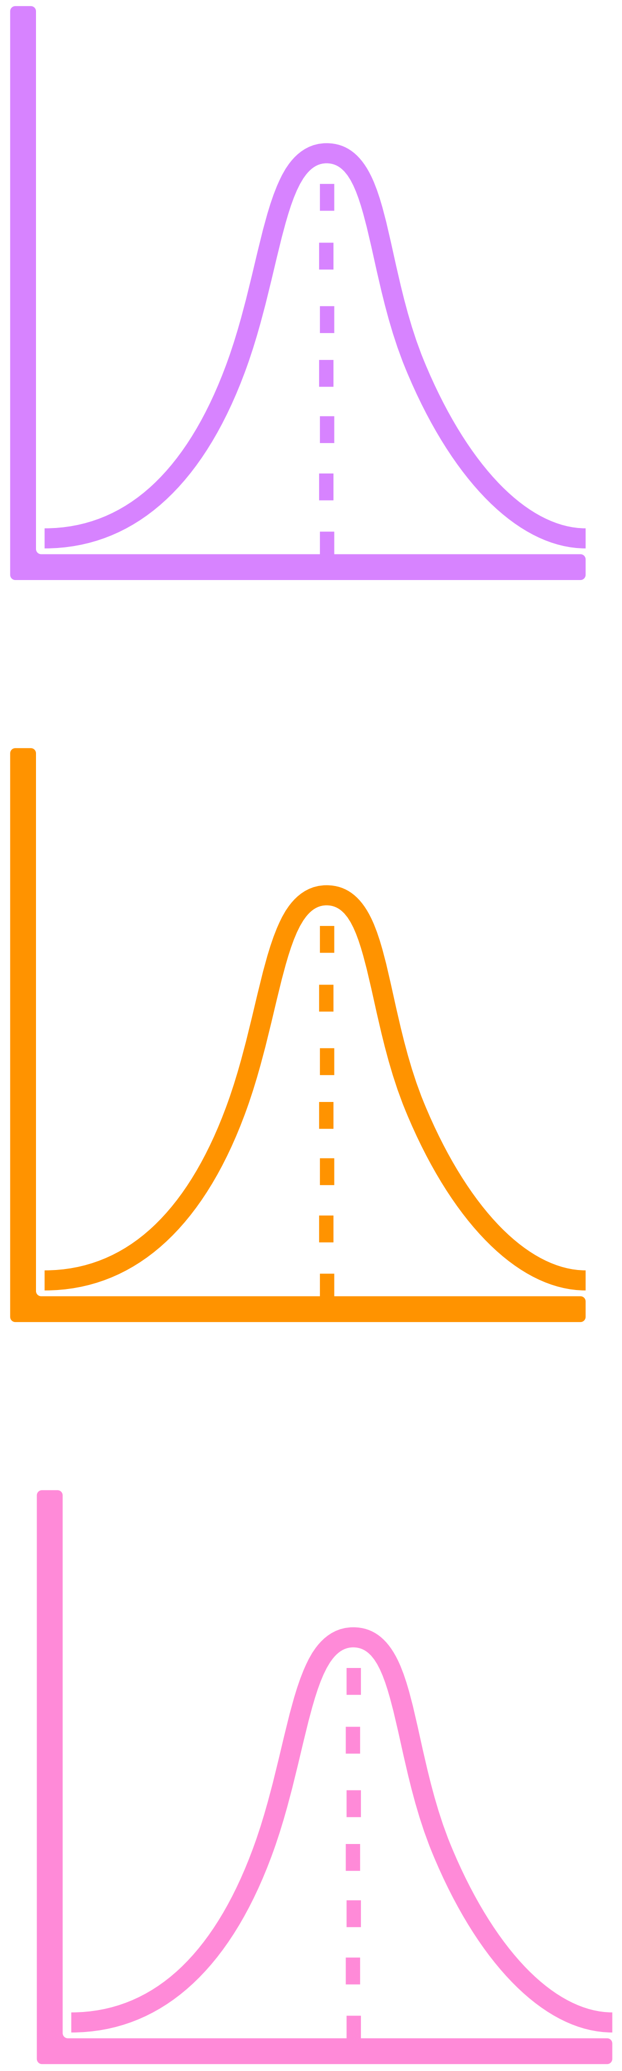

$\mathscr{D}_1$ P_1

$\mathscr{D}_2$ P_2

$\mathscr{D}_3$ P_3

- Nonanticipatory
- Preemption allowed
- No release dates
- Single machine

Goal: nonanticipatory policy minimizing

$$\mathbb{E}\left[\sum_{j=1}^{n} C_j\right]$$

Problem Definition

independent (not necessarily identical)

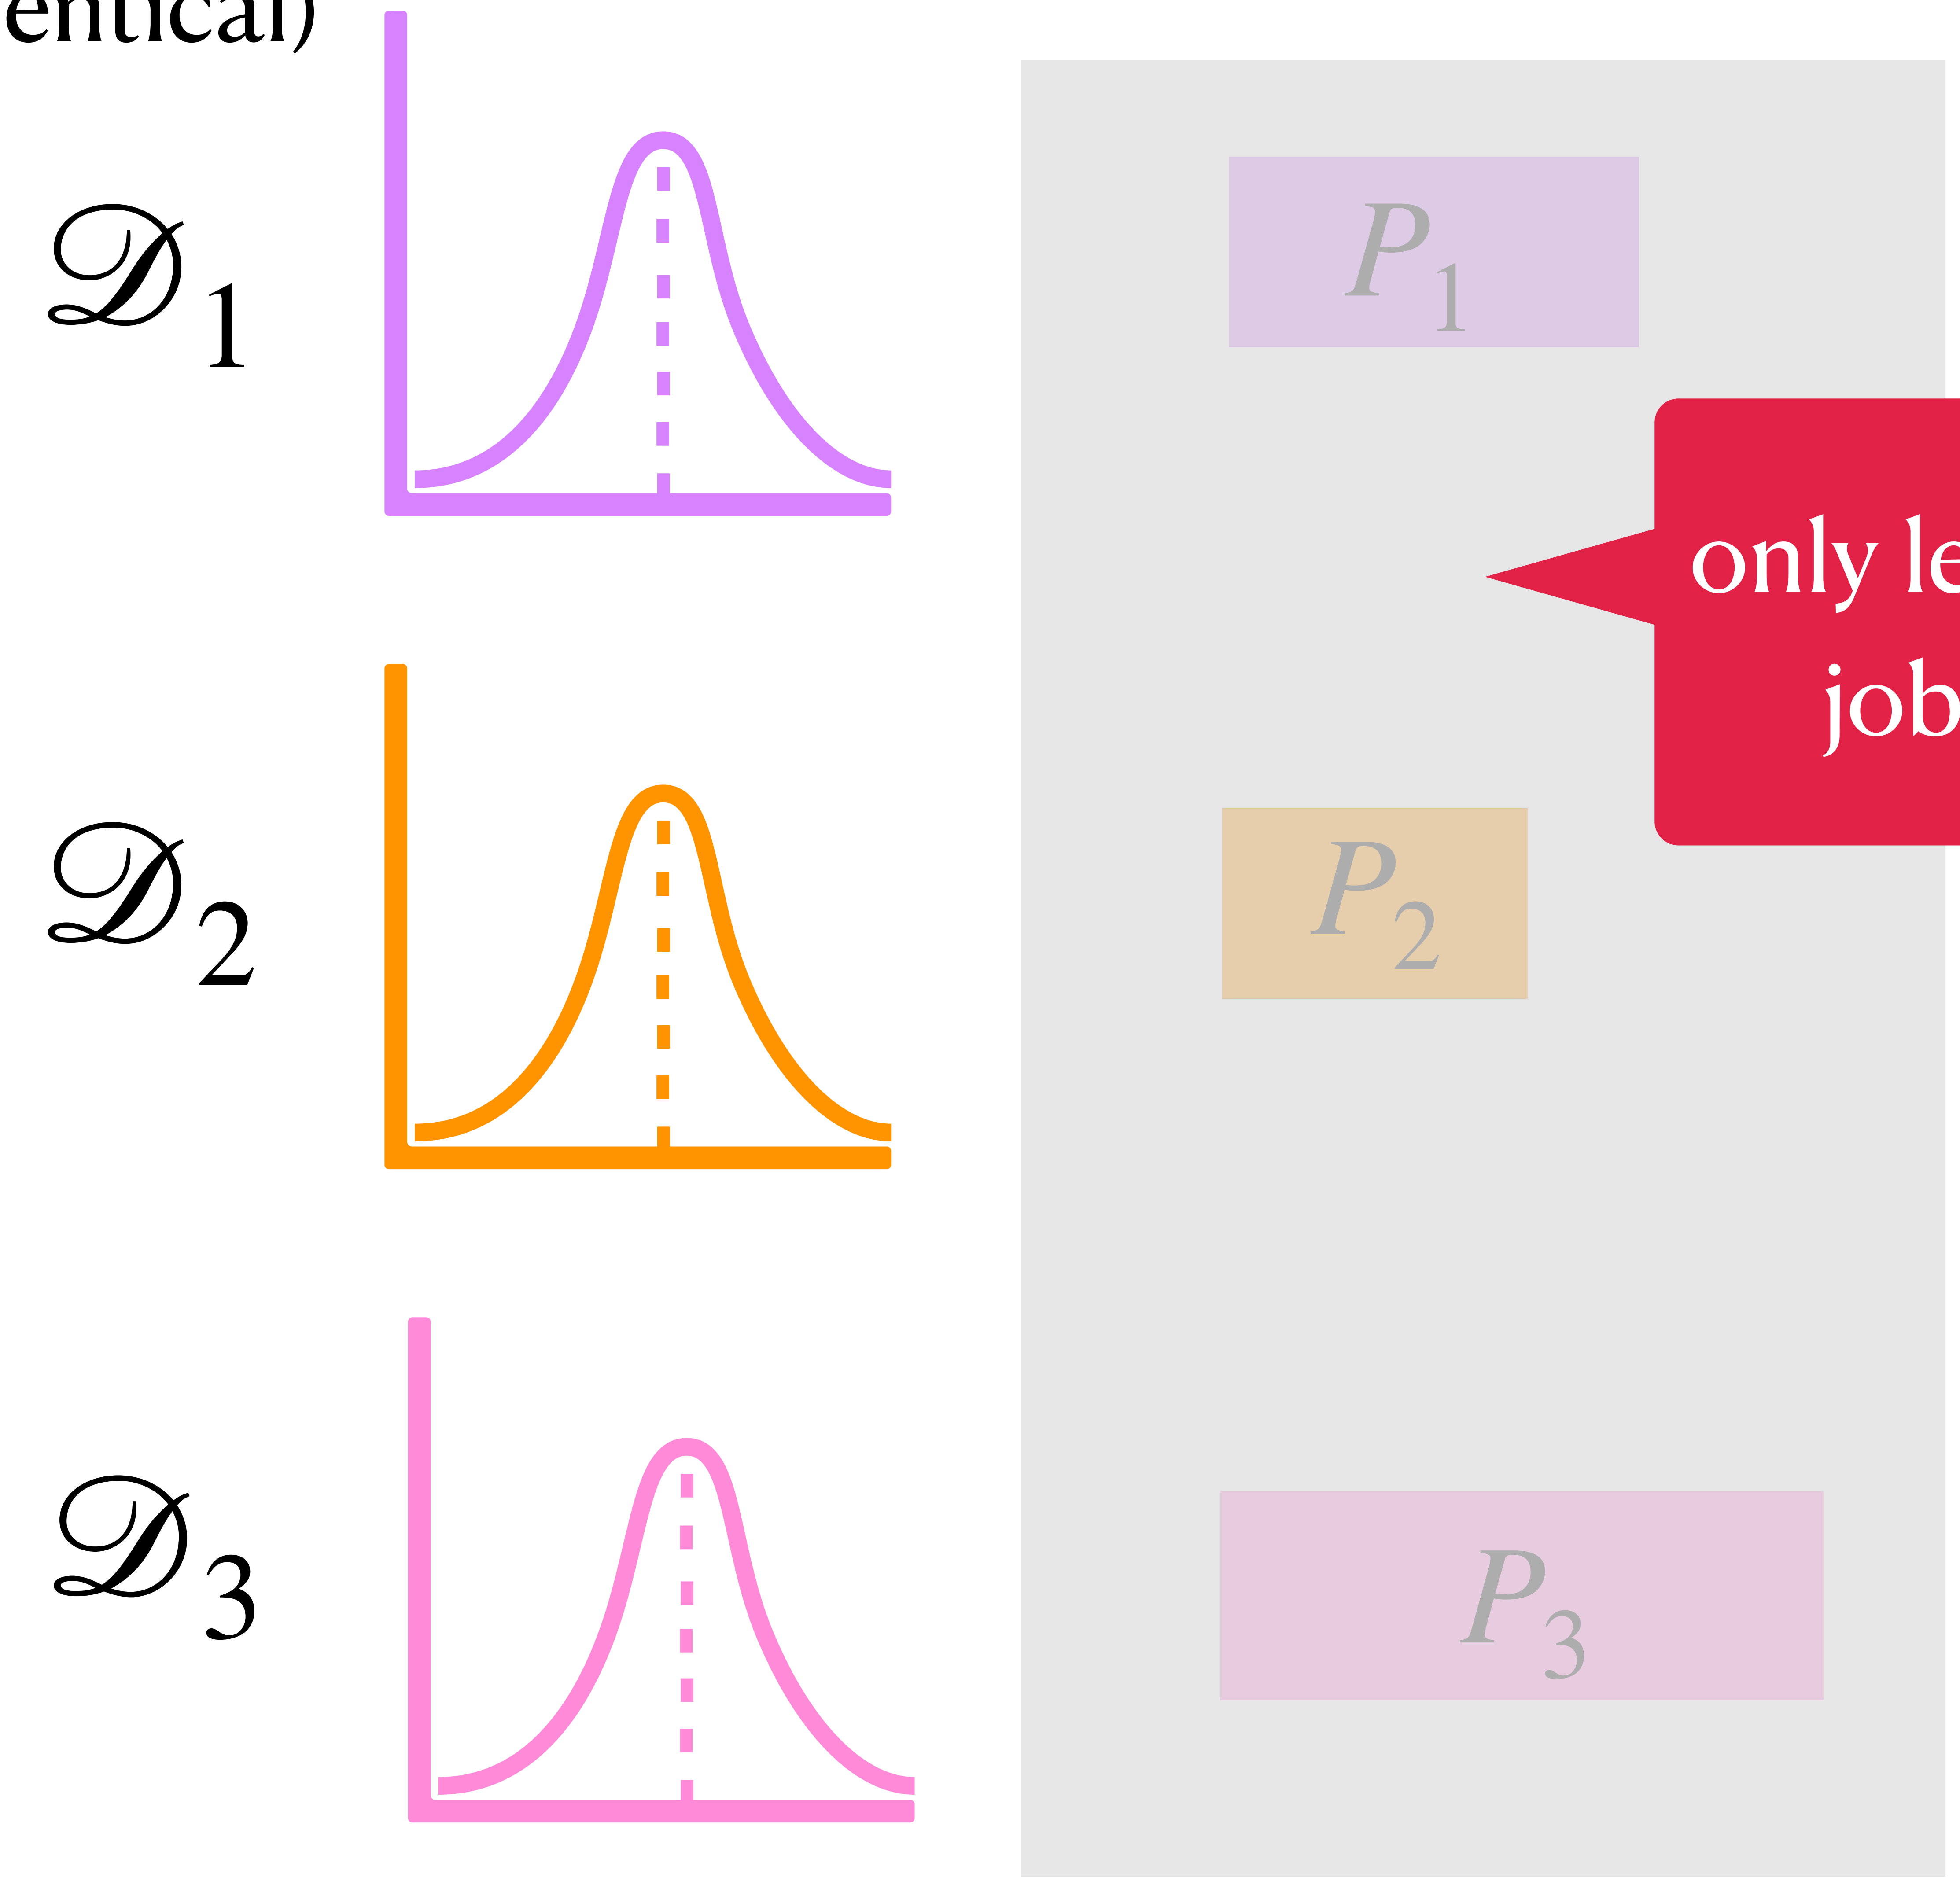

only learn size when job completes

- Nonanticipatory
- Preemption allowed
- No release dates
- Single machine

Goal: nonanticipatory policy minimizing

$$\mathbb{E}\left[\sum_{j=1}^{n} C_j\right]$$

Optimal Policies

independent
(not necessarily
identical)

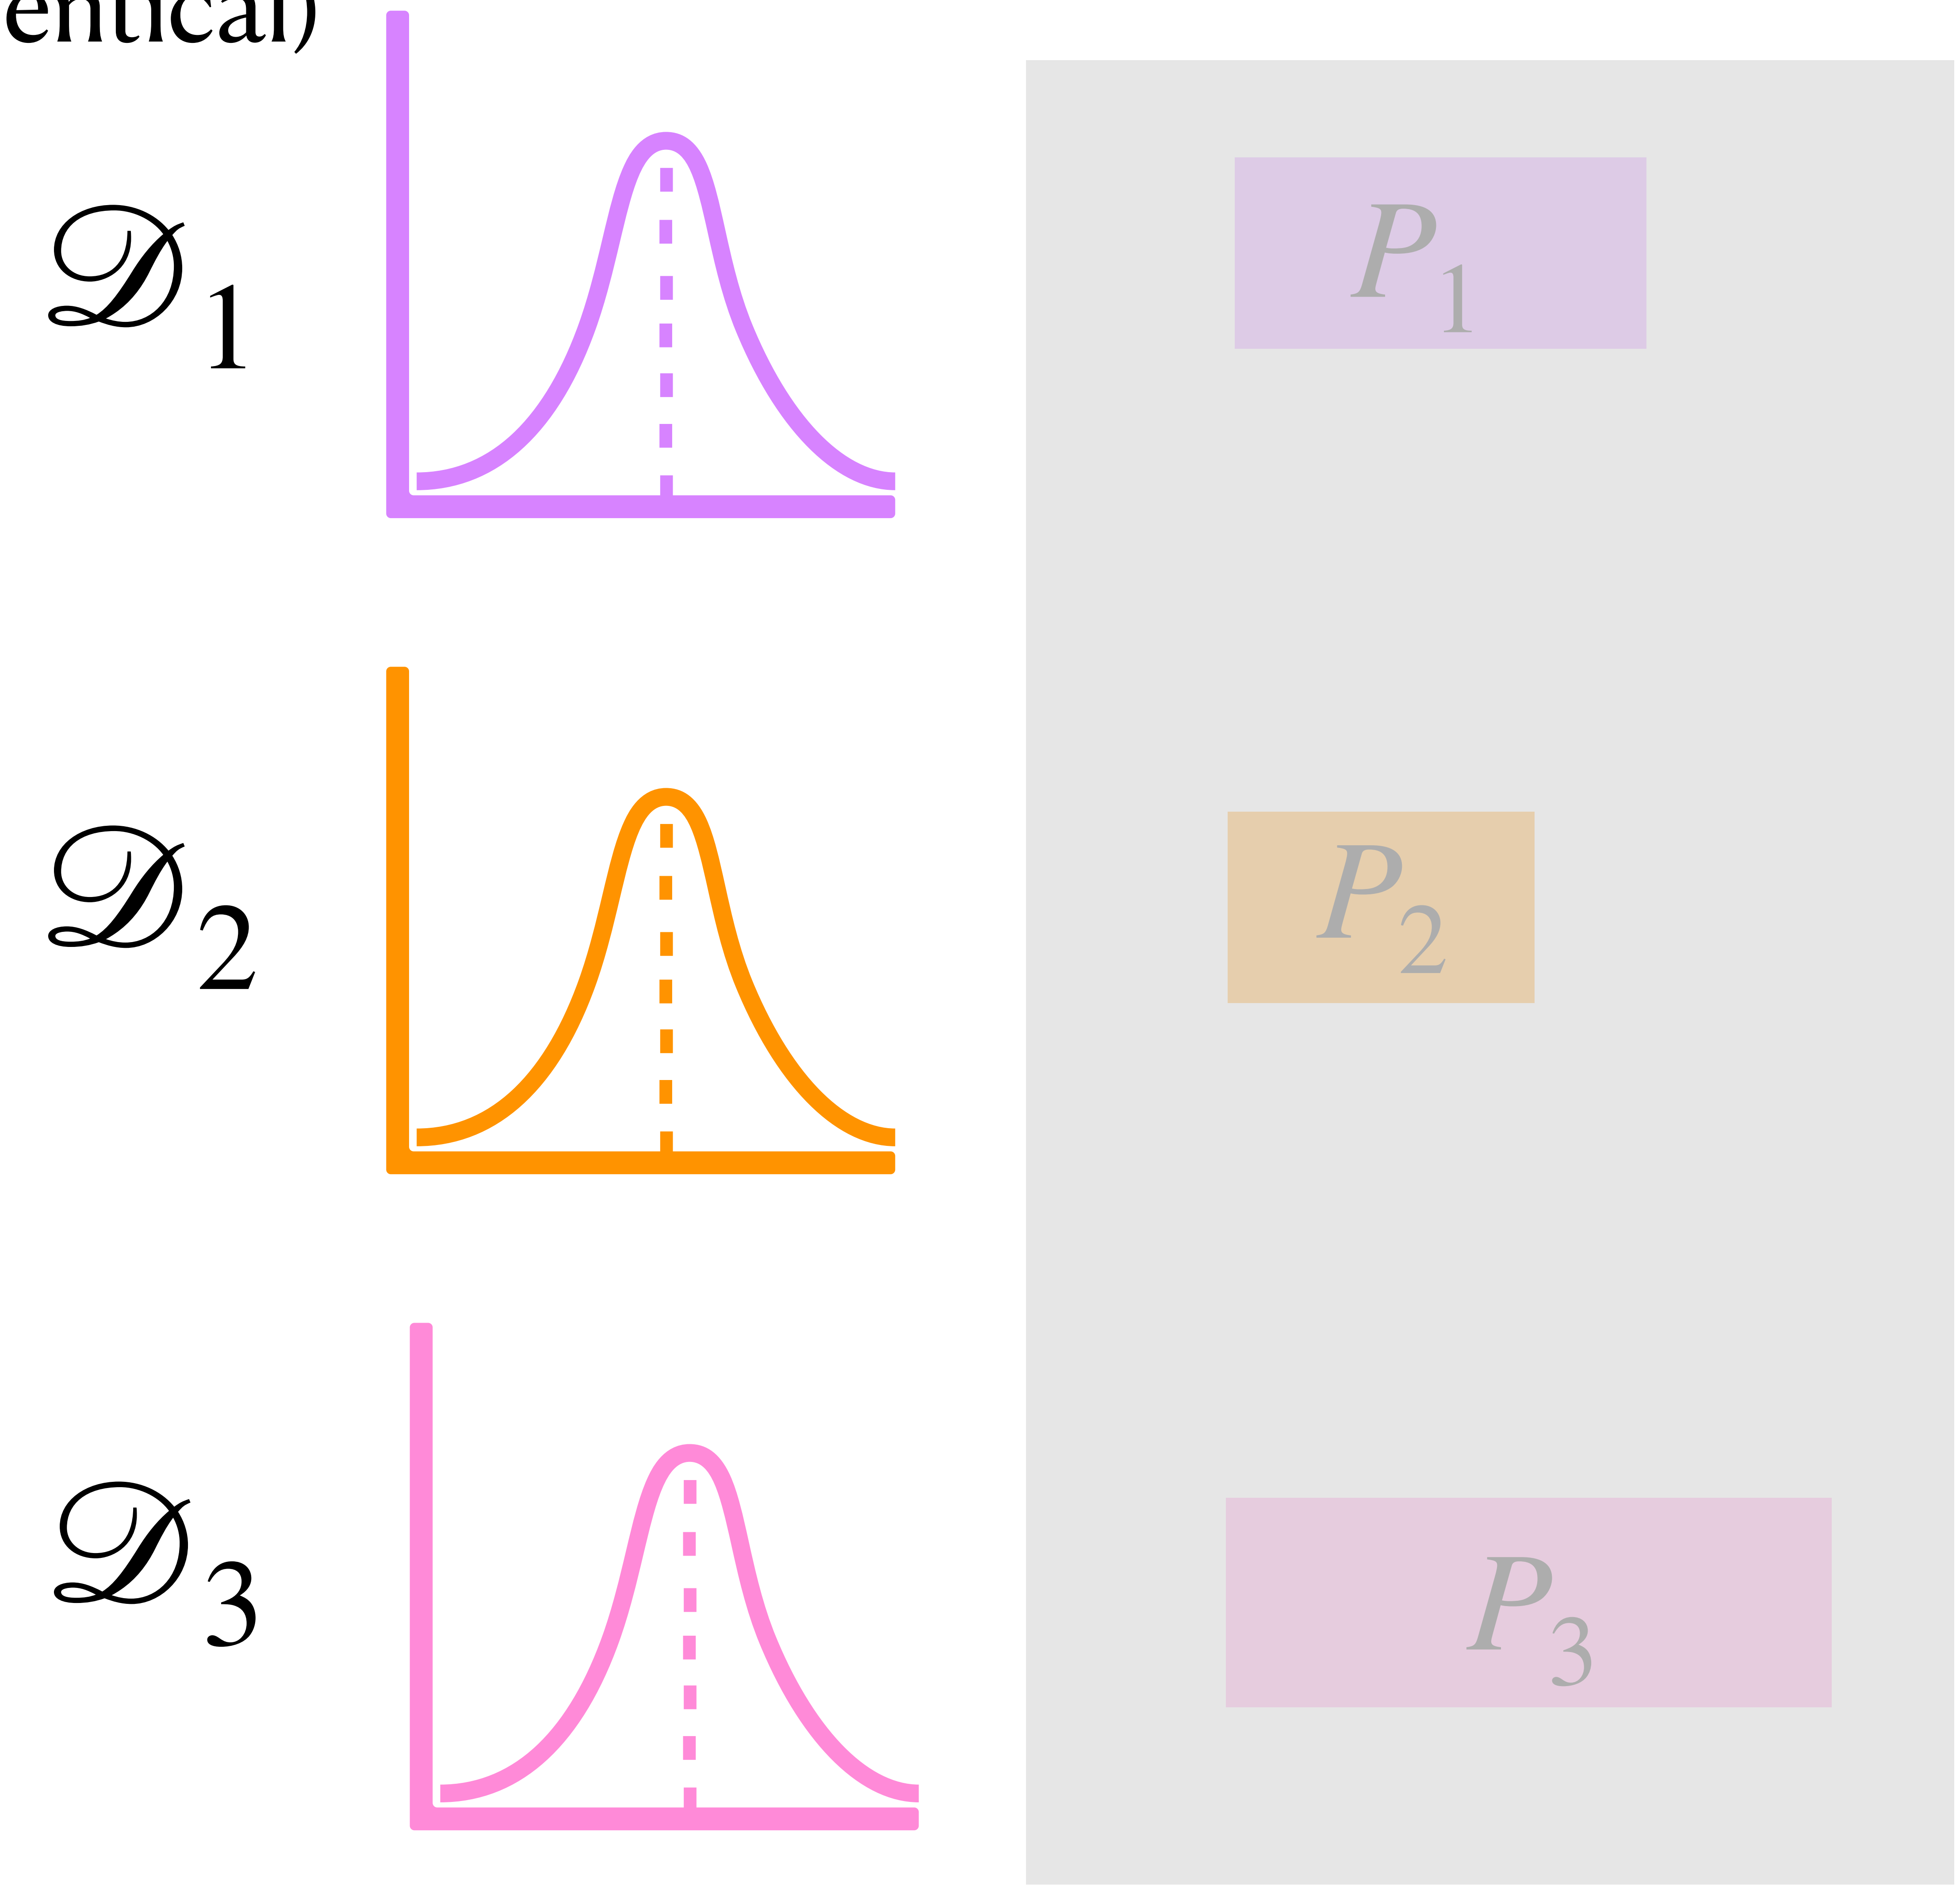

Goal: nonanticipatory policy minimizing

$$\mathbb{E}\left[\sum_{j=1}^{n} C_j\right]$$

Optimal Policies

independent
(not necessarily
identical)

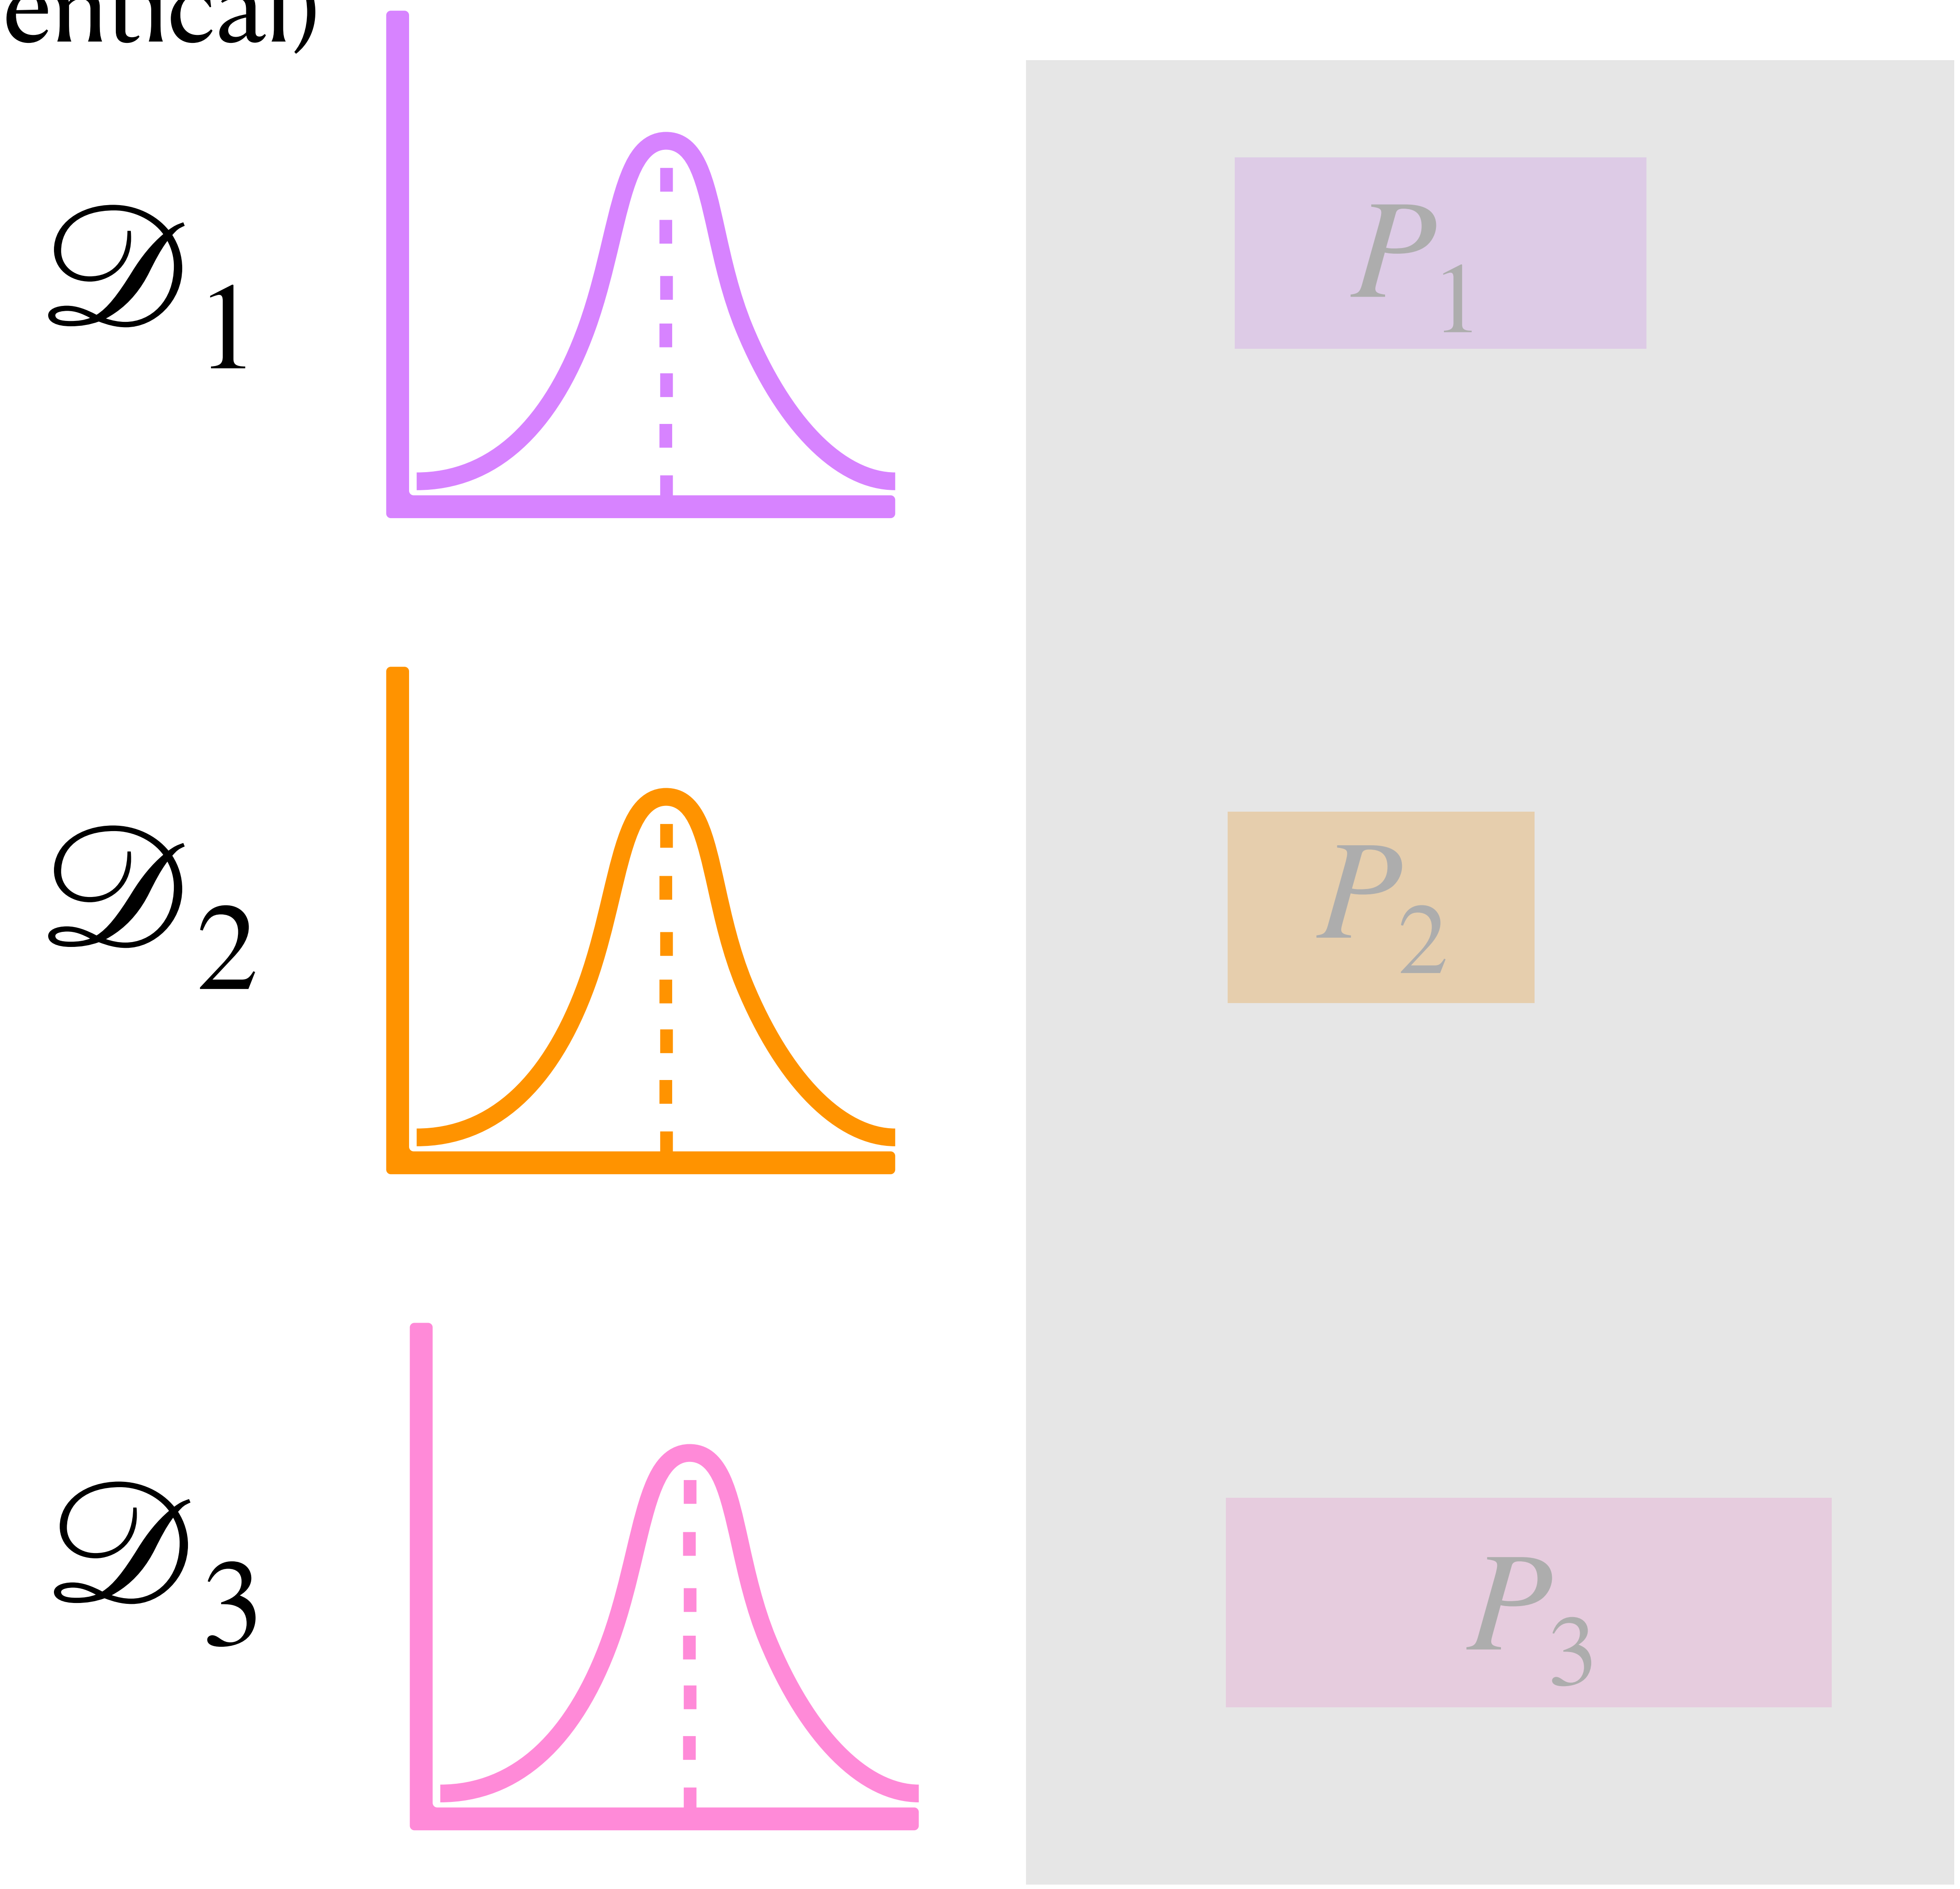

Goal: nonanticipatory policy minimizing

$$\mathbb{E}\left[\sum_{j=1}^{n} C_j\right]$$

- **Deterministic:** SPT

Optimal Policies

independent
(not necessarily
identical)

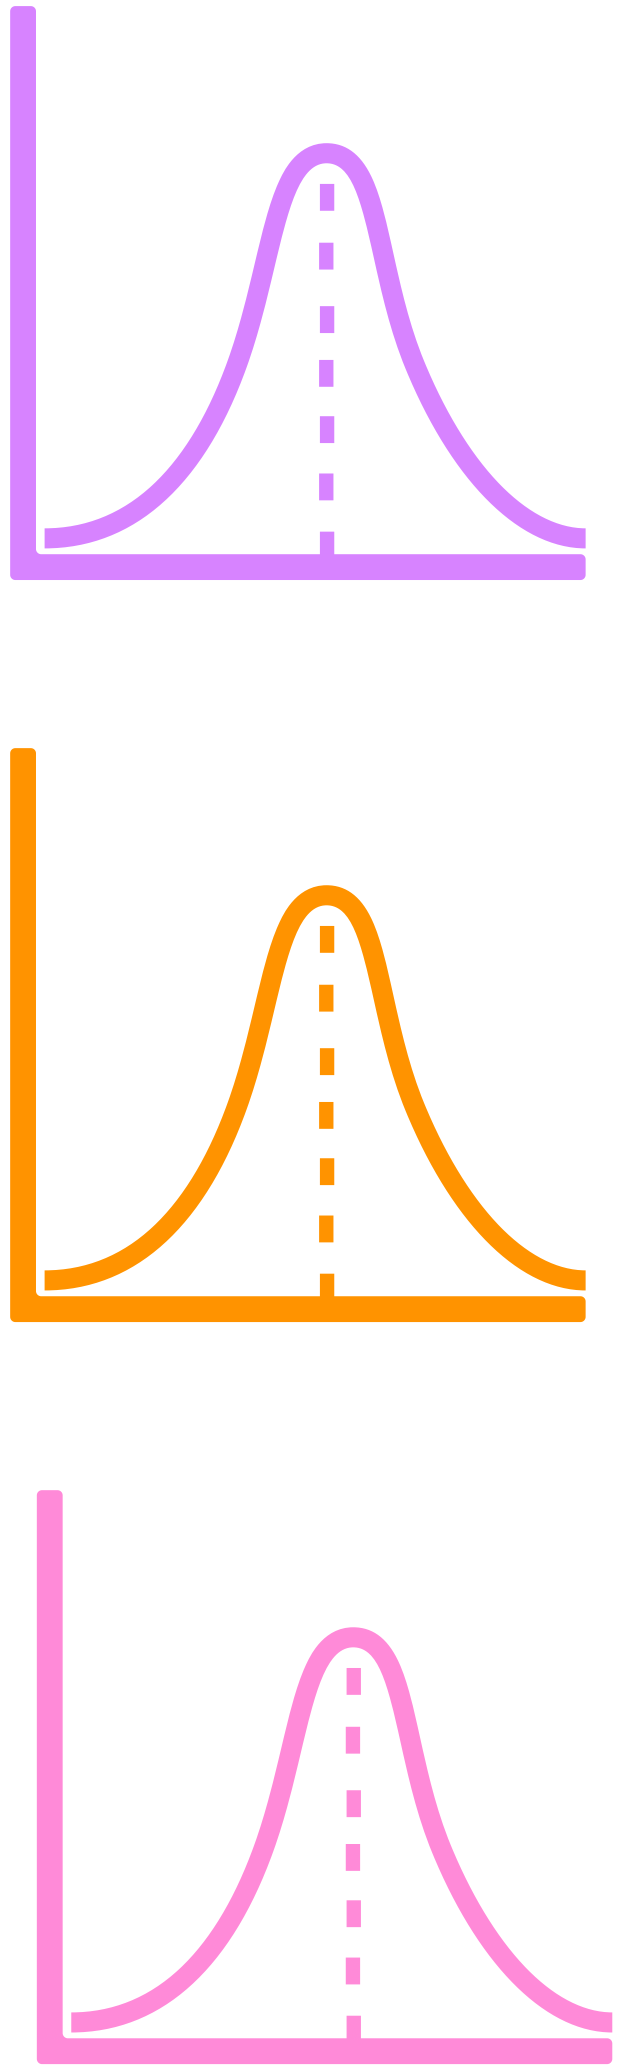

$\mathscr{D}_1$

$\mathscr{D}_2$

$\mathscr{D}_3$

P_1

P_2

P_3

Goal: nonanticipatory policy minimizing

$$\mathbb{E}\left[\sum_{j=1}^{n} C_j\right]$$

- **Deterministic:** SPT
- **Decreasing hazard rates:** SERPT

Optimal Policies

independent
(not necessarily
identical)

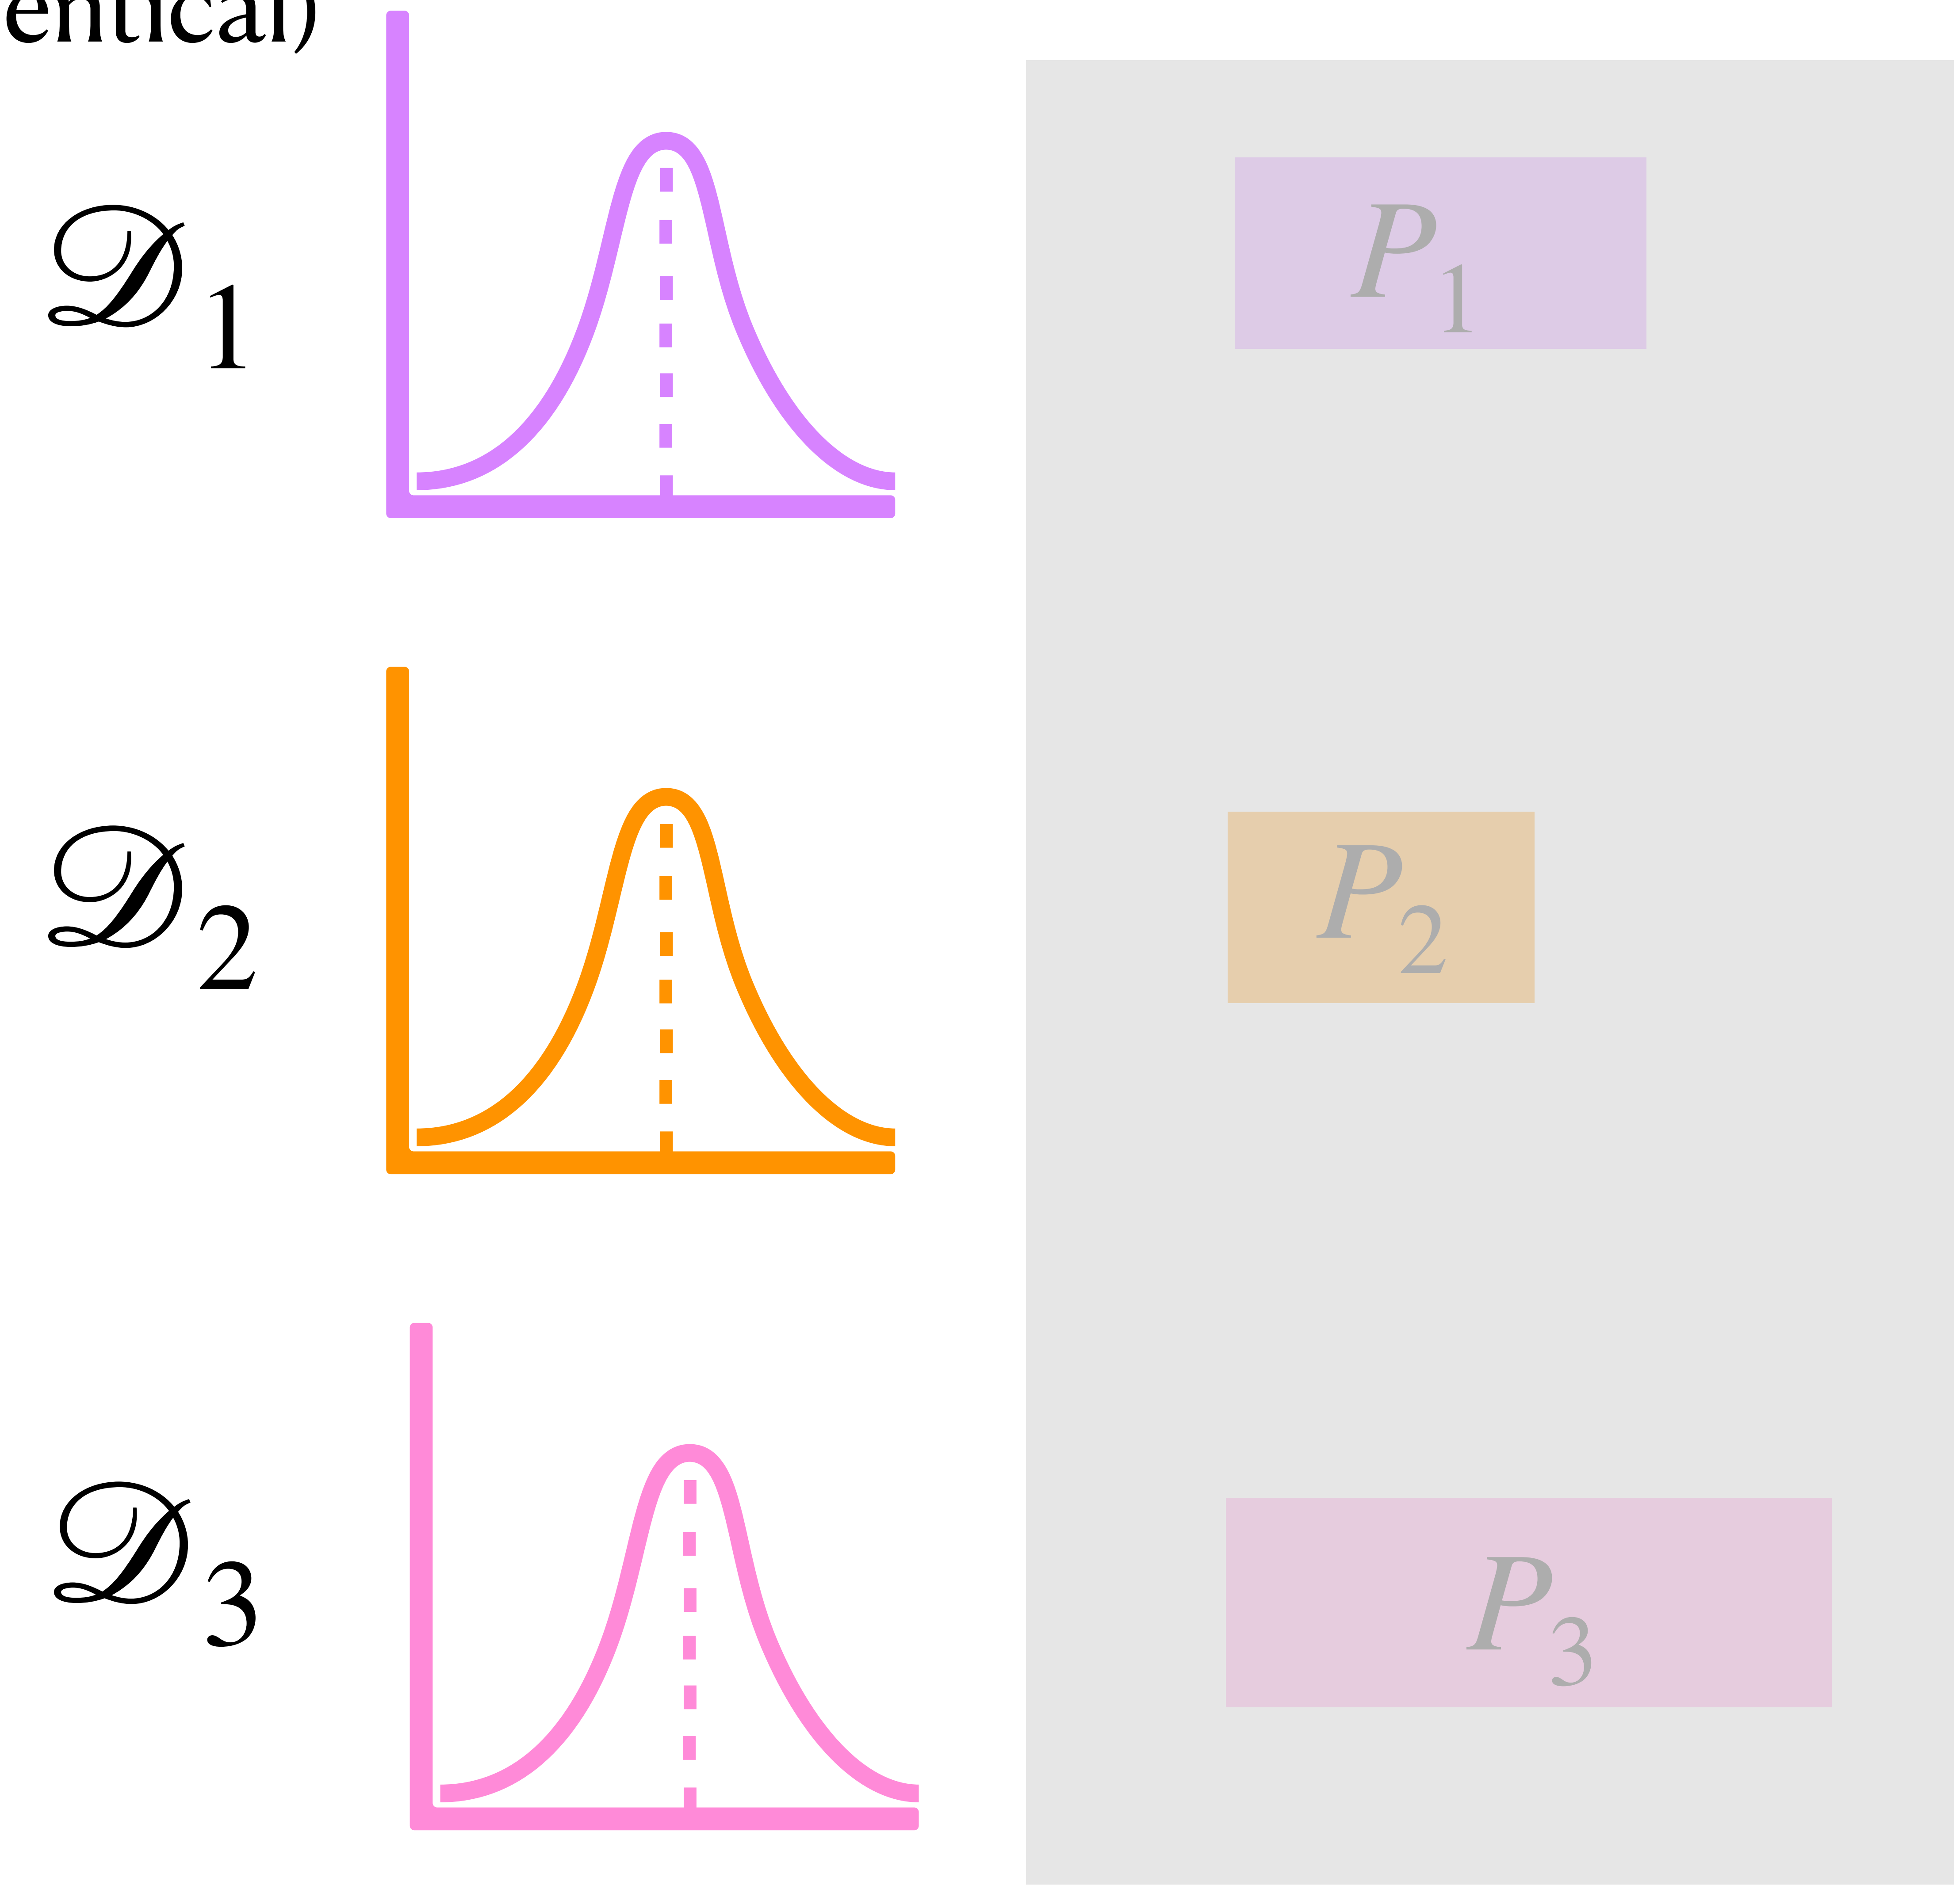

Goal: nonanticipatory policy minimizing

$$\mathbb{E}\left[\sum_{j=1}^{n} C_j\right]$$

- **Deterministic:** SPT
- **Decreasing hazard rates:** SERPT
- **Increasing hazard rates:** SEPT

Optimal Policies

independent
(not necessarily
identical)

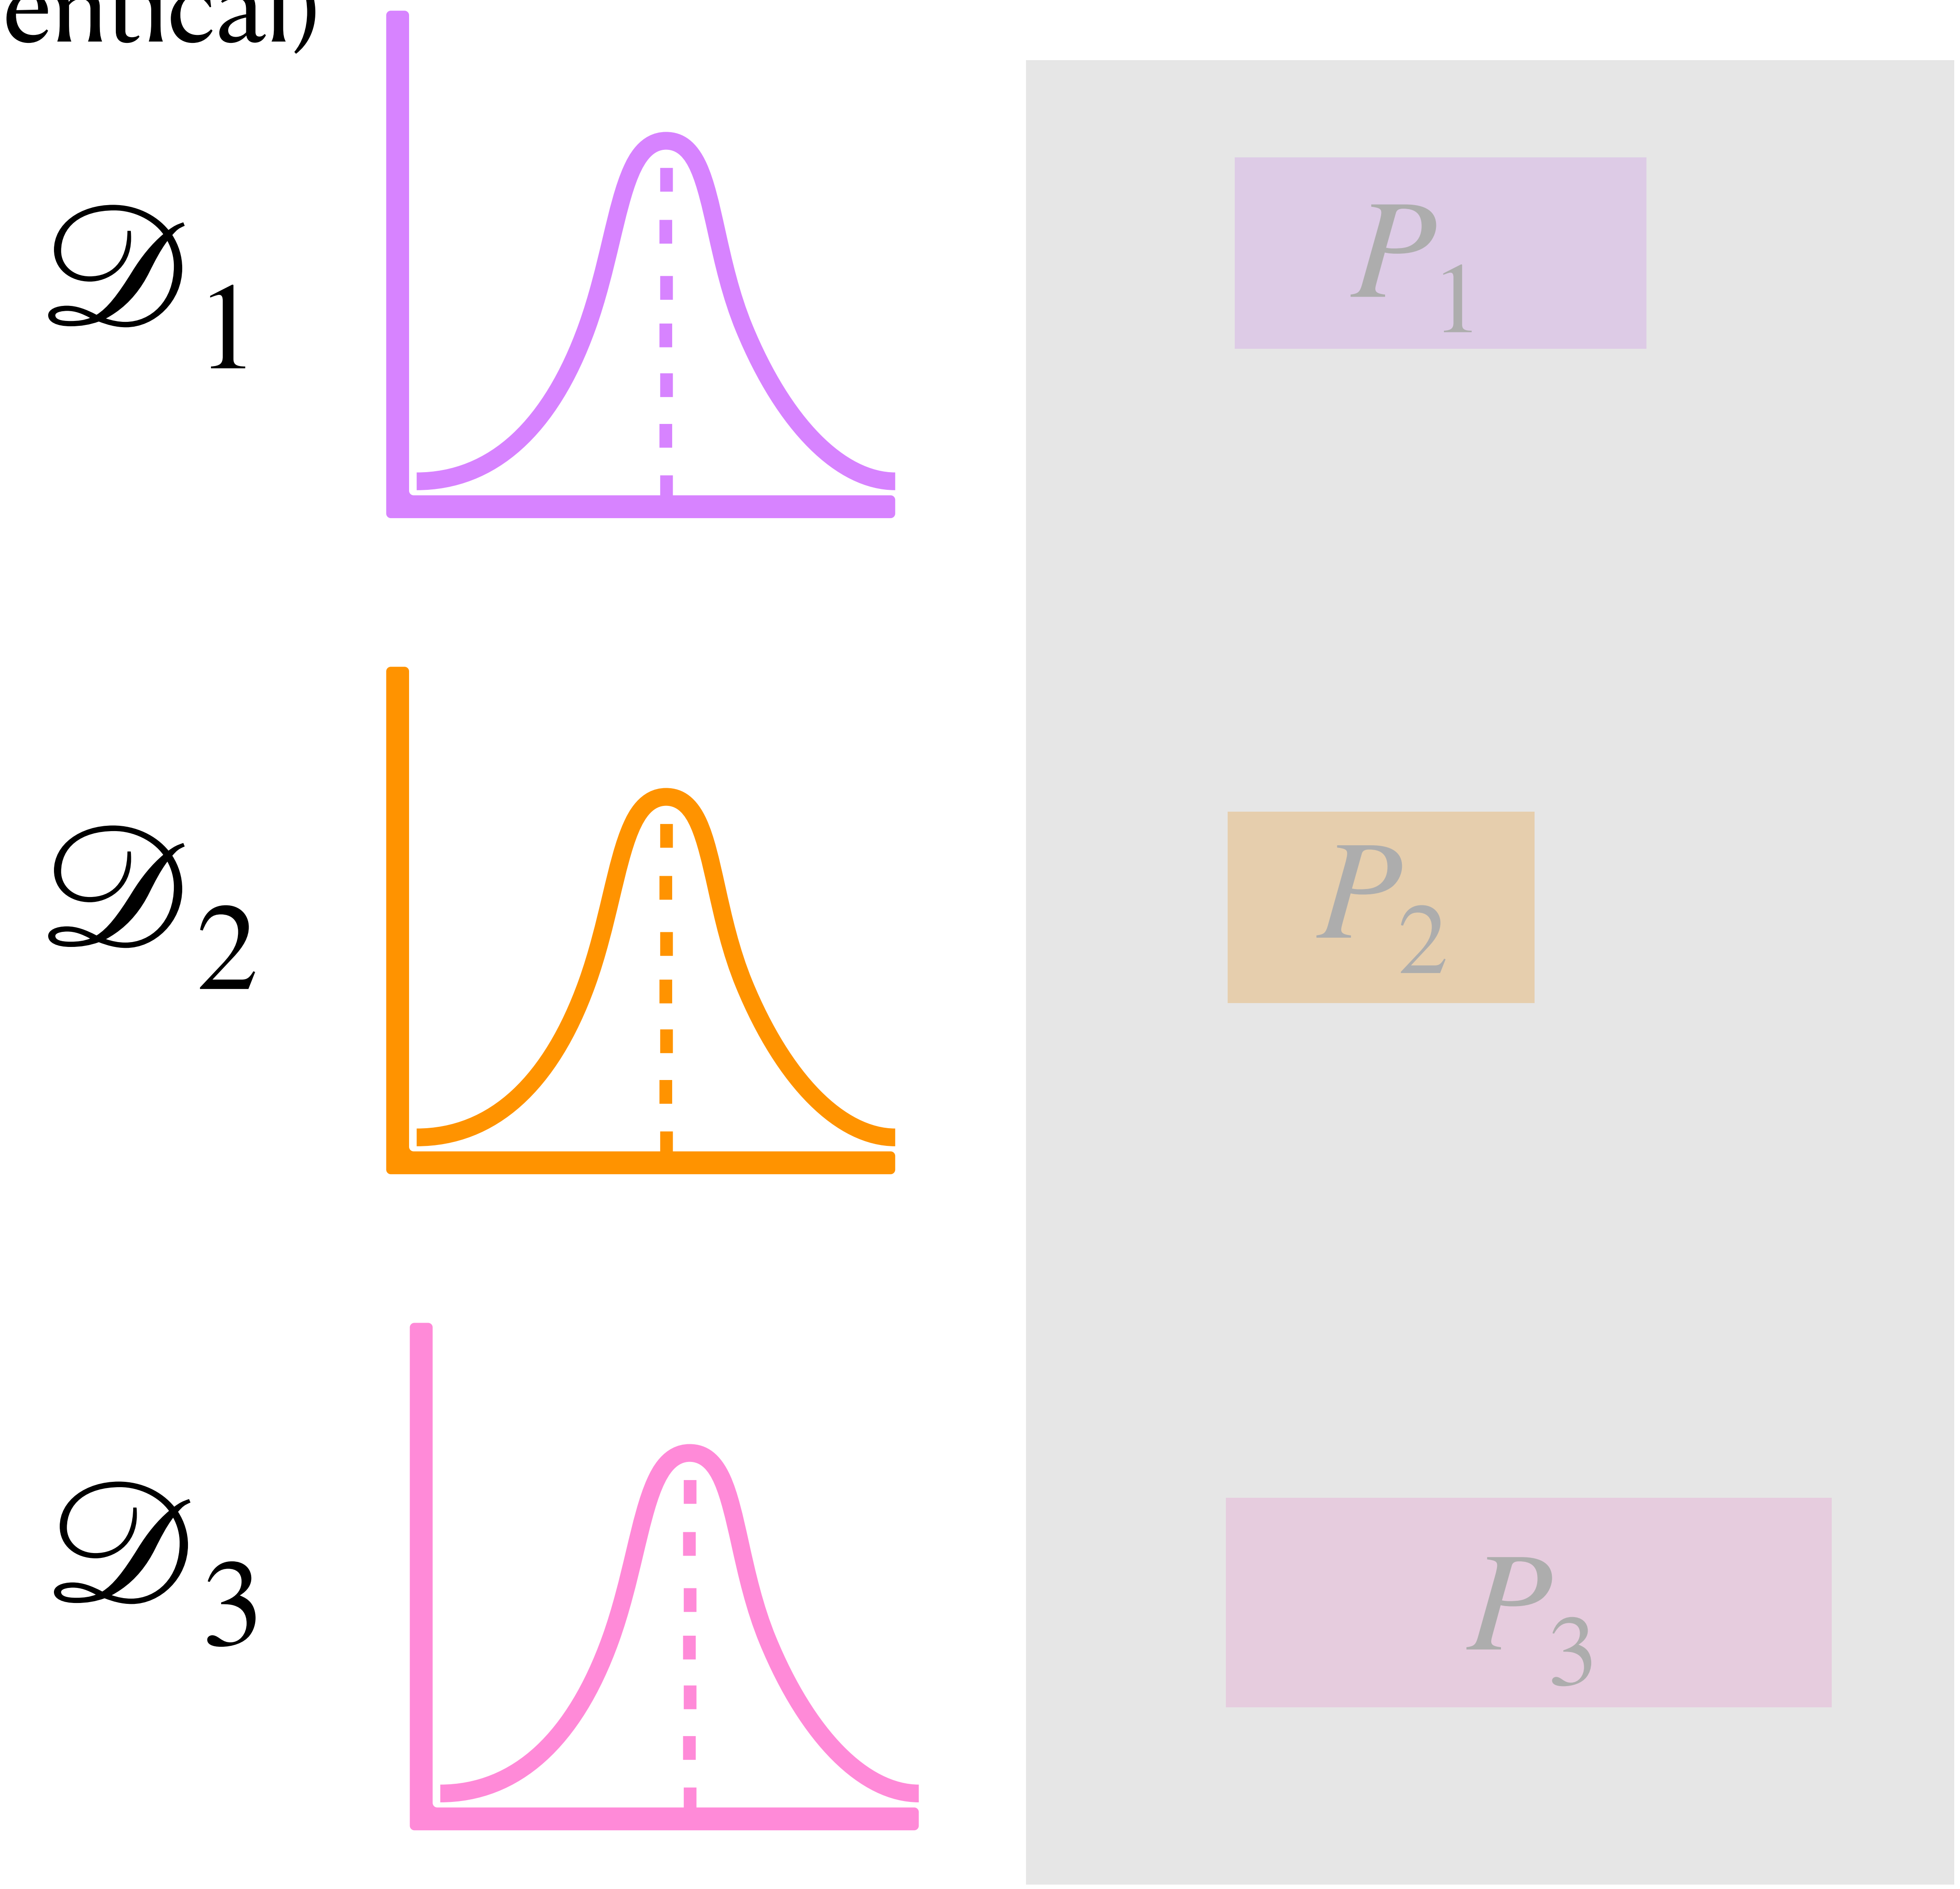

Goal: nonanticipatory policy minimizing

$$\mathbb{E}\left[\sum_{j=1}^{n} C_j\right]$$

- **Deterministic:** SPT
- **Decreasing hazard rates:** SERPT
- **Increasing hazard rates:** SEPT
- **In general:** Gittins Index Priority Policy

Optimal Policies

independent
(not necessarily
identical)

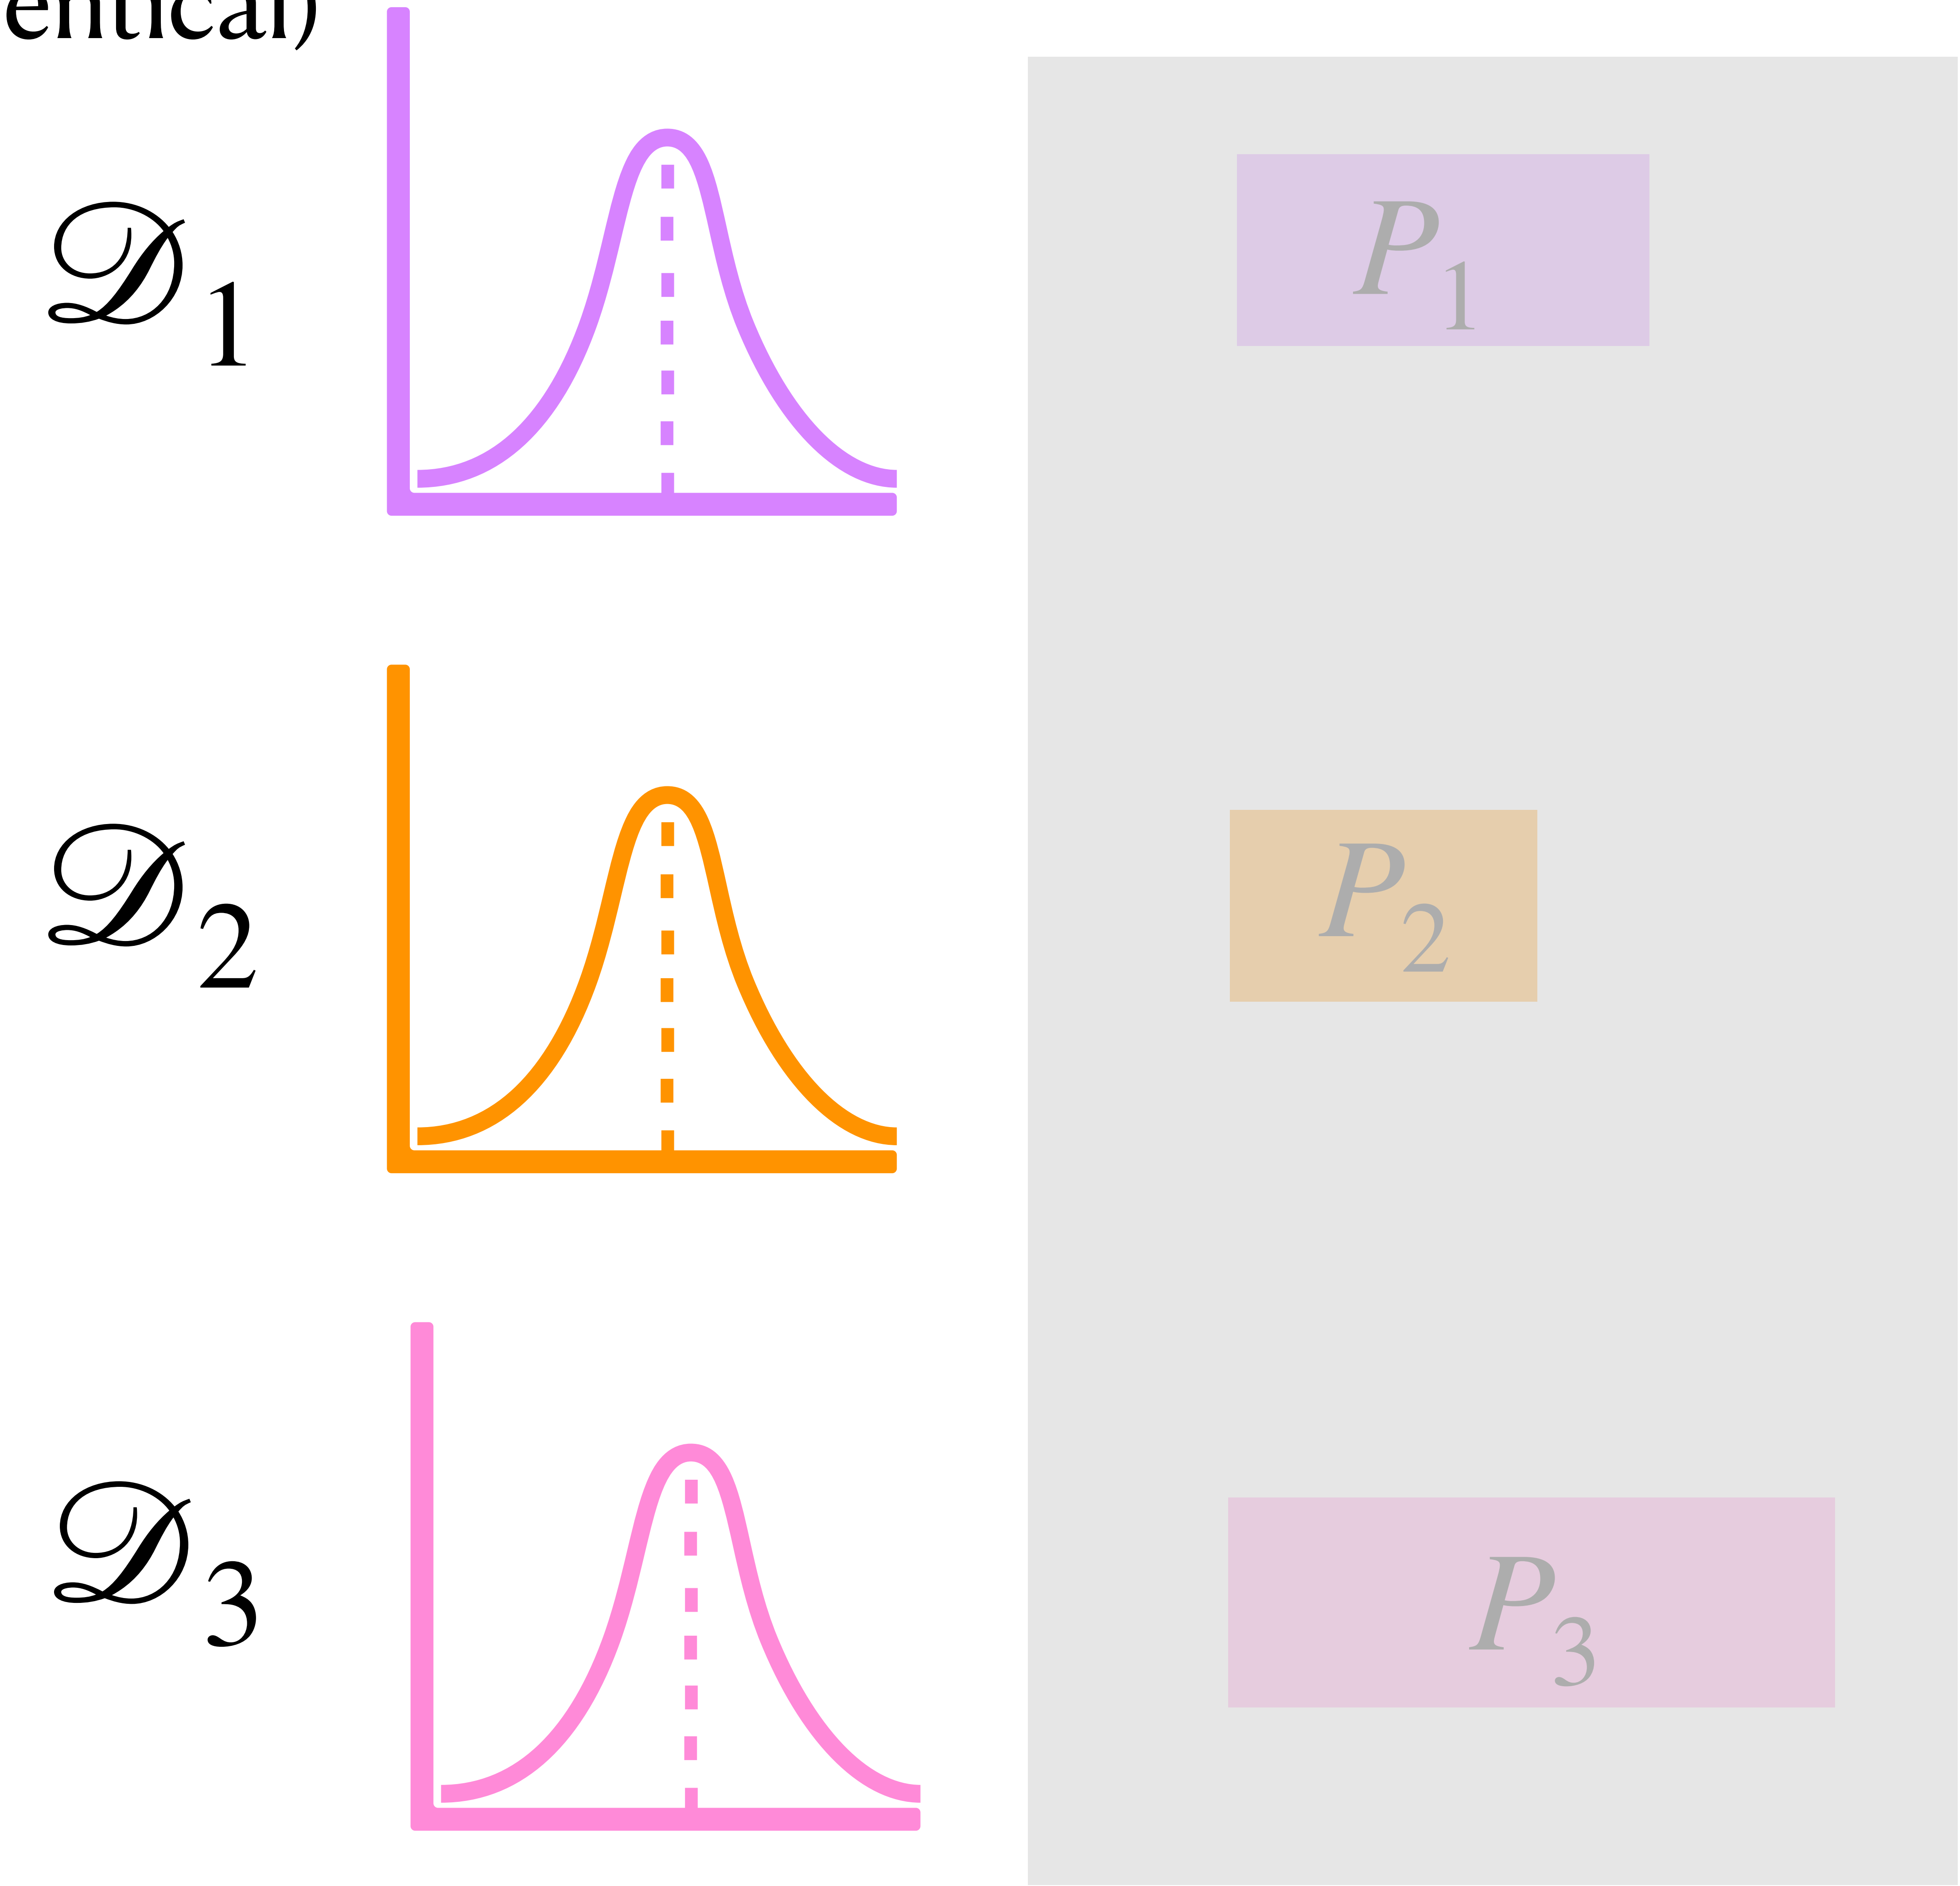

Goal: nonanticipatory policy minimizing

$$\mathbb{E}\left[\sum_{j=1}^{n} C_j\right]$$

- **Deterministic:** SPT
- **Decreasing hazard rates:** SERPT
- **Increasing hazard rates:** SEPT
- **In general:** Gittins Index Priority Policy

Optimal Policies

independent
(not necessarily
identical)

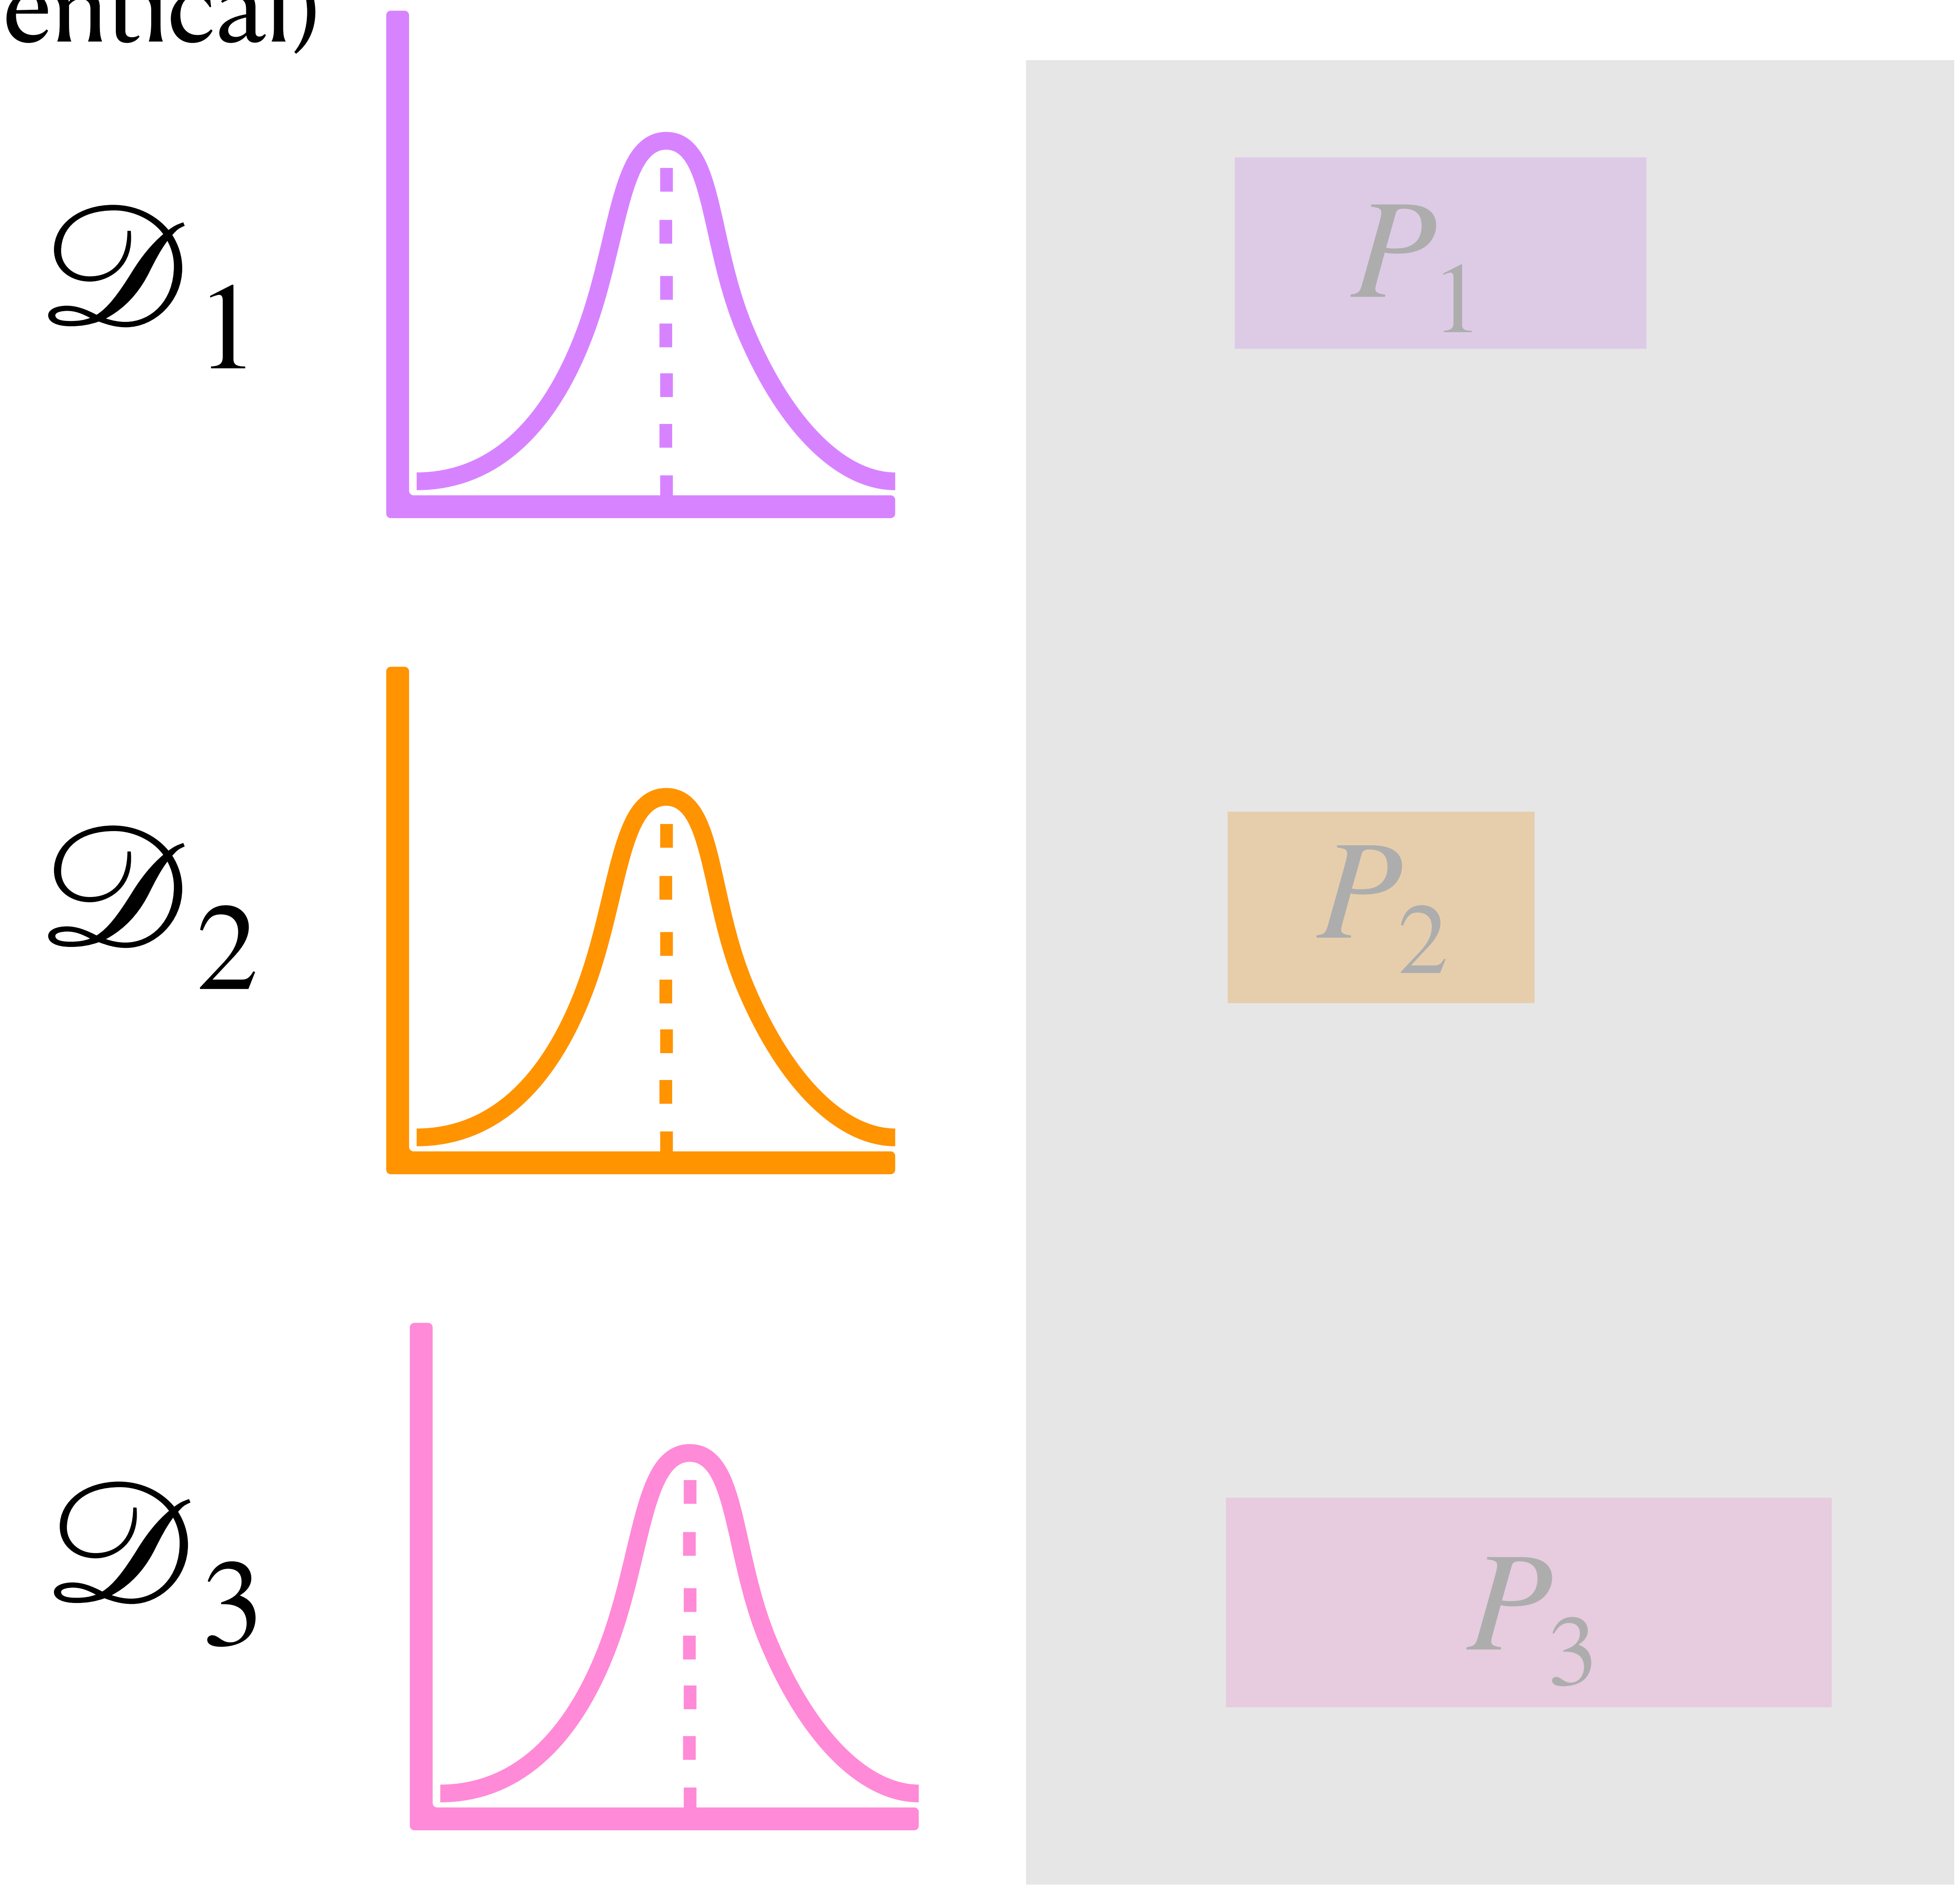

Goal: nonanticipatory policy minimizing

$$\mathbb{E}\left[\sum_{j=1}^{n} C_j\right]$$

- **Deterministic:** SPT
- **Decreasing hazard rates:** SERPT
- **Increasing hazard rates:** SEPT
- **In general:** Gittins Index Priority Policy

Optimal Policies

independent
(not necessarily identical)

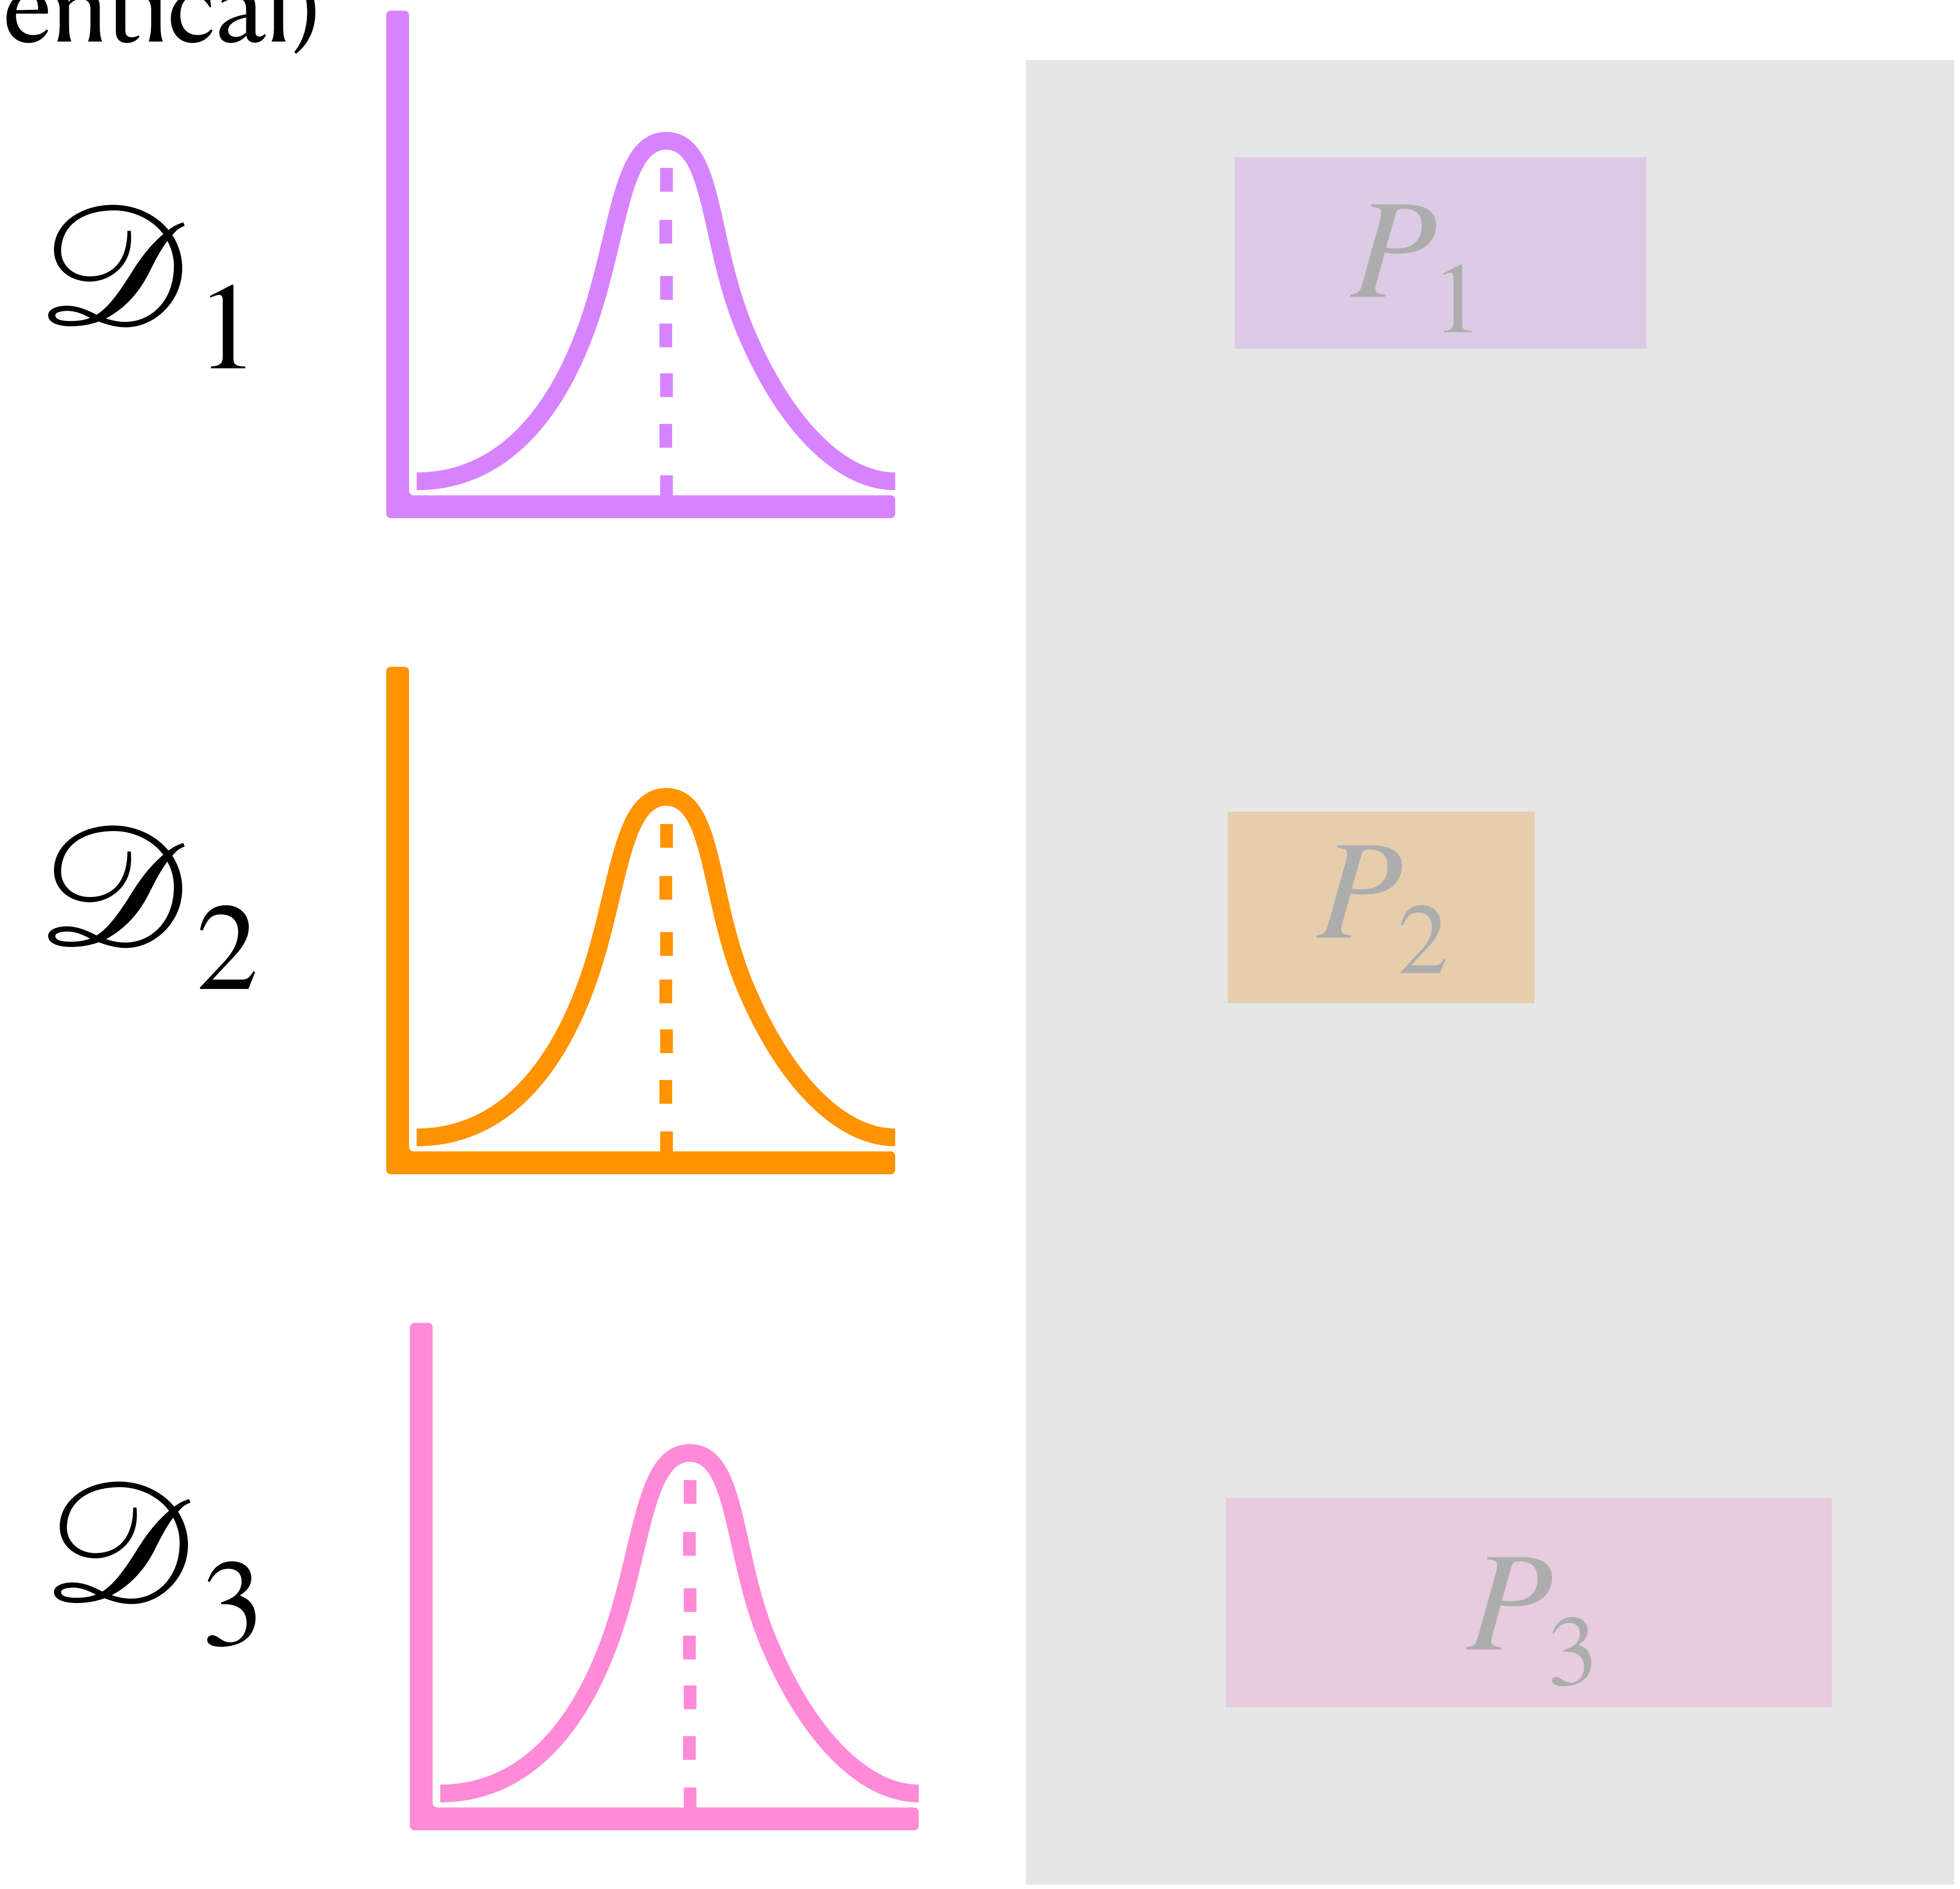

Goal: nonanticipatory policy minimizing

$$\mathbb{E}\left[\sum_{j=1}^{n} C_j\right]$$

- **Deterministic:** SPT
- **Decreasing hazard rates:** SERPT
- **Increasing hazard rates:** SEPT
- **In general:** Gittins Index Priority Policy

Policy Robustness

Goal: nonanticipatory policy minimizing

$$\mathbb{E}\left[\sum_{j=1}^{n} C_j\right]$$

- **Deterministic:** SPT
- **Decreasing hazard rates:** SERPT
- **Increasing hazard rates:** SEPT
- **In general:** Gittins Index Priority Policy

Policy Robustness

Robust: SEPT

Goal: nonanticipatory policy minimizing

$$\mathbb{E}\left[\sum_{j=1}^{n} C_j\right]$$

- **Deterministic:** SPT
- **Decreasing hazard rates:** SERPT
- **Increasing hazard rates:** SEPT
- **In general:** Gittins Index Priority Policy

Policy Robustness

Robust: SEPT

Why: cost expressed in terms of unconditional expected values

Goal: nonanticipatory policy minimizing

$$\mathbb{E}\left[\sum_{j=1}^{n} C_j\right]$$

- **Deterministic:** SPT
- **Decreasing hazard rates:** SERPT
- **Increasing hazard rates:** SEPT
- **In general:** Gittins Index Priority Policy

Policy Robustness

Robust: SEPT

Why: cost expressed in terms of unconditional expected values

$$\sum_{j=1}^{n} (n-j) \cdot \mathbb{E}[P_j]$$

Goal: nonanticipatory policy minimizing

$$\mathbb{E}\left[\sum_{j=1}^{n} C_j\right]$$

- **Deterministic:** SPT
- **Decreasing hazard rates:** SERPT
- **Increasing hazard rates:** SEPT
- **In general:** Gittins Index Priority Policy

Policy Robustness

Robust: SEPT

Why: cost expressed in terms of unconditional expected values

$$\sum_{j=1}^{n} (n-j) \cdot \mathbb{E}[P_j]$$

Not Robust: Gittins 🛑

Goal: nonanticipatory policy minimizing

$$\mathbb{E}\left[\sum_{j=1}^{n} C_j\right]$$

- **Deterministic:** SPT
- **Decreasing hazard rates:** SERPT
- **Increasing hazard rates:** SEPT
- **In general:** Gittins Index Priority Policy

Policy Robustness

Robust: SEPT

Why: cost expressed in terms of unconditional expected values

$$\sum_{j=1}^{n} (n-j) \cdot \mathbb{E}[P_j]$$

Not Robust: Gittins 🛑

Why: policy based on conditional expectations and probabilities that are highly sensitive to error

Goal: nonanticipatory policy minimizing

$$\mathbb{E}\left[\sum_{j=1}^{n} C_j\right]$$

- **Deterministic:** SPT
- **Decreasing hazard rates:** SERPT
- **Increasing hazard rates:** SEPT
- **In general:** Gittins Index Priority Policy

Gittins Index Policy

Robust: SEPT

Why: cost expressed in terms of unconditional expected values

$$\sum_{j=1}^{n} (n-j) \cdot \mathbb{E}[P_j]$$

Not Robust: Gittins

Why: policy based on conditional expectations and probabilities that are highly sensitive to error

Gittins Index Policy

finite
support

Robust: SEPT

Why: cost expressed in terms of unconditional expected values

$$\sum_{j=1}^{n} (n-j) \cdot \mathbb{E}[P_j]$$

Not Robust: Gittins

Why: policy based on conditional expectations and probabilities that are highly sensitive to error

Gittins Index Policy

finite support

Robust: SEPT

Why: cost expressed in terms of unconditional expected values

$$\sum_{j=1}^{n} (n-j) \cdot \mathbb{E}[P_j]$$

Not Robust: Gittins

Why: policy based on conditional expectations and probabilities that are highly sensitive to error

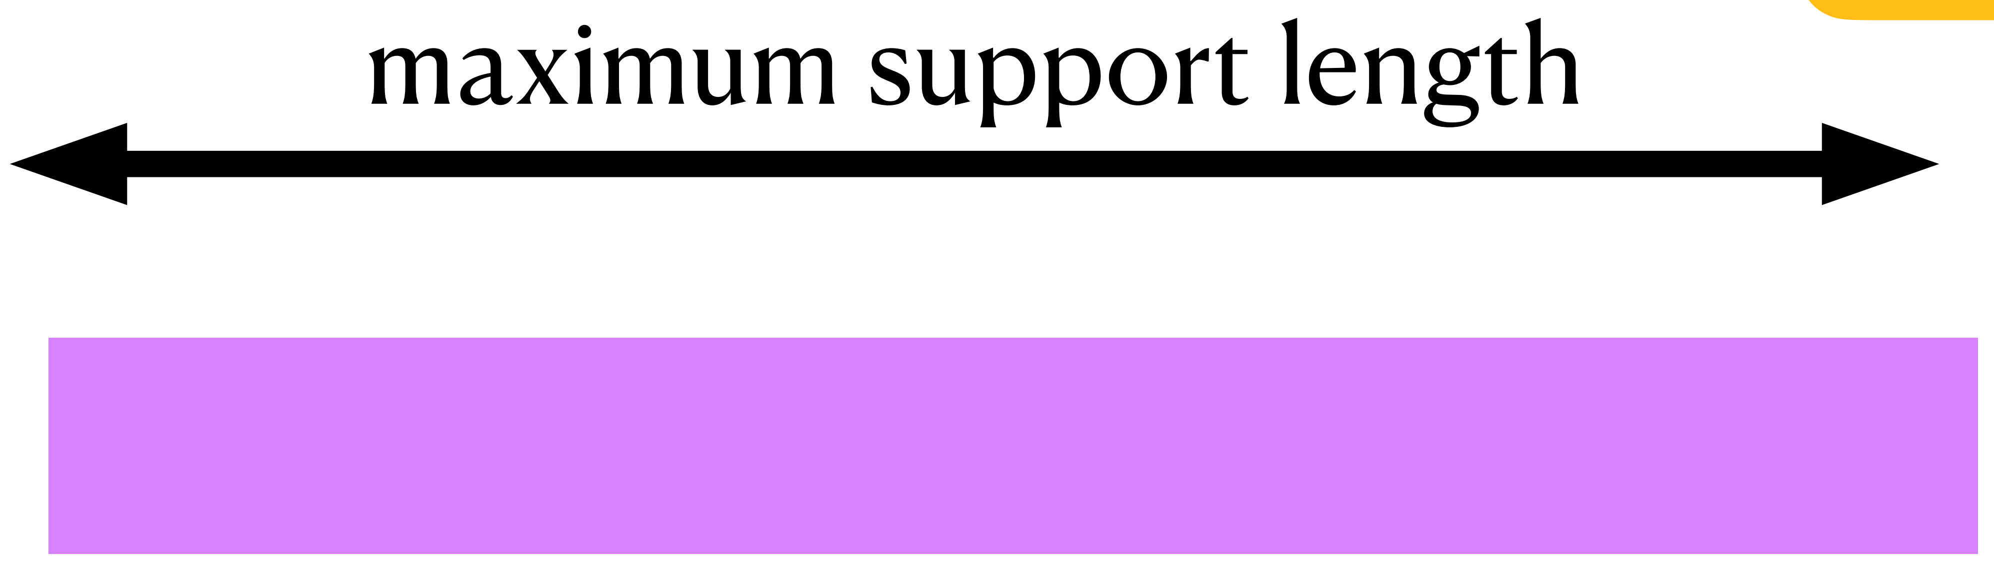

maximum support length

Gittins Index Policy

finite support

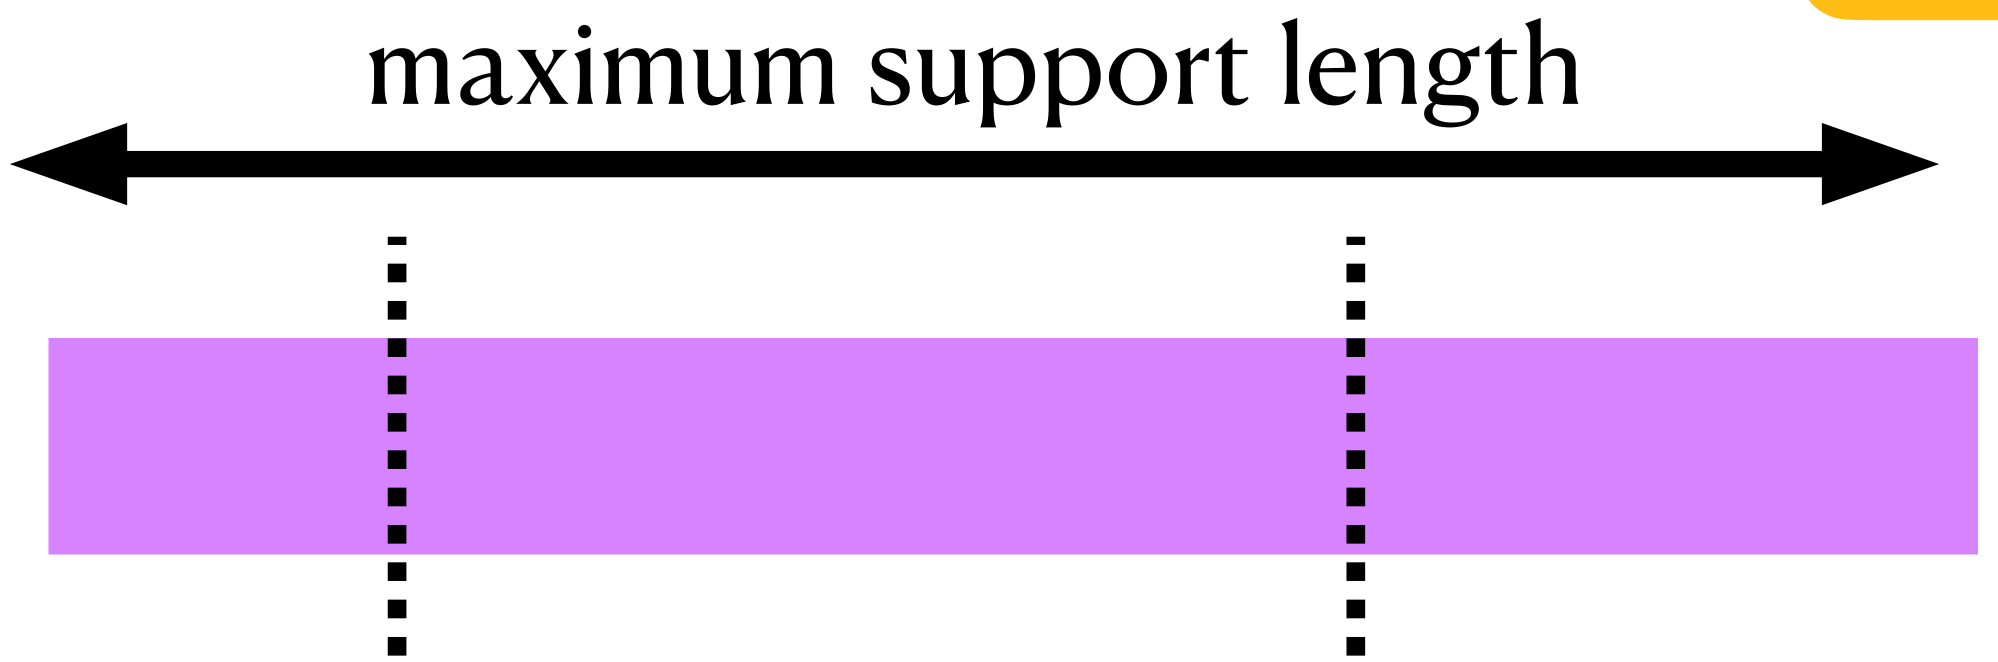

Robust: SEPT

Why: cost expressed in terms of unconditional expected values

$$\sum_{j=1}^{n} (n-j) \cdot \mathbb{E}[P_j]$$

Not Robust: Gittins 🛑

Why: policy based on conditional expectations and probabilities that are highly sensitive to error

Gittins Index Policy

finite
support

Robust: SEPT

Why: cost expressed in terms of unconditional expected values

$$\sum_{j=1}^{n} (n - j) \cdot \mathbb{E}[P_j]$$

Not Robust: Gittins

Why: policy based on conditional expectations and probabilities that are highly sensitive to error

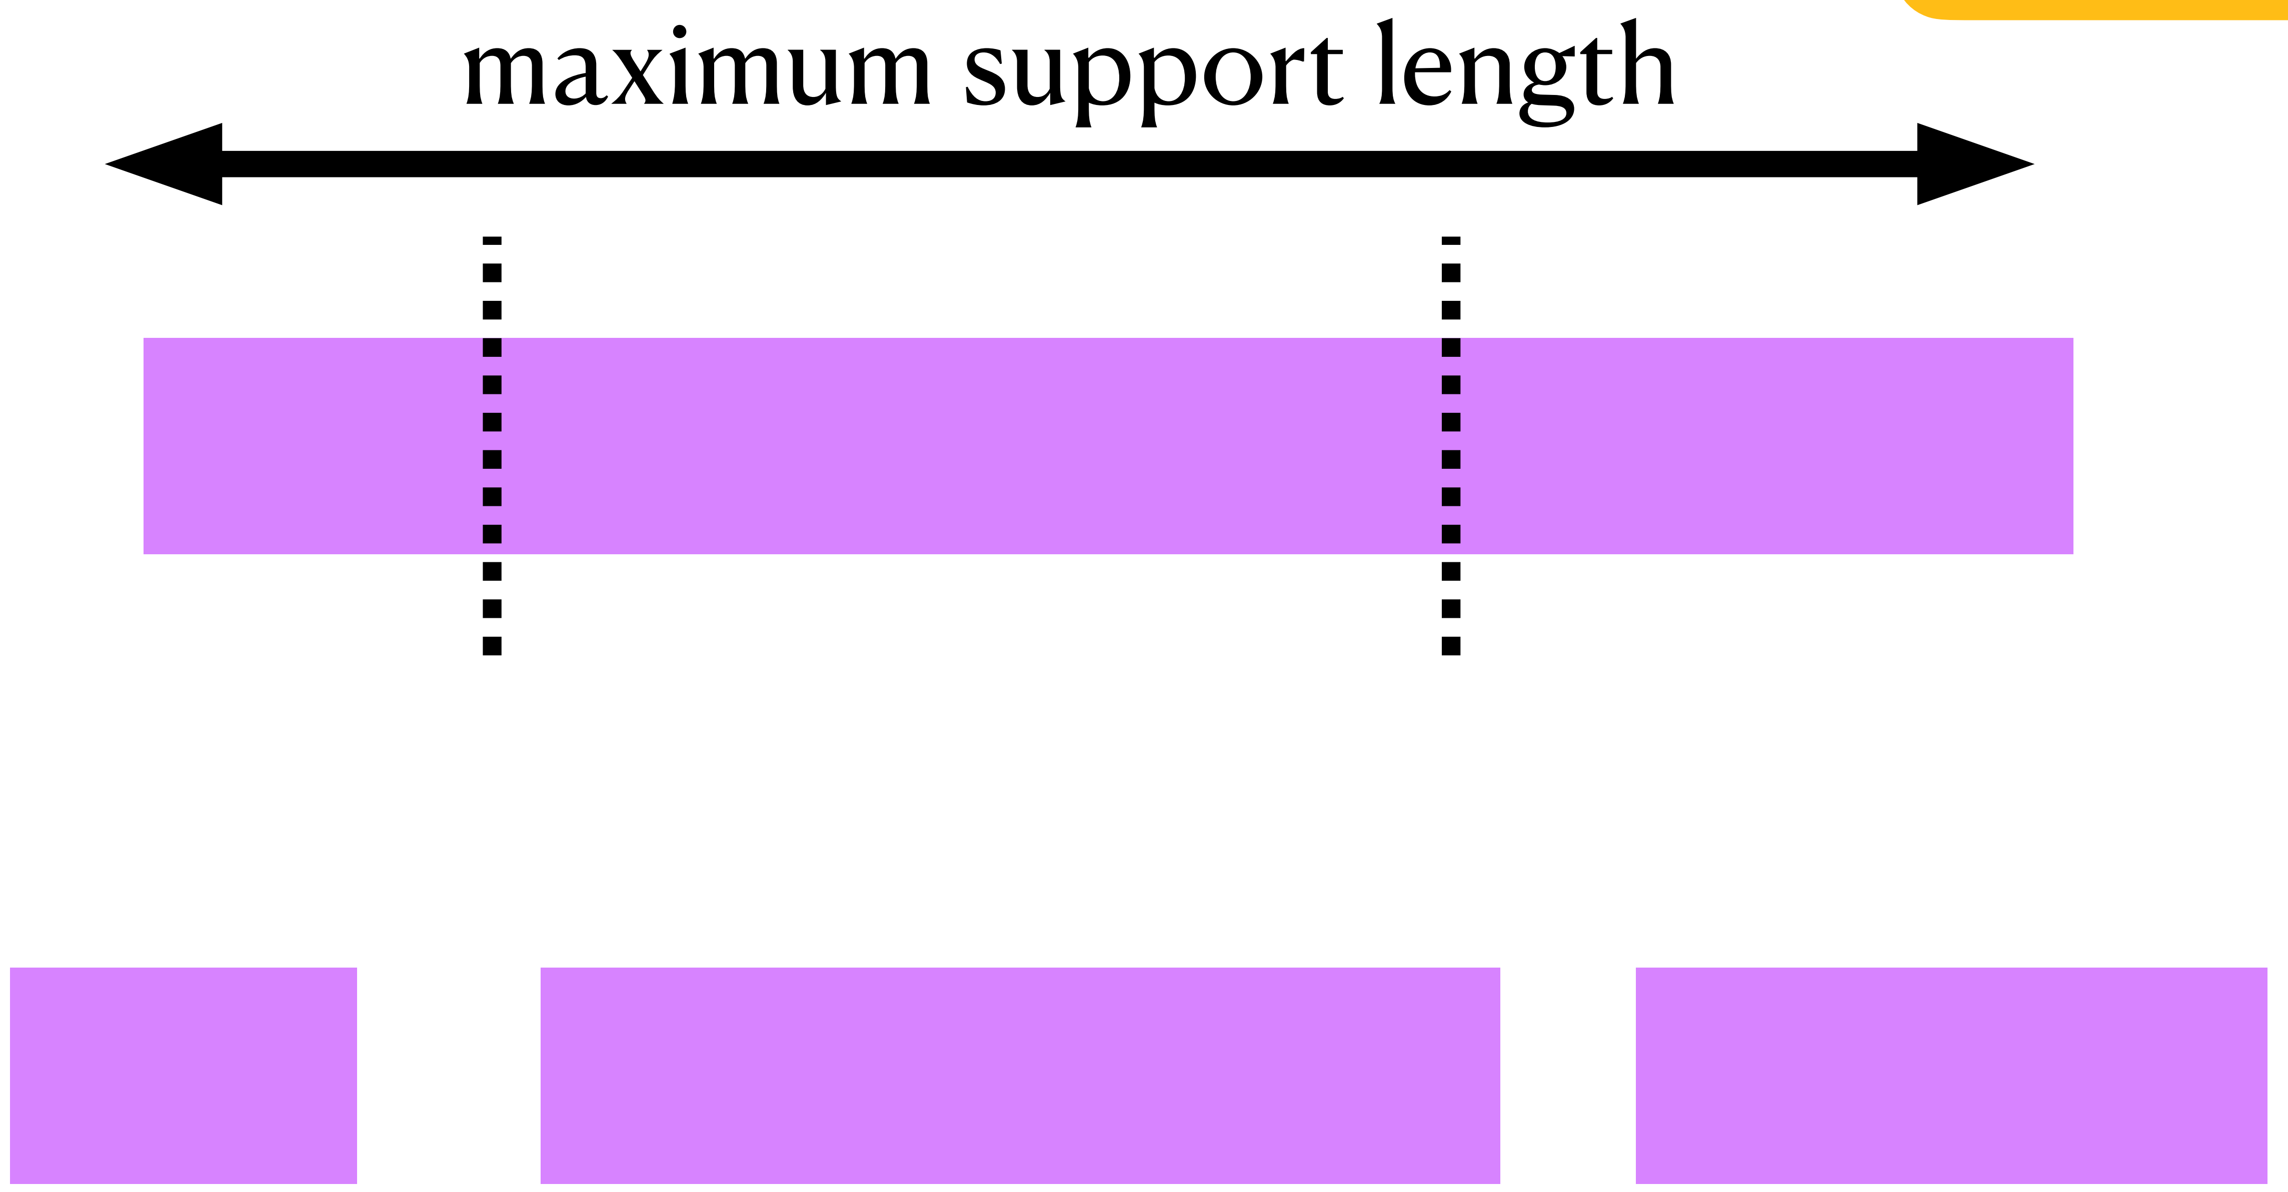

Gittins Index Policy

finite support

Robust: SEPT

Why: cost expressed in terms of unconditional expected values

$$\sum_{j=1}^{n} (n-j) \cdot \mathbb{E}[P_j]$$

Not Robust: Gittins

Why: policy based on conditional expectations and probabilities that are highly sensitive to error

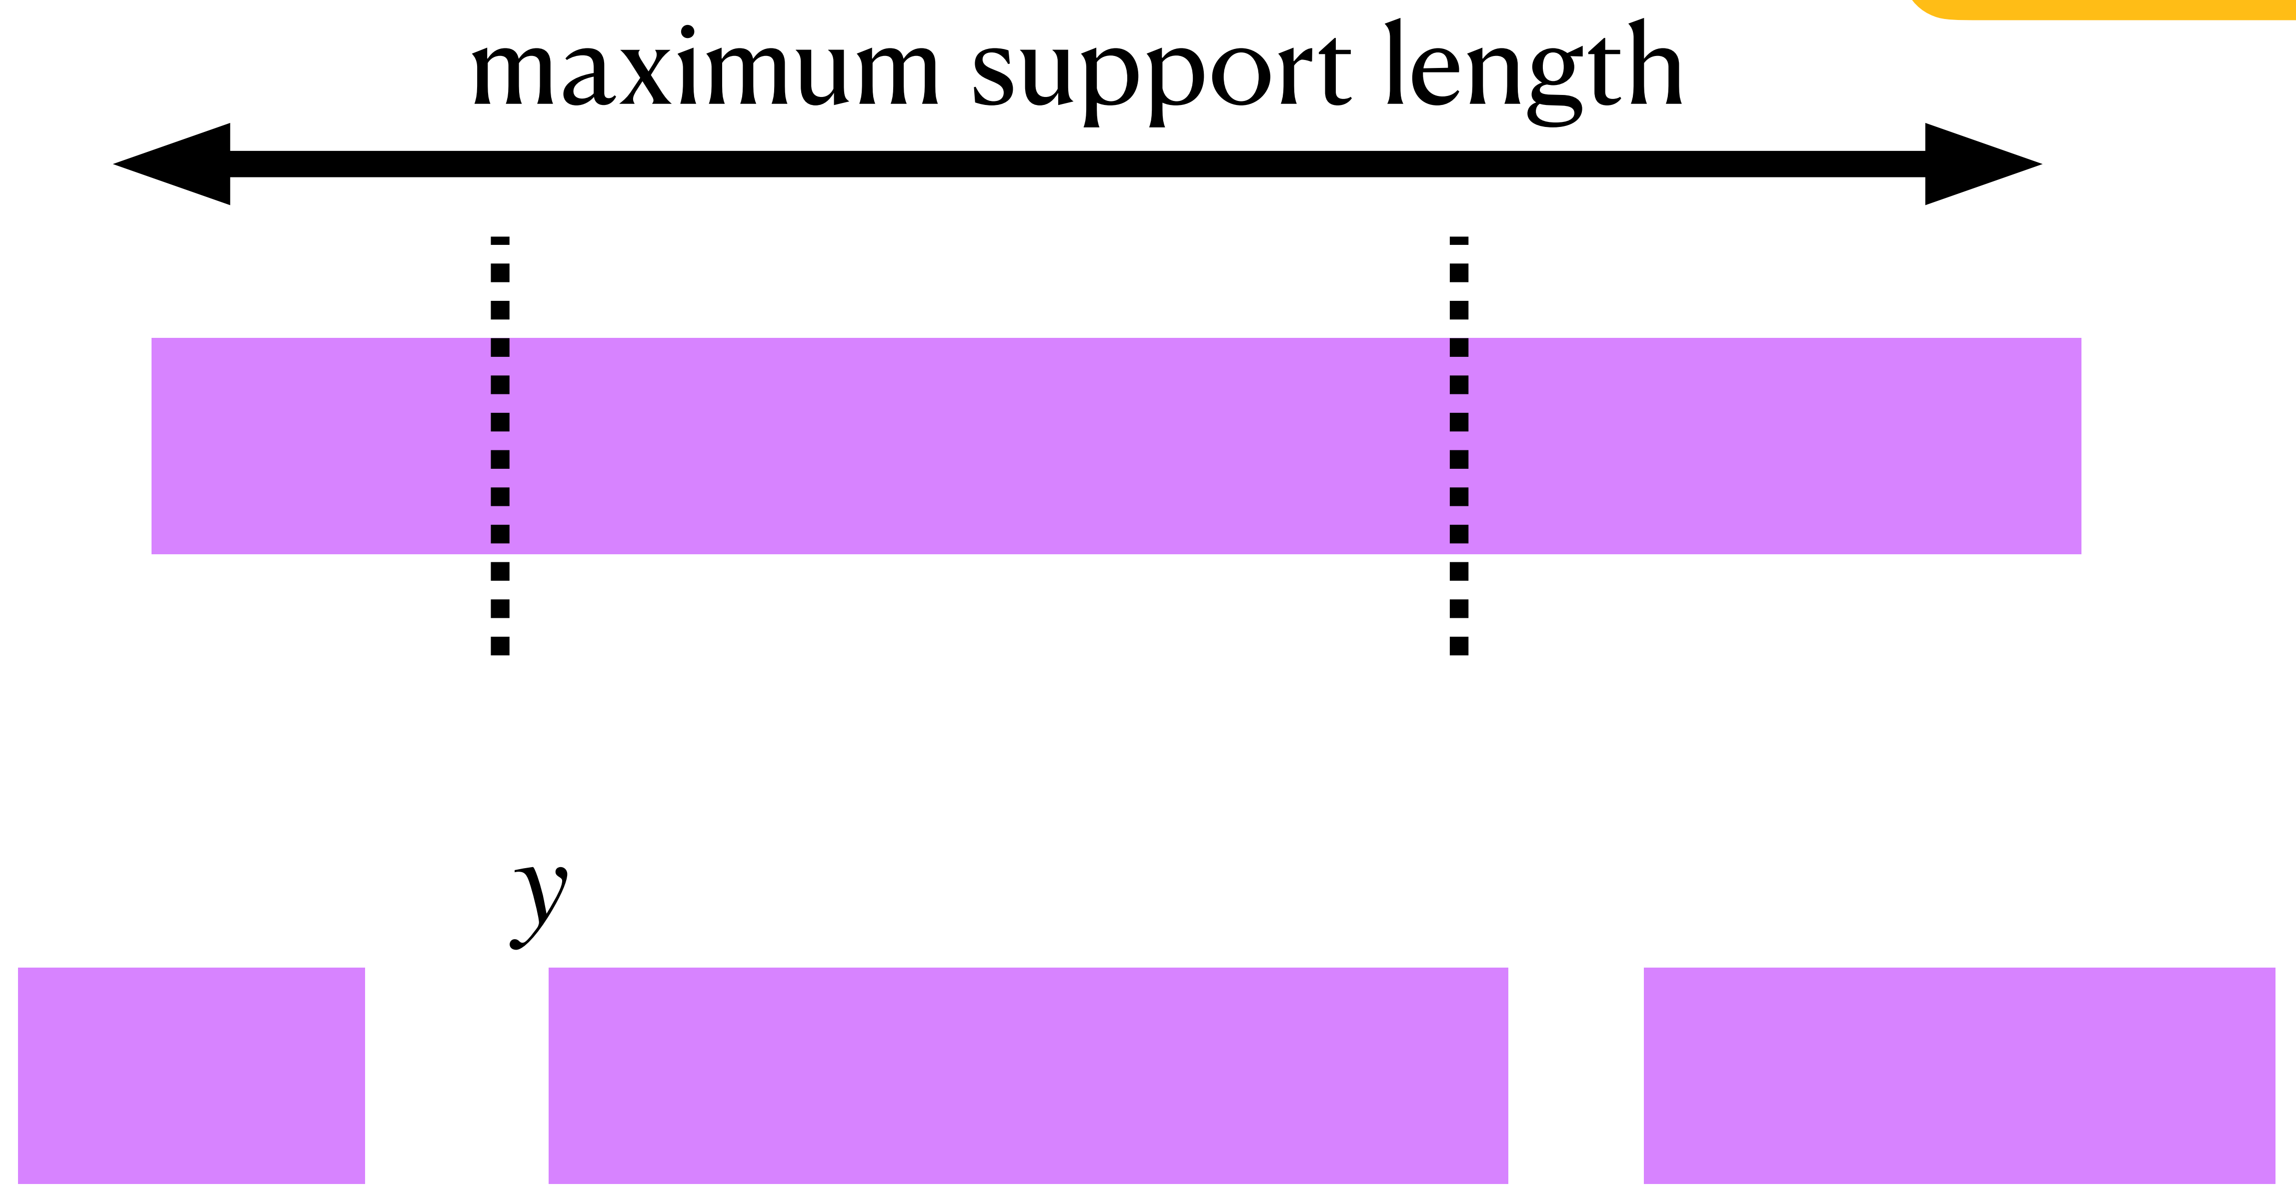

Gittins Index Policy

finite support

Robust: SEPT

Why: cost expressed in terms of unconditional expected values

$$\sum_{j=1}^{n} (n-j) \cdot \mathbb{E}[P_j]$$

Not Robust: Gittins

Why: policy based on conditional expectations and probabilities that are highly sensitive to error

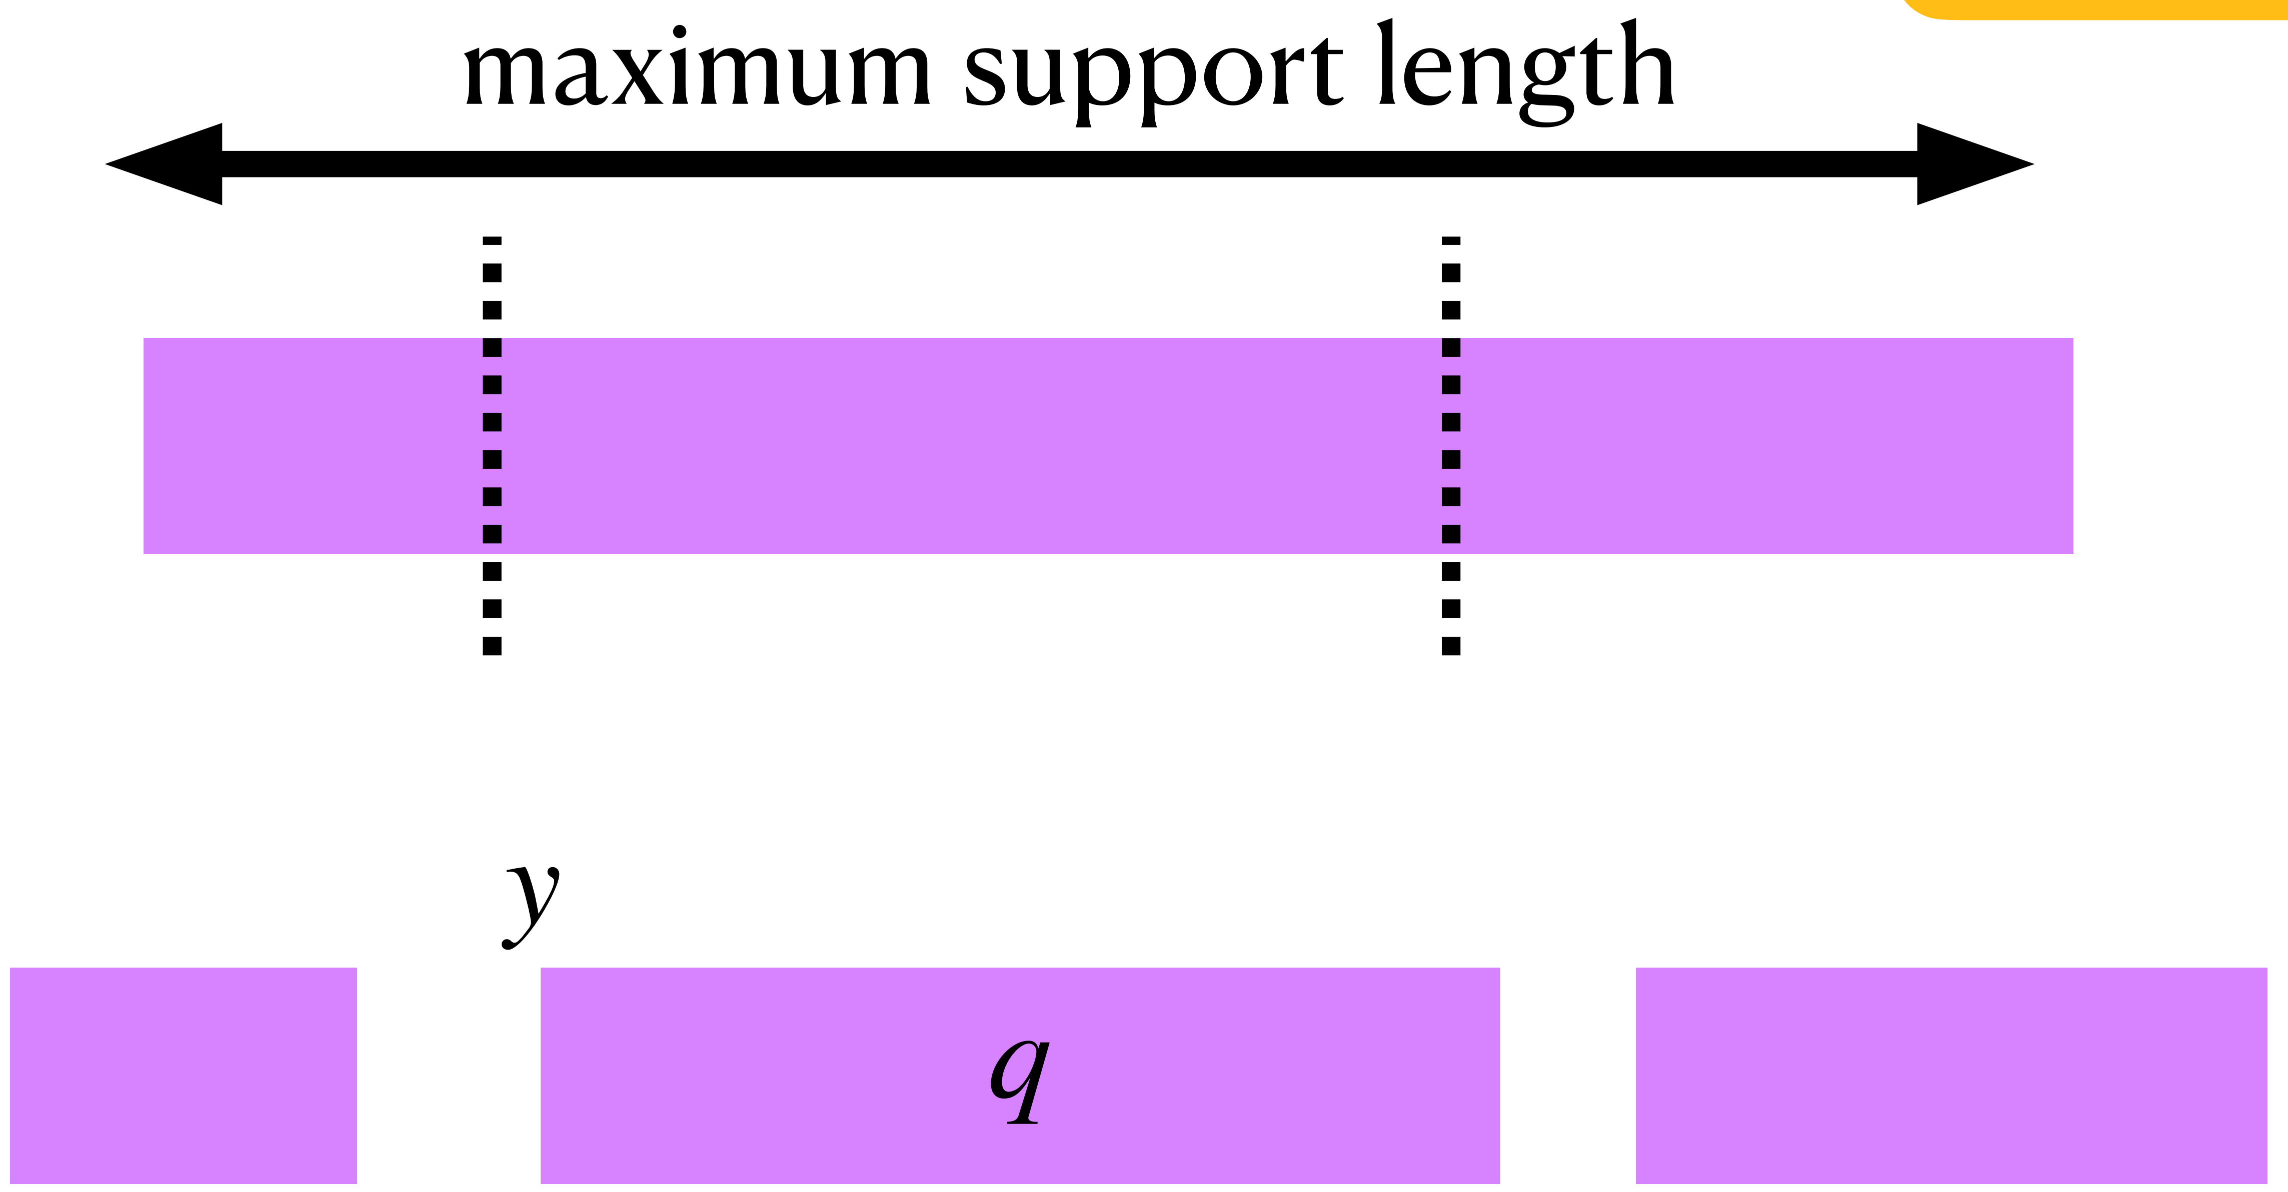

Gittins Index Policy

finite support

Robust: SEPT

Why: cost expressed in terms of unconditional expected values

$$\sum_{j=1}^{n} (n-j) \cdot \mathbb{E}[P_j]$$

Not Robust: Gittins

Why: policy based on conditional expectations and probabilities that are highly sensitive to error

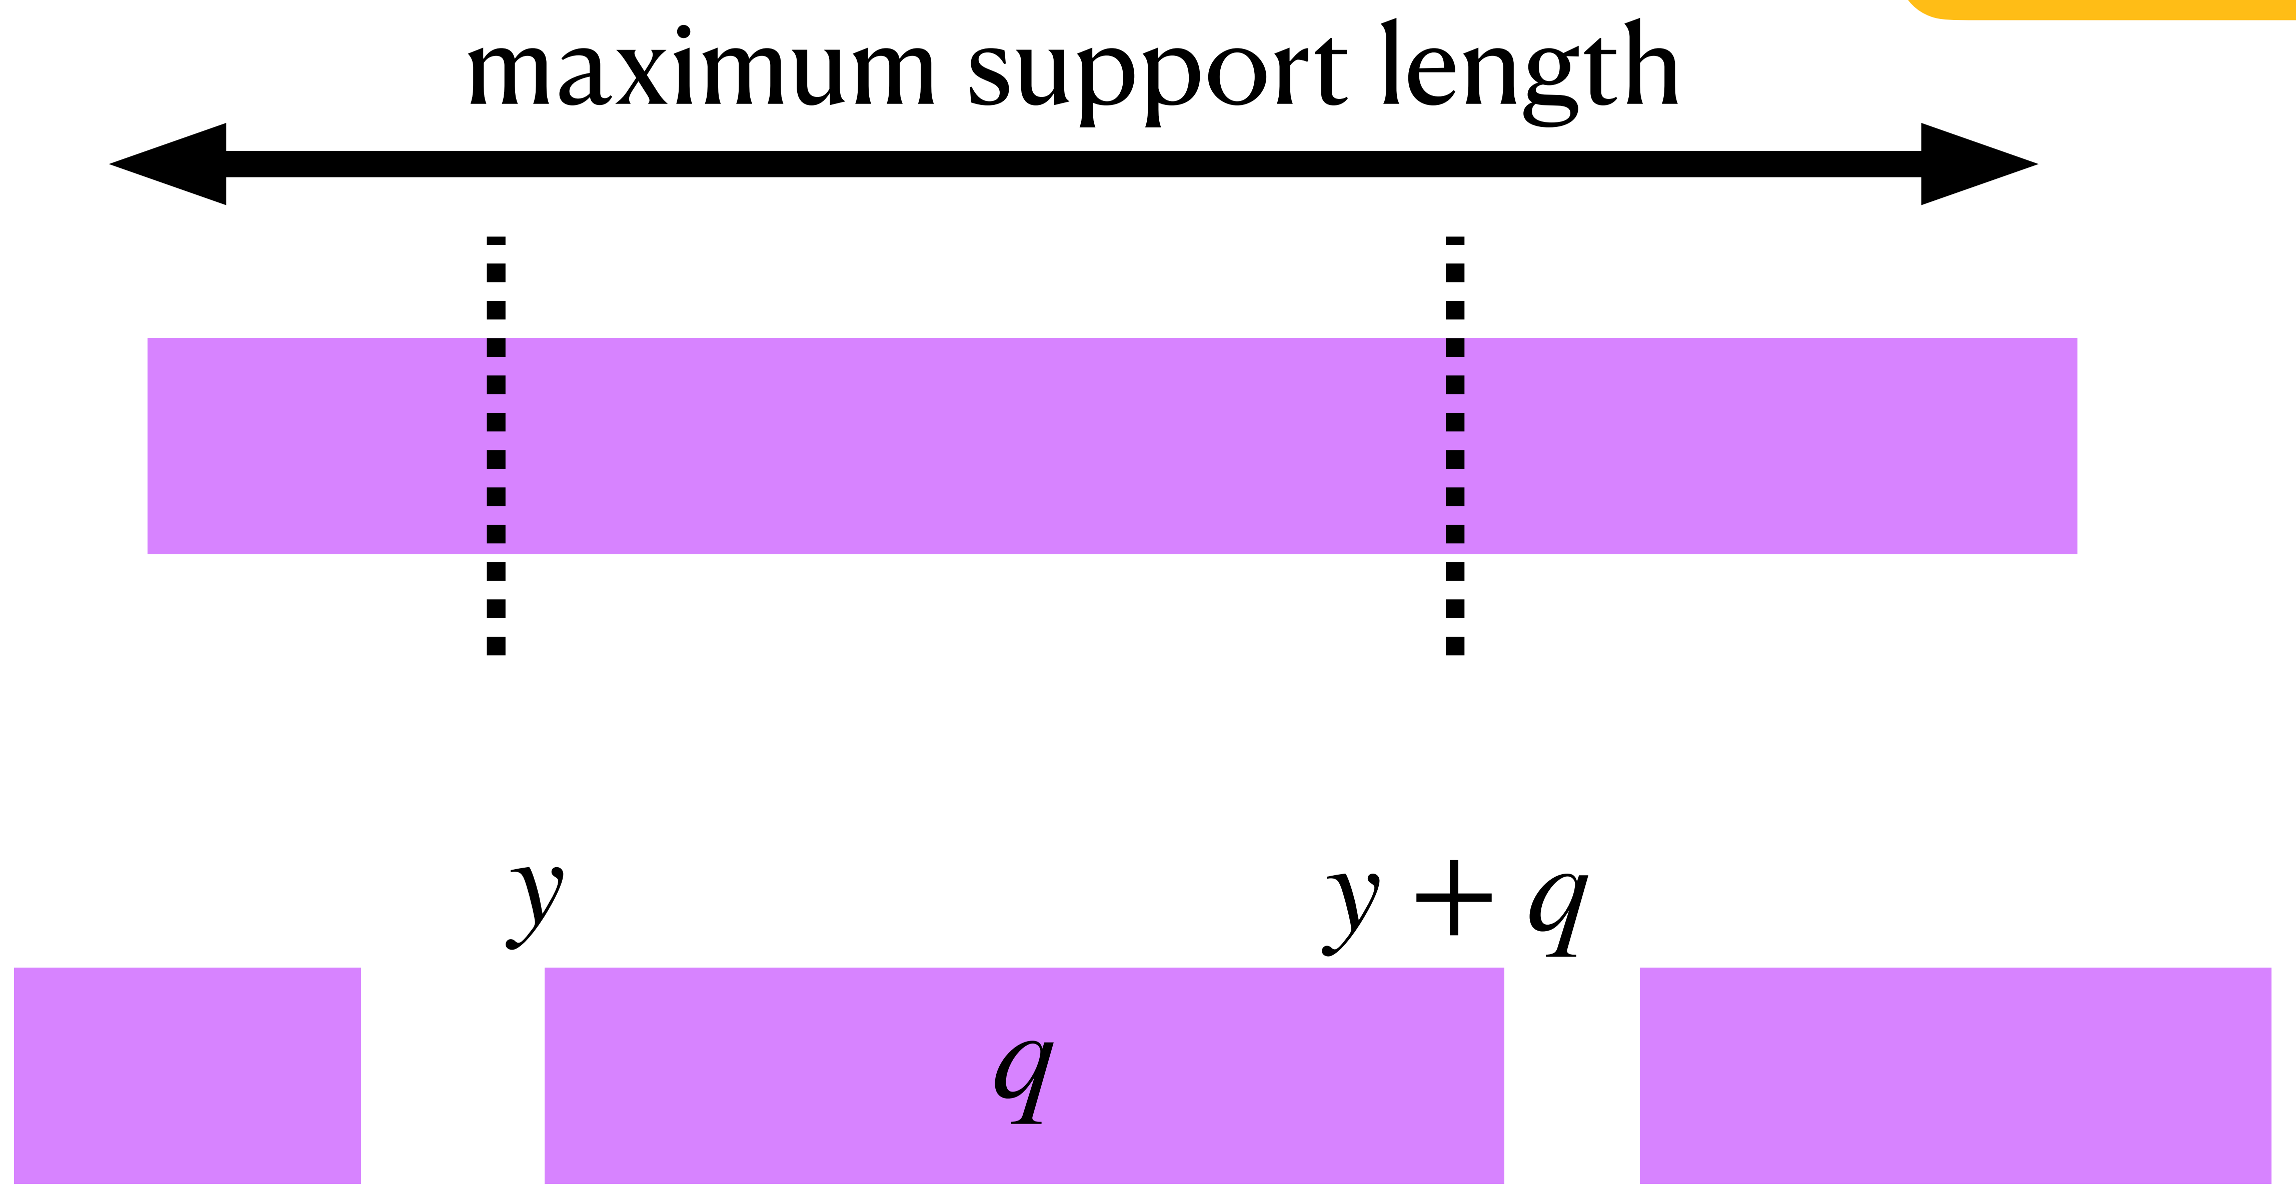

Gittins Index Policy

finite support

Robust: SEPT

Why: cost expressed in terms of unconditional expected values

$$\sum_{j=1}^{n} (n-j) \cdot \mathbb{E}[P_j]$$

Not Robust: Gittins

Why: policy based on conditional expectations and probabilities that are highly sensitive to error

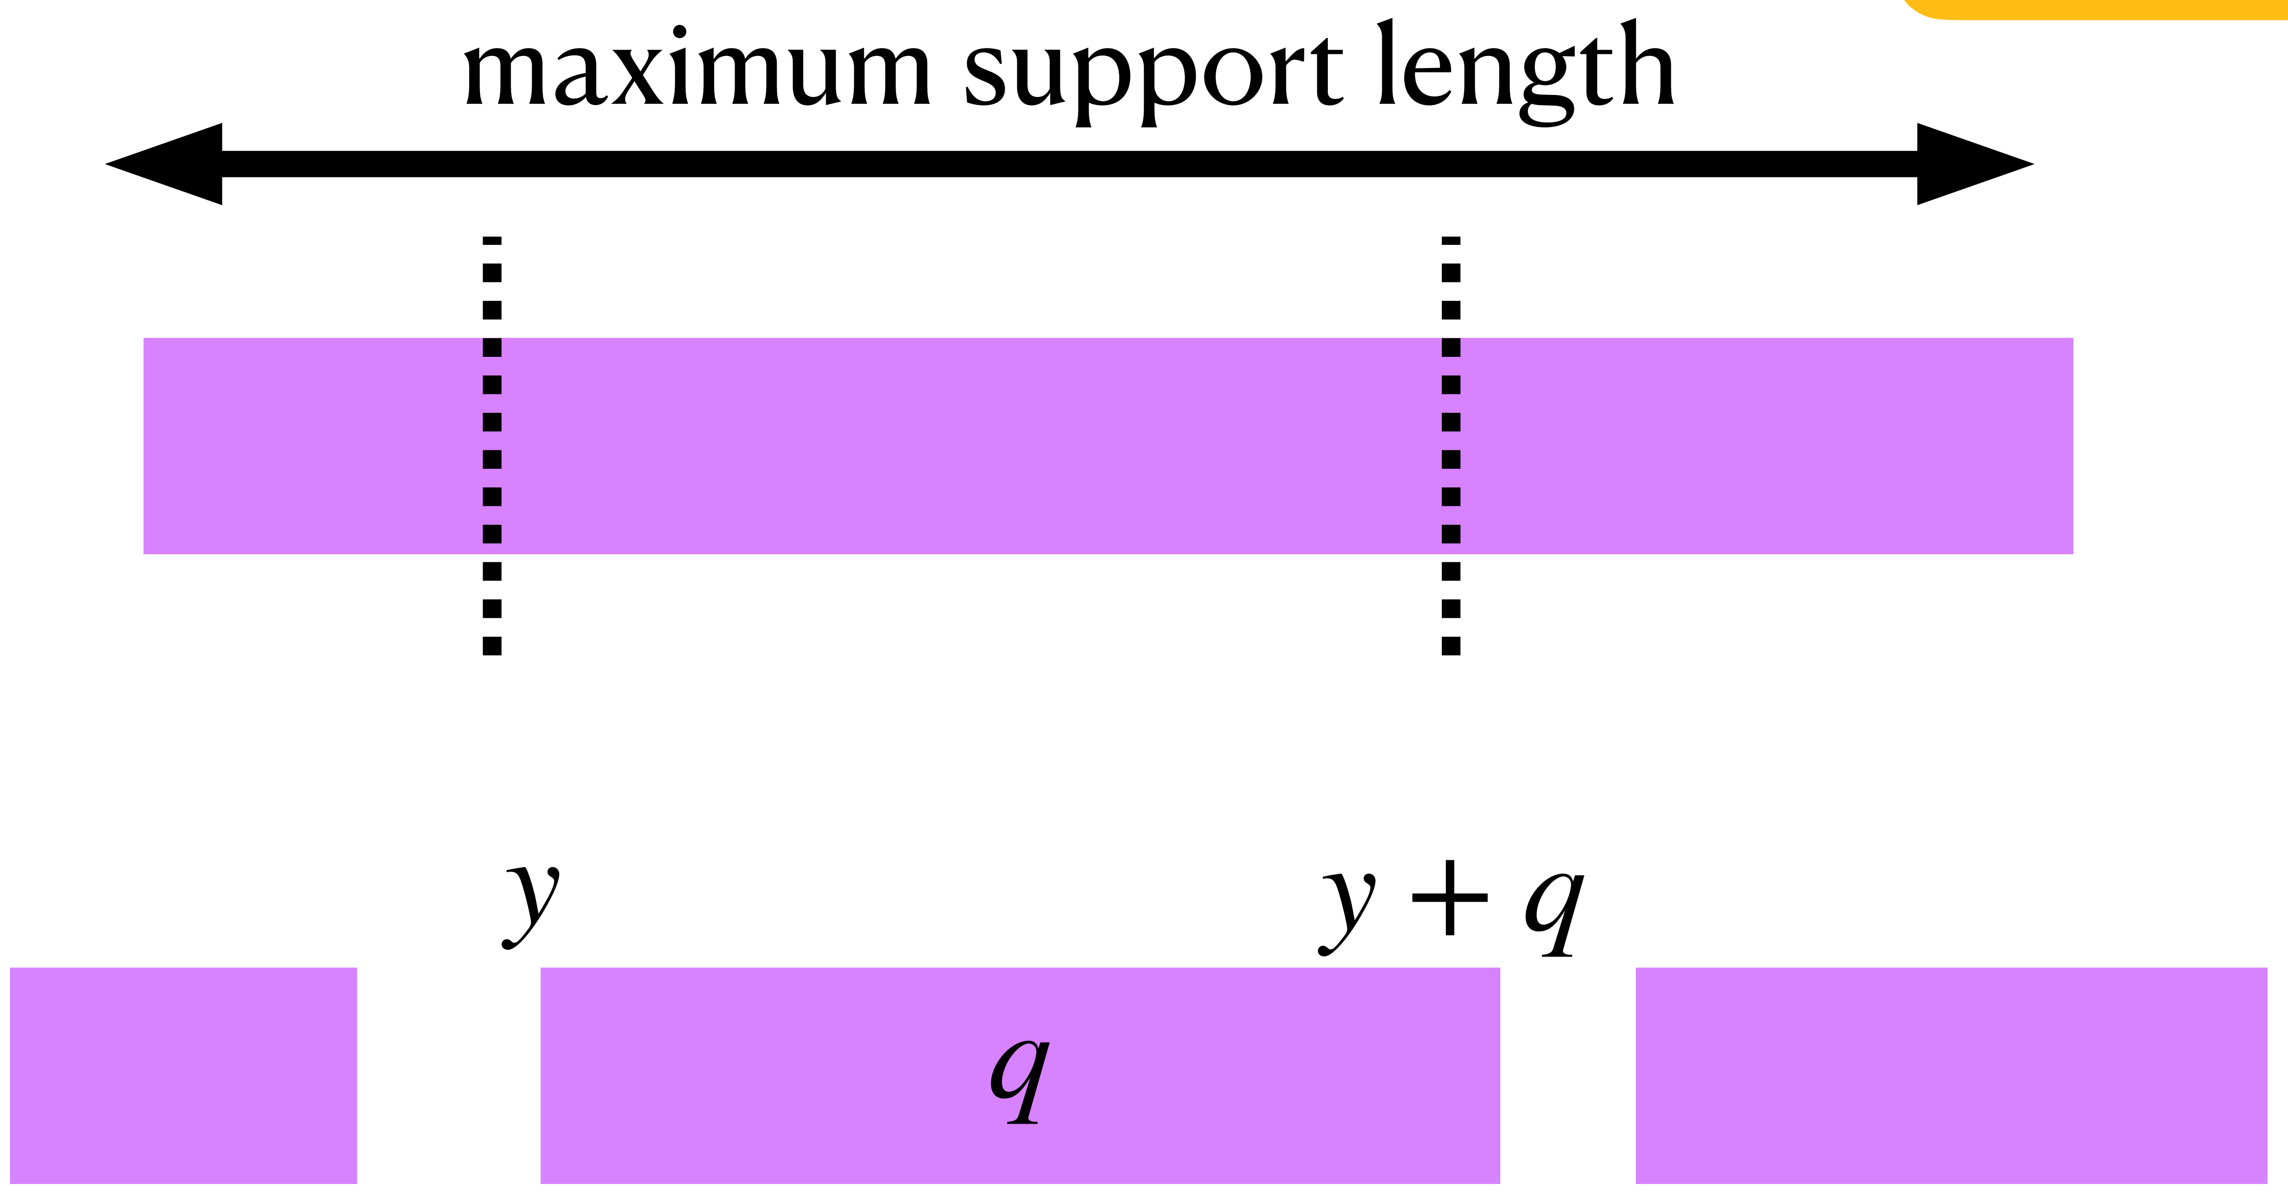

$$\max_{q \geq 0} \frac{\mathbb{P}(P_j - y \leq q \mid P_j > y)}{\mathbb{E}[\min\{P_j - y, q\} \mid P_j > y]}$$

Gittins Index Policy

finite support

Robust: SEPT

Why: cost expressed in terms of unconditional expected values

$$\sum_{j=1}^{n} (n-j) \cdot \mathbb{E}[P_j]$$

Not Robust: Gittins

Why: policy based on conditional expectations and probabilities that are highly sensitive to error

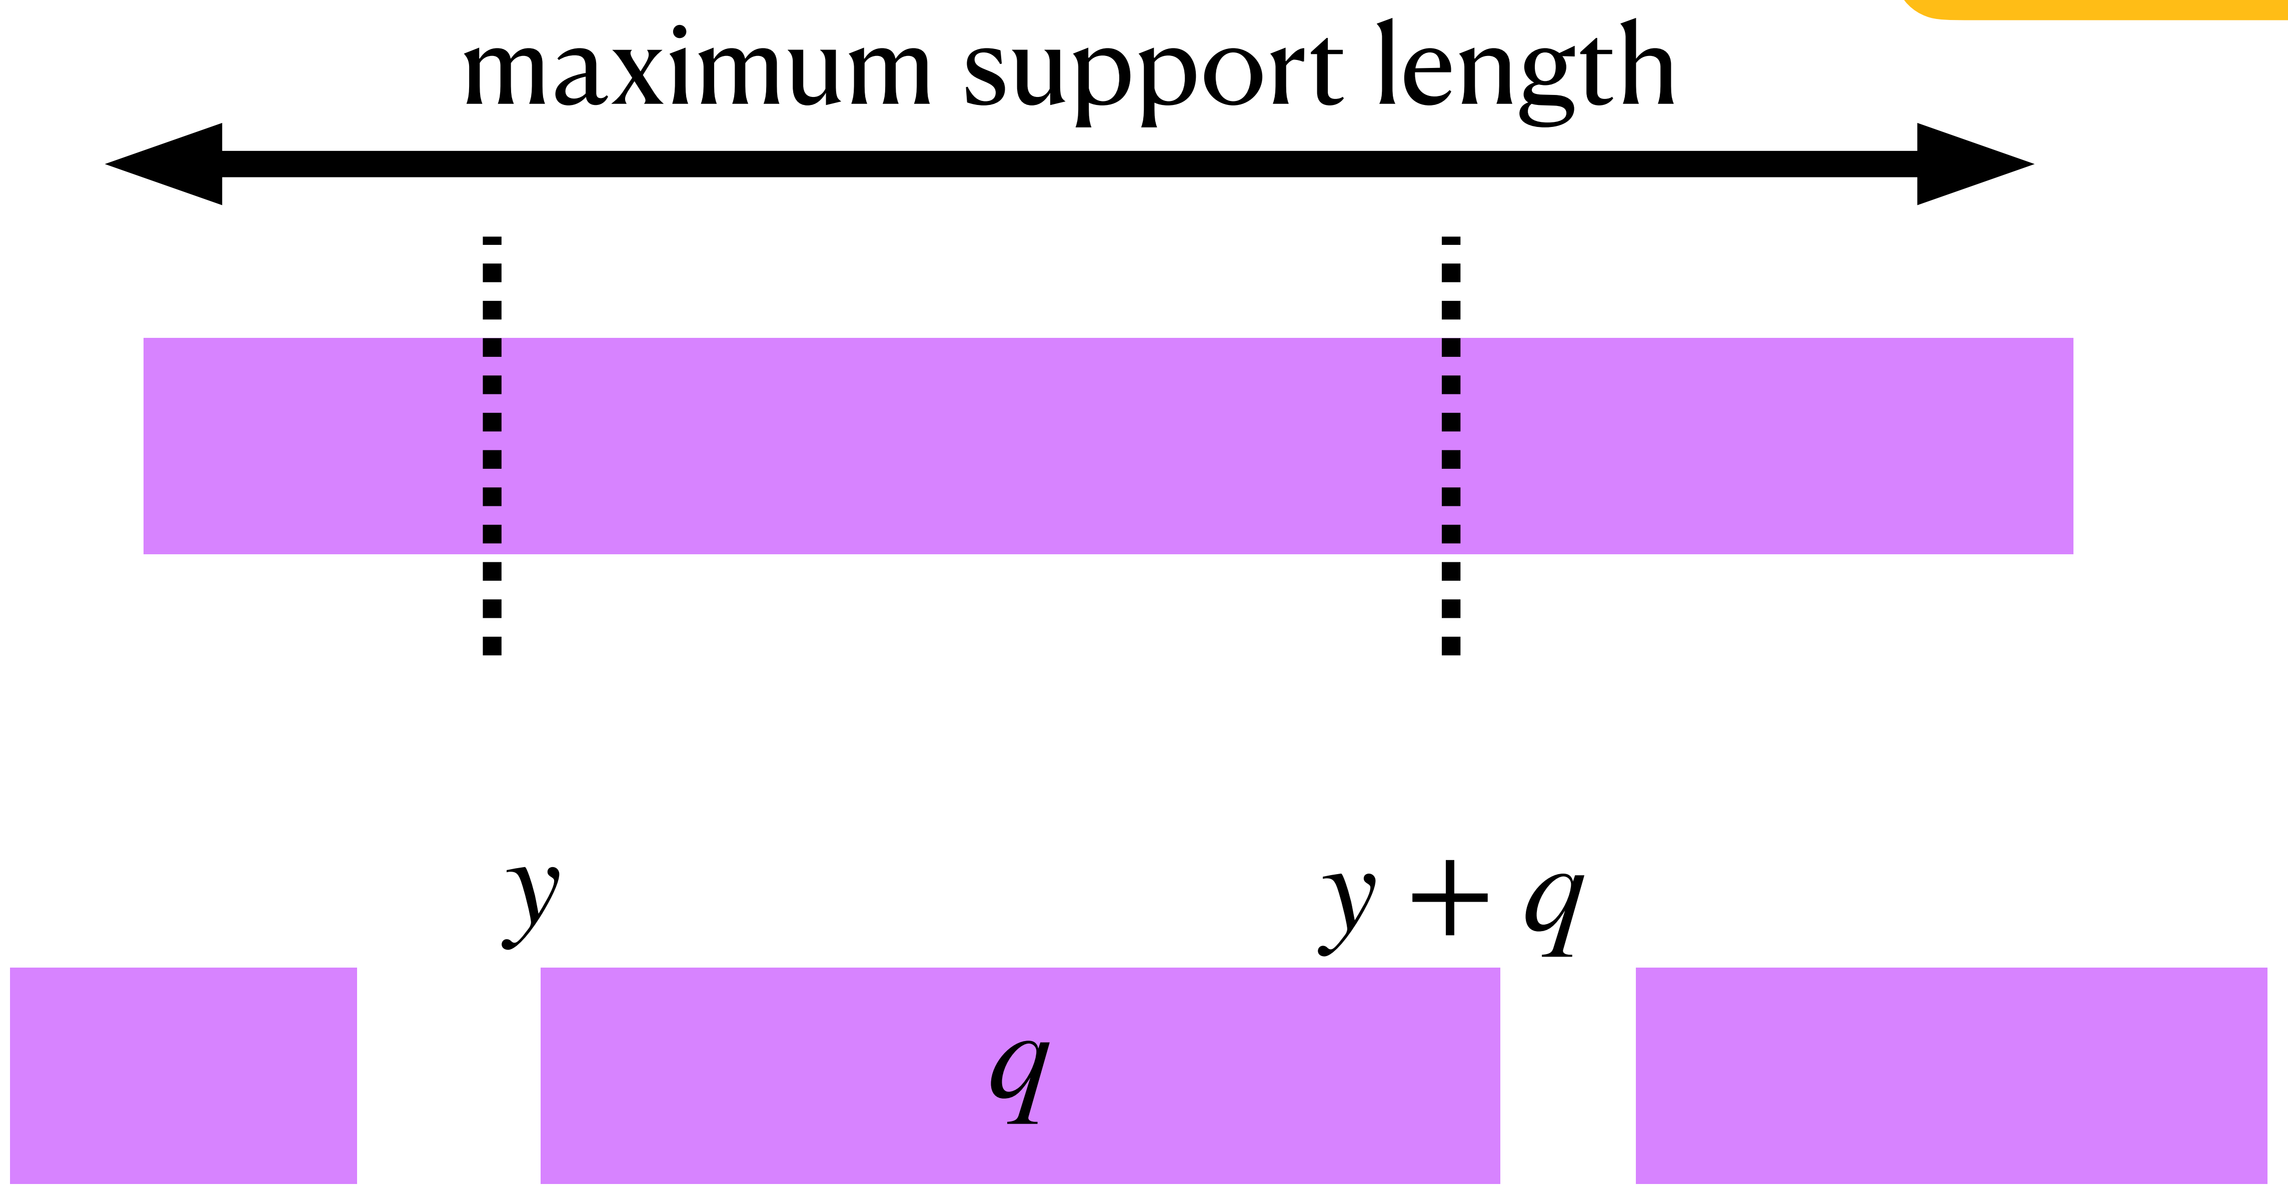

prob. of finishing in interval, given not finished now

$$\max_{q \geq 0} \frac{\mathbb{P}(P_j - y \leq q \mid P_j > y)}{\mathbb{E}[\min\{P_j - y, q\} \mid P_j > y]}$$

Gittins Index Policy

finite
support

Robust: SEPT

Why: cost expressed in terms of unconditional expected values

$$\sum_{j=1}^{n} (n-j) \cdot \mathbb{E}[P_j]$$

Not Robust: Gittins

Why: policy based on conditional expectations and probabilities that are highly sensitive to error

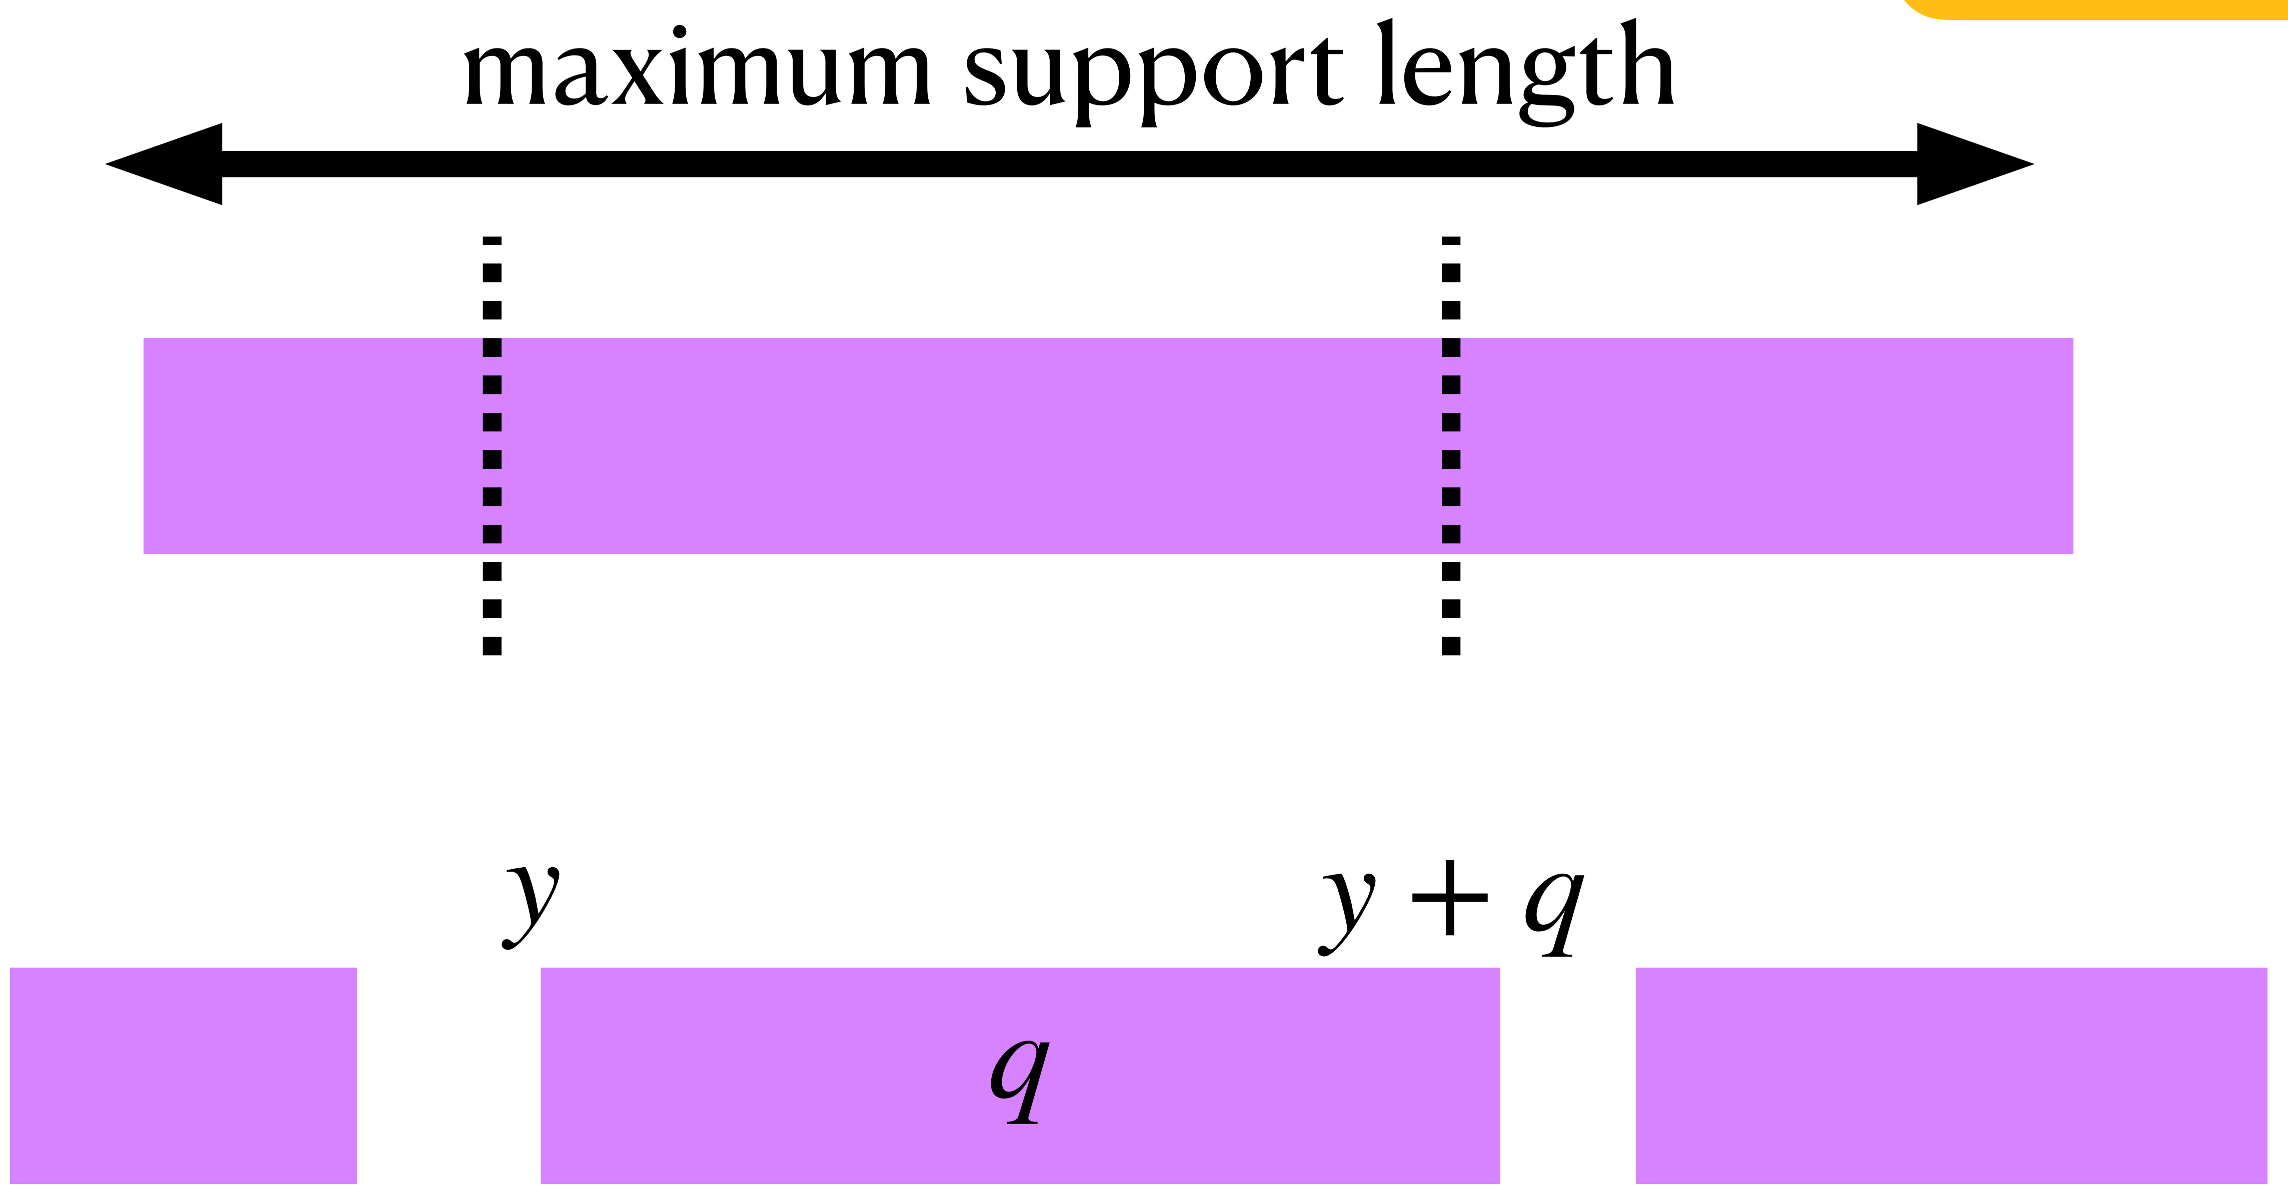

prob. of finishing in interval, given not finished now

$$\max_{q \geq 0} \frac{\mathbb{P}(P_j - y \leq q \mid P_j > y)}{\mathbb{E}[\min\{P_j - y, q\} \mid P_j > y]}$$

expected time investment, given not finished now

Gittins Index Policy

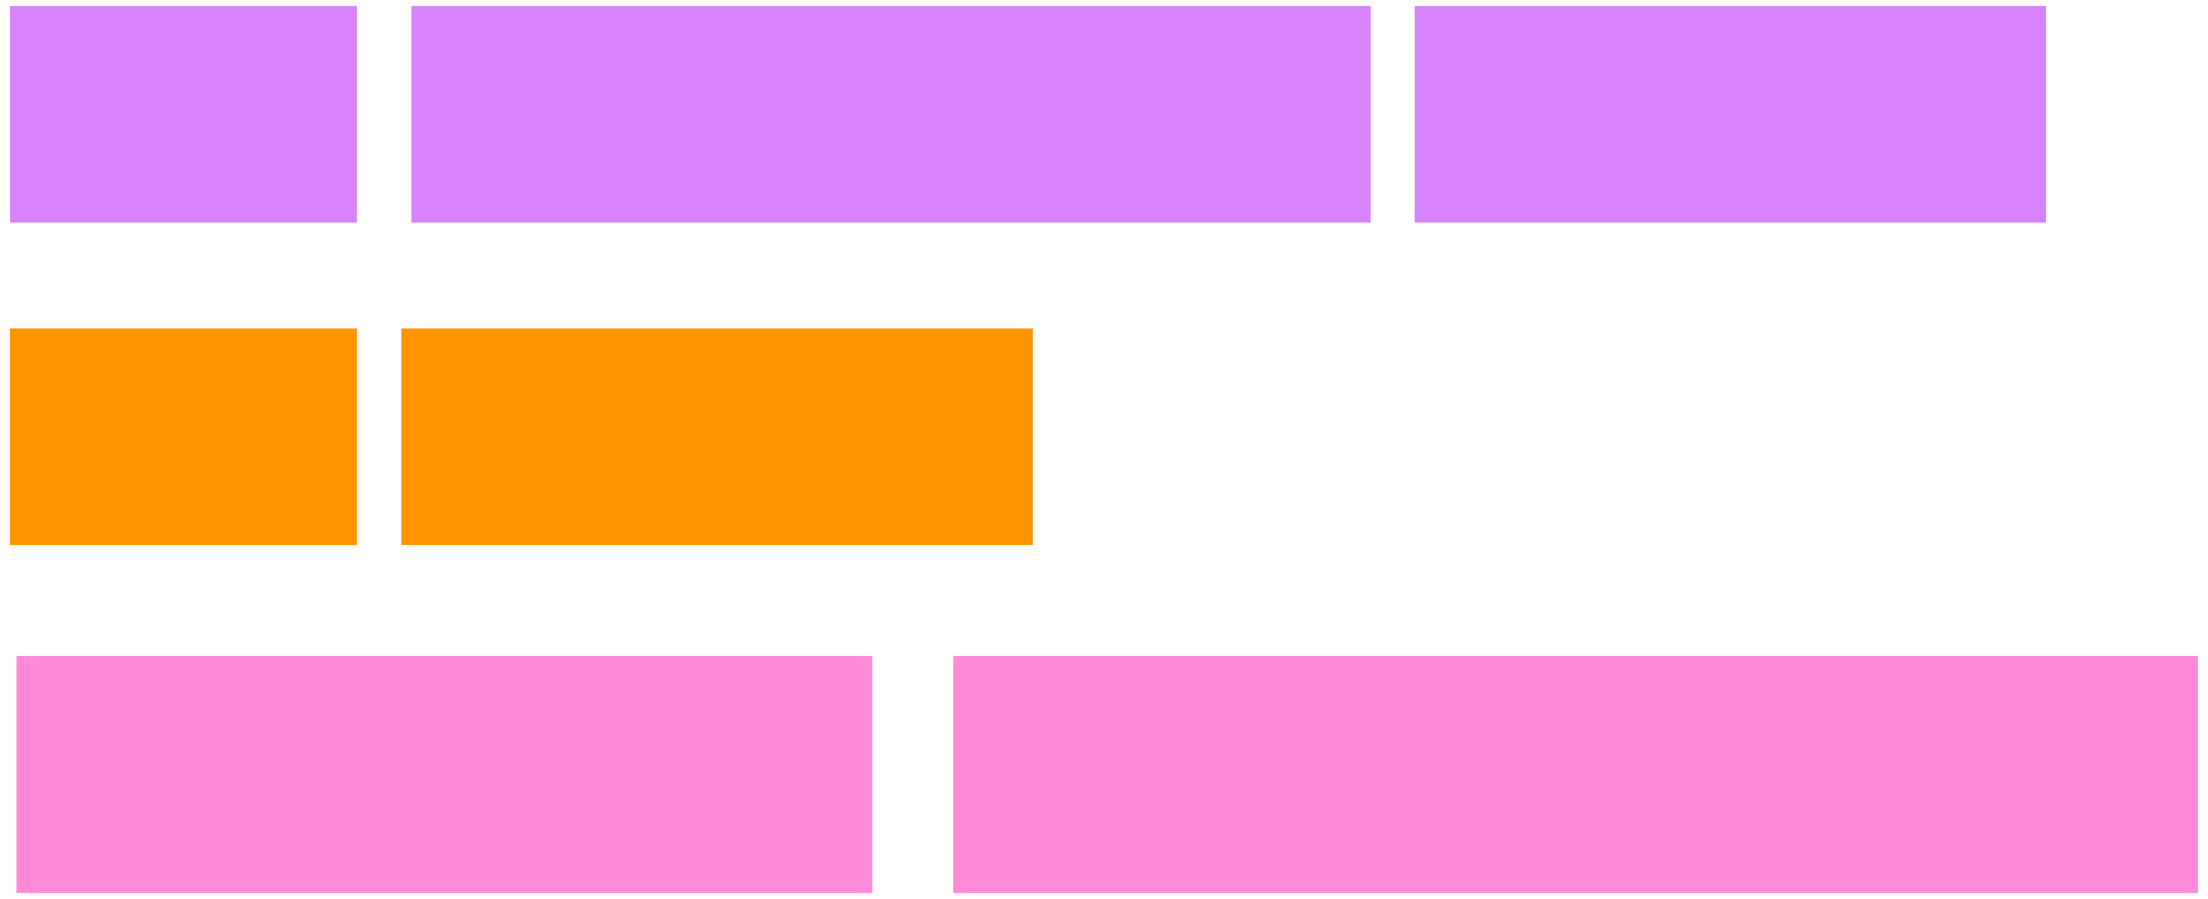

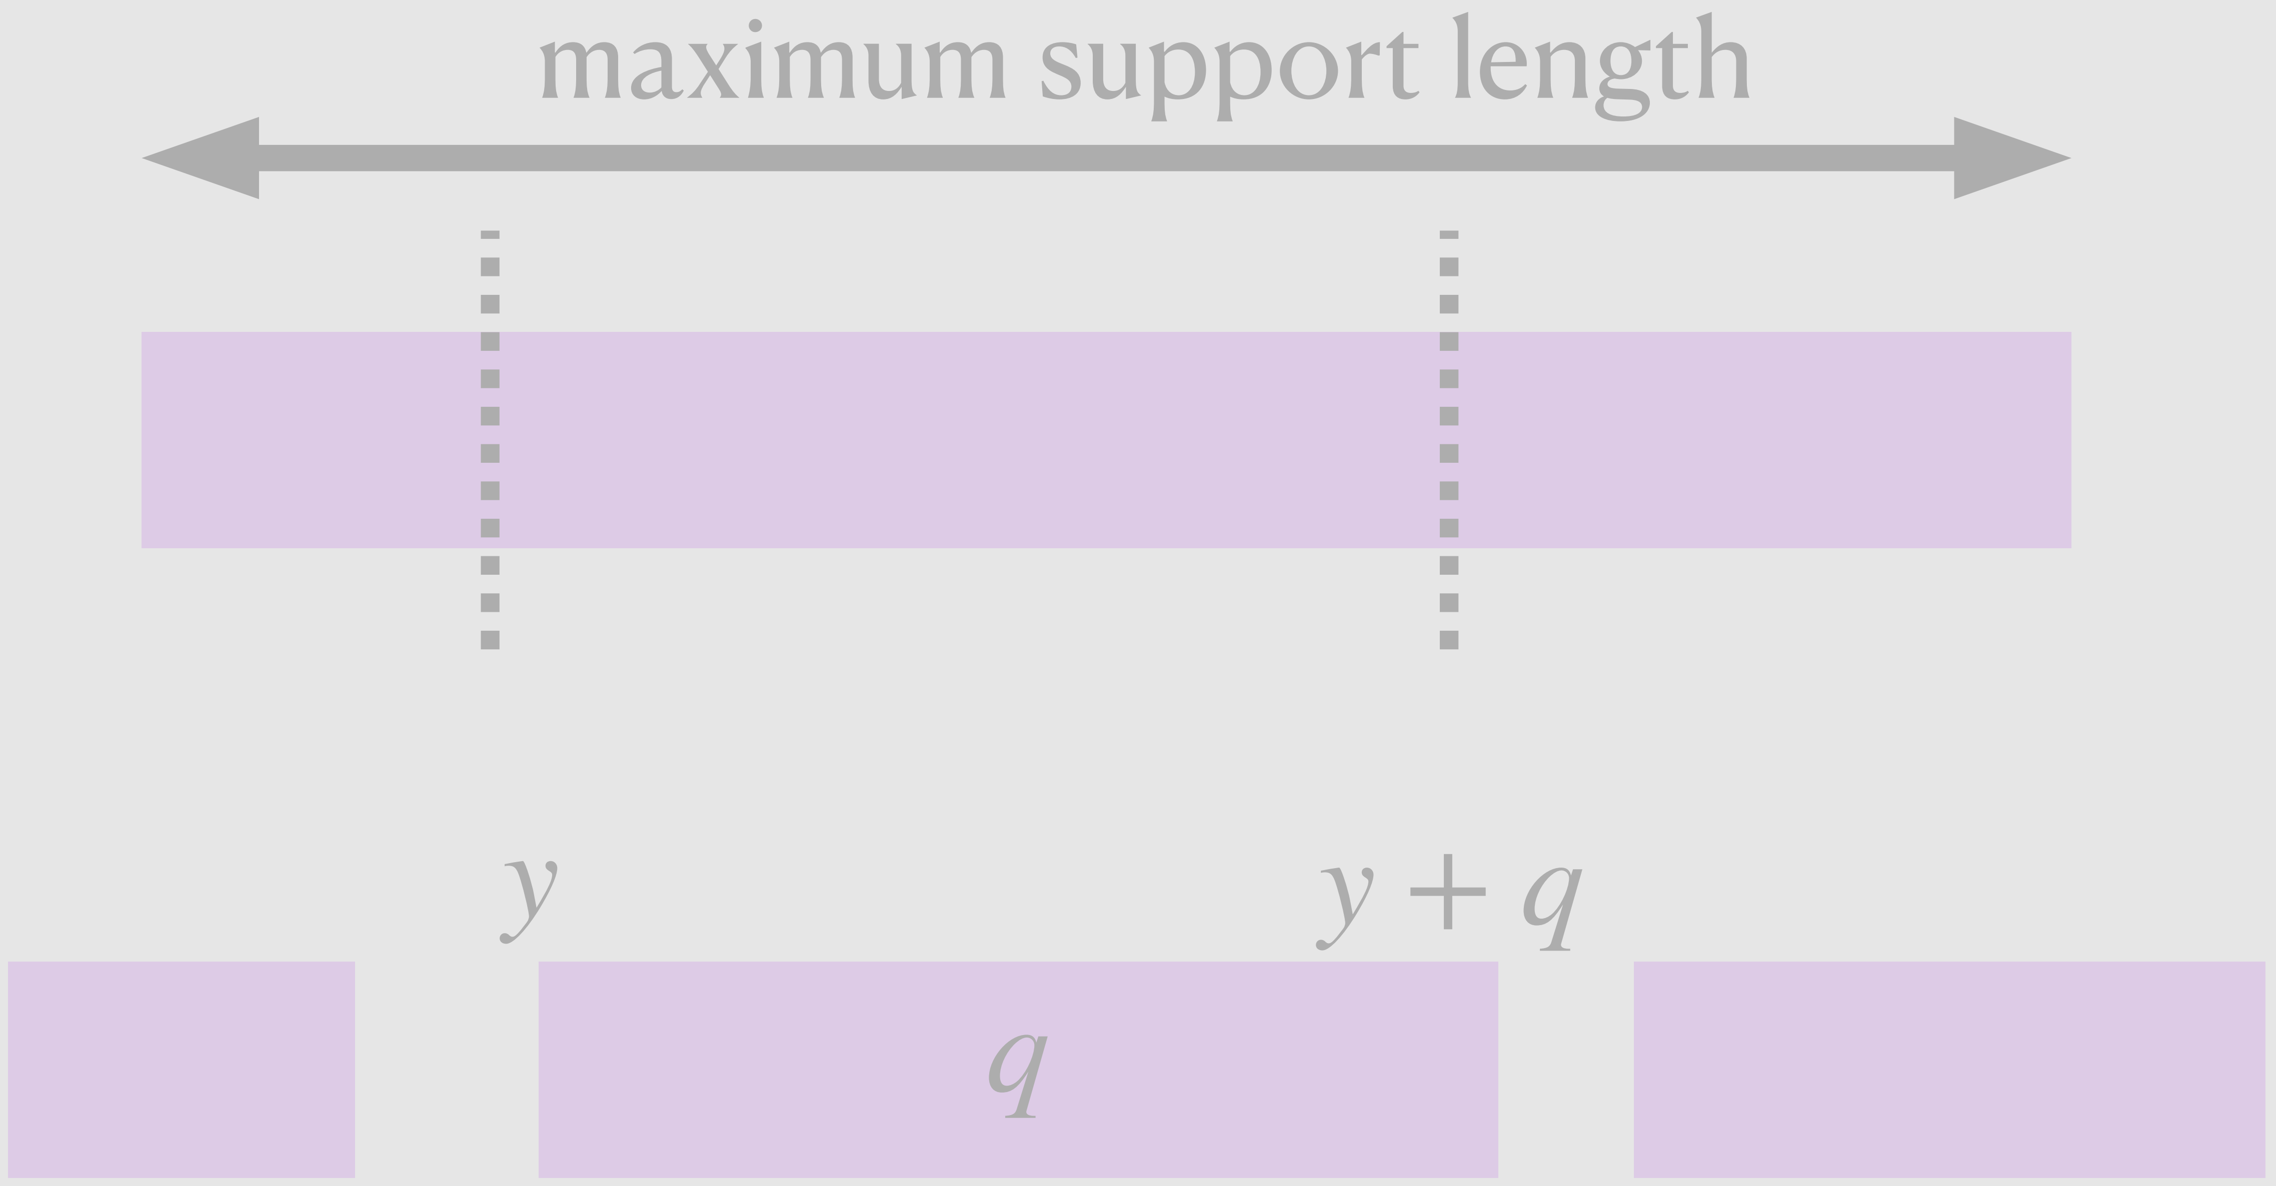

prob. of finishing in interval, given not finished now

$$\max_{q\geq 0} \frac{\mathbb{P}(P_j - y \leq q \mid P_j > y)}{\mathbb{E}[\min\{P_j - y, q\} \mid P_j > y]}$$

expected time investment, given not finished now

Gittins Index Policy

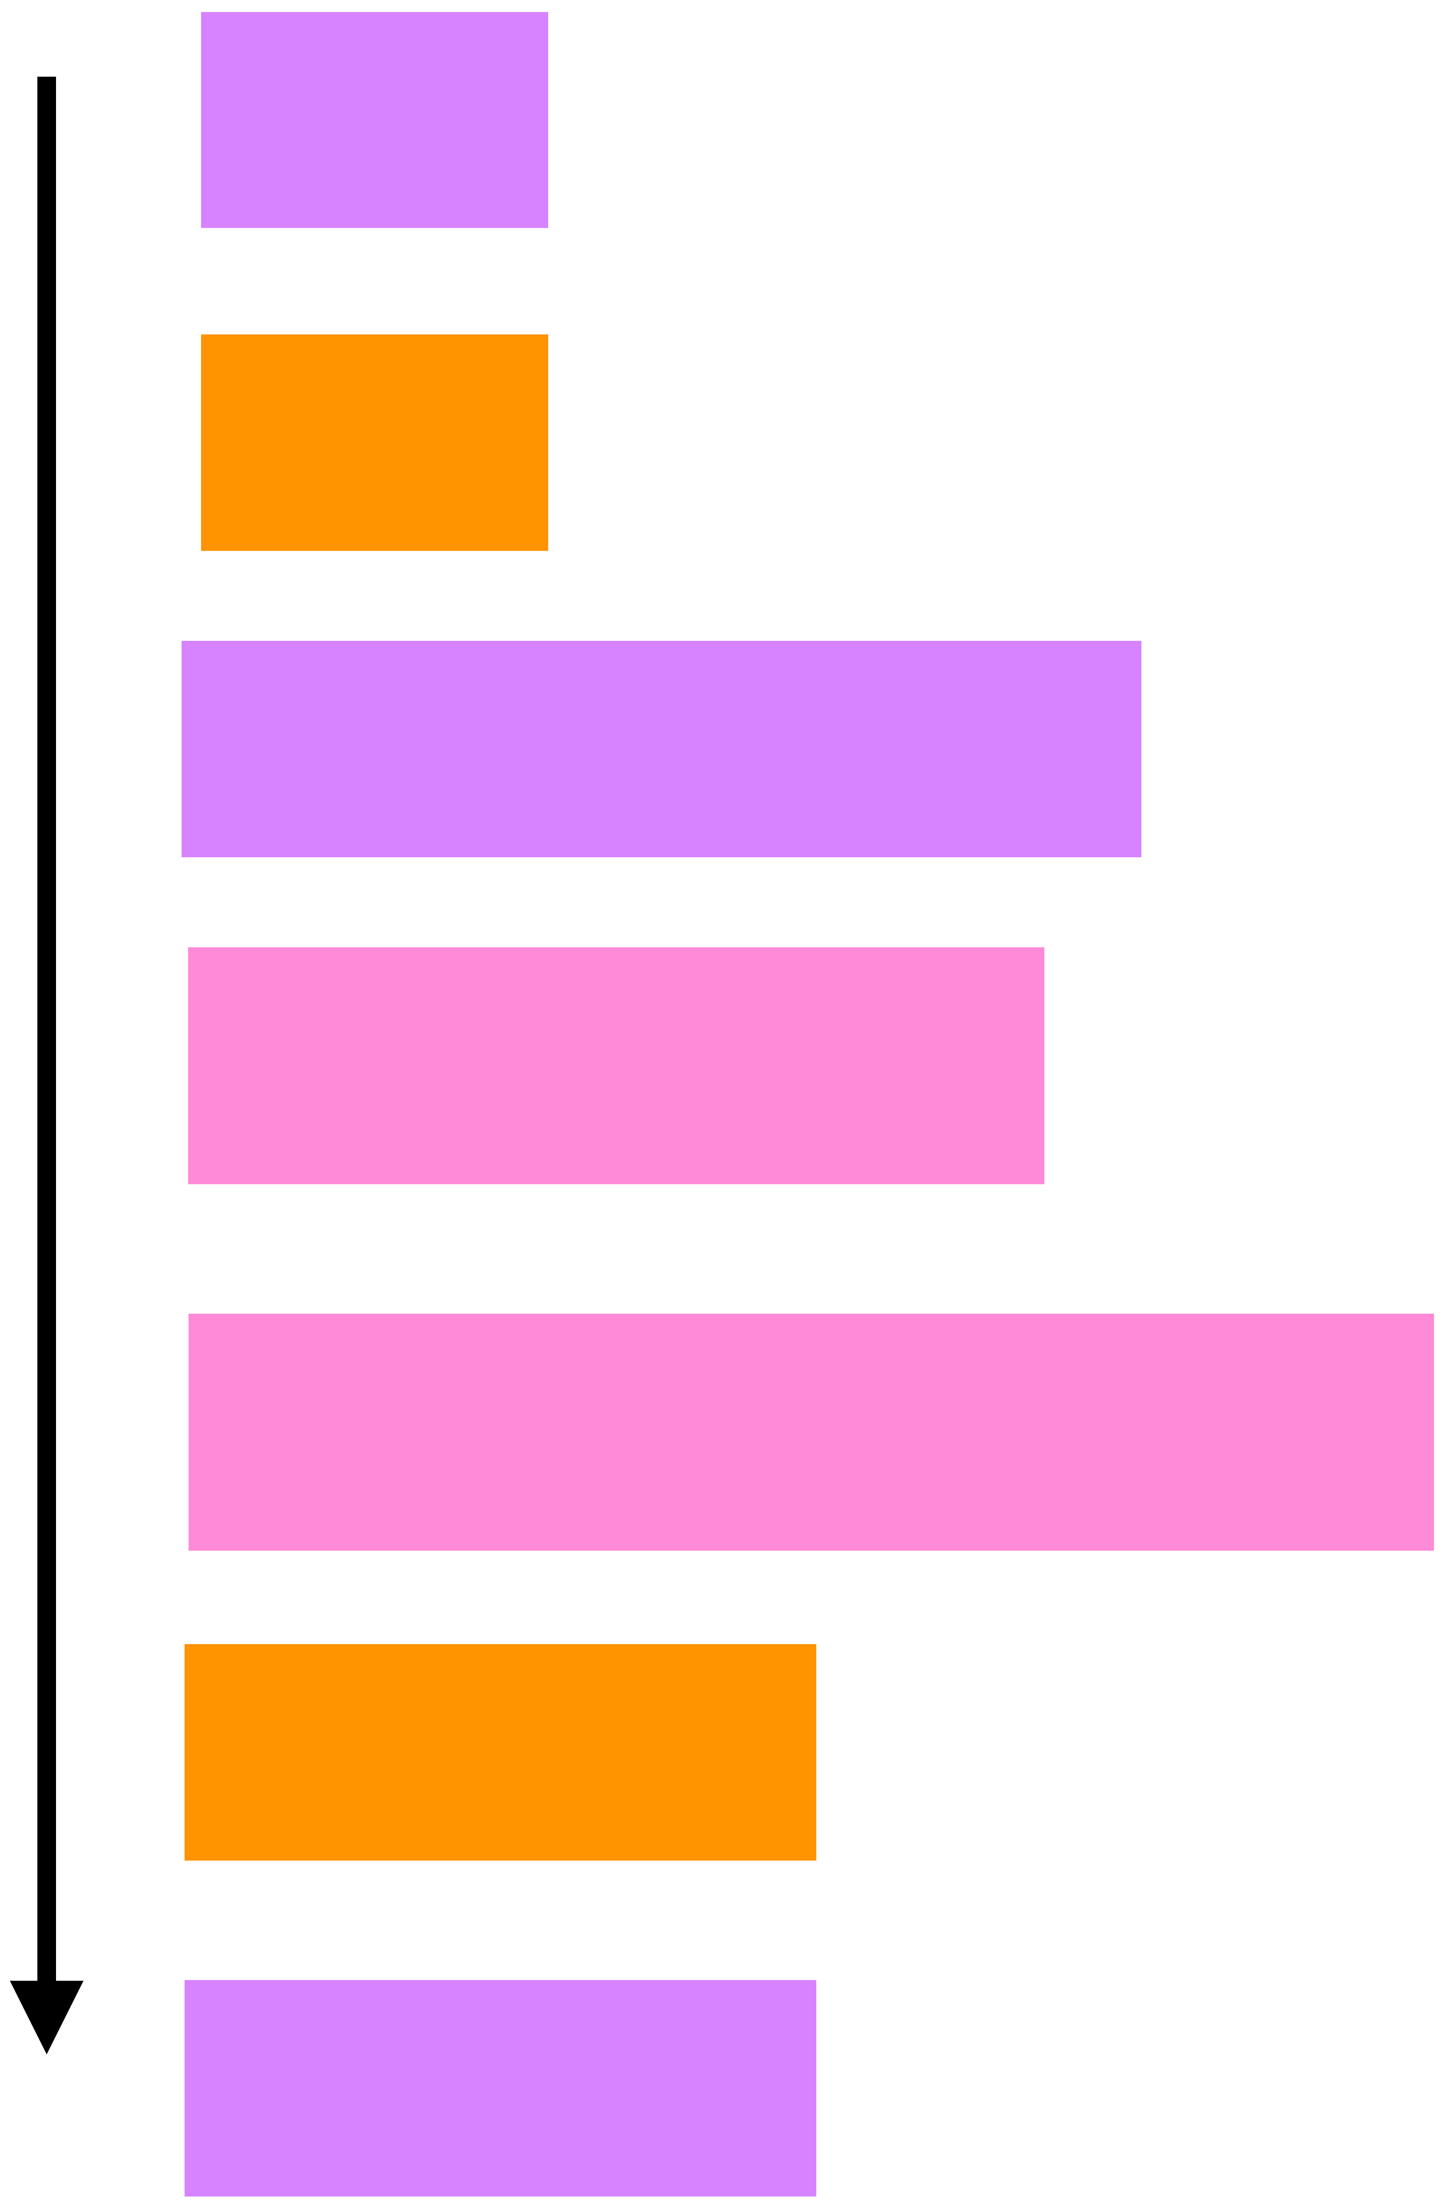

run in
decreasing
order of
rank

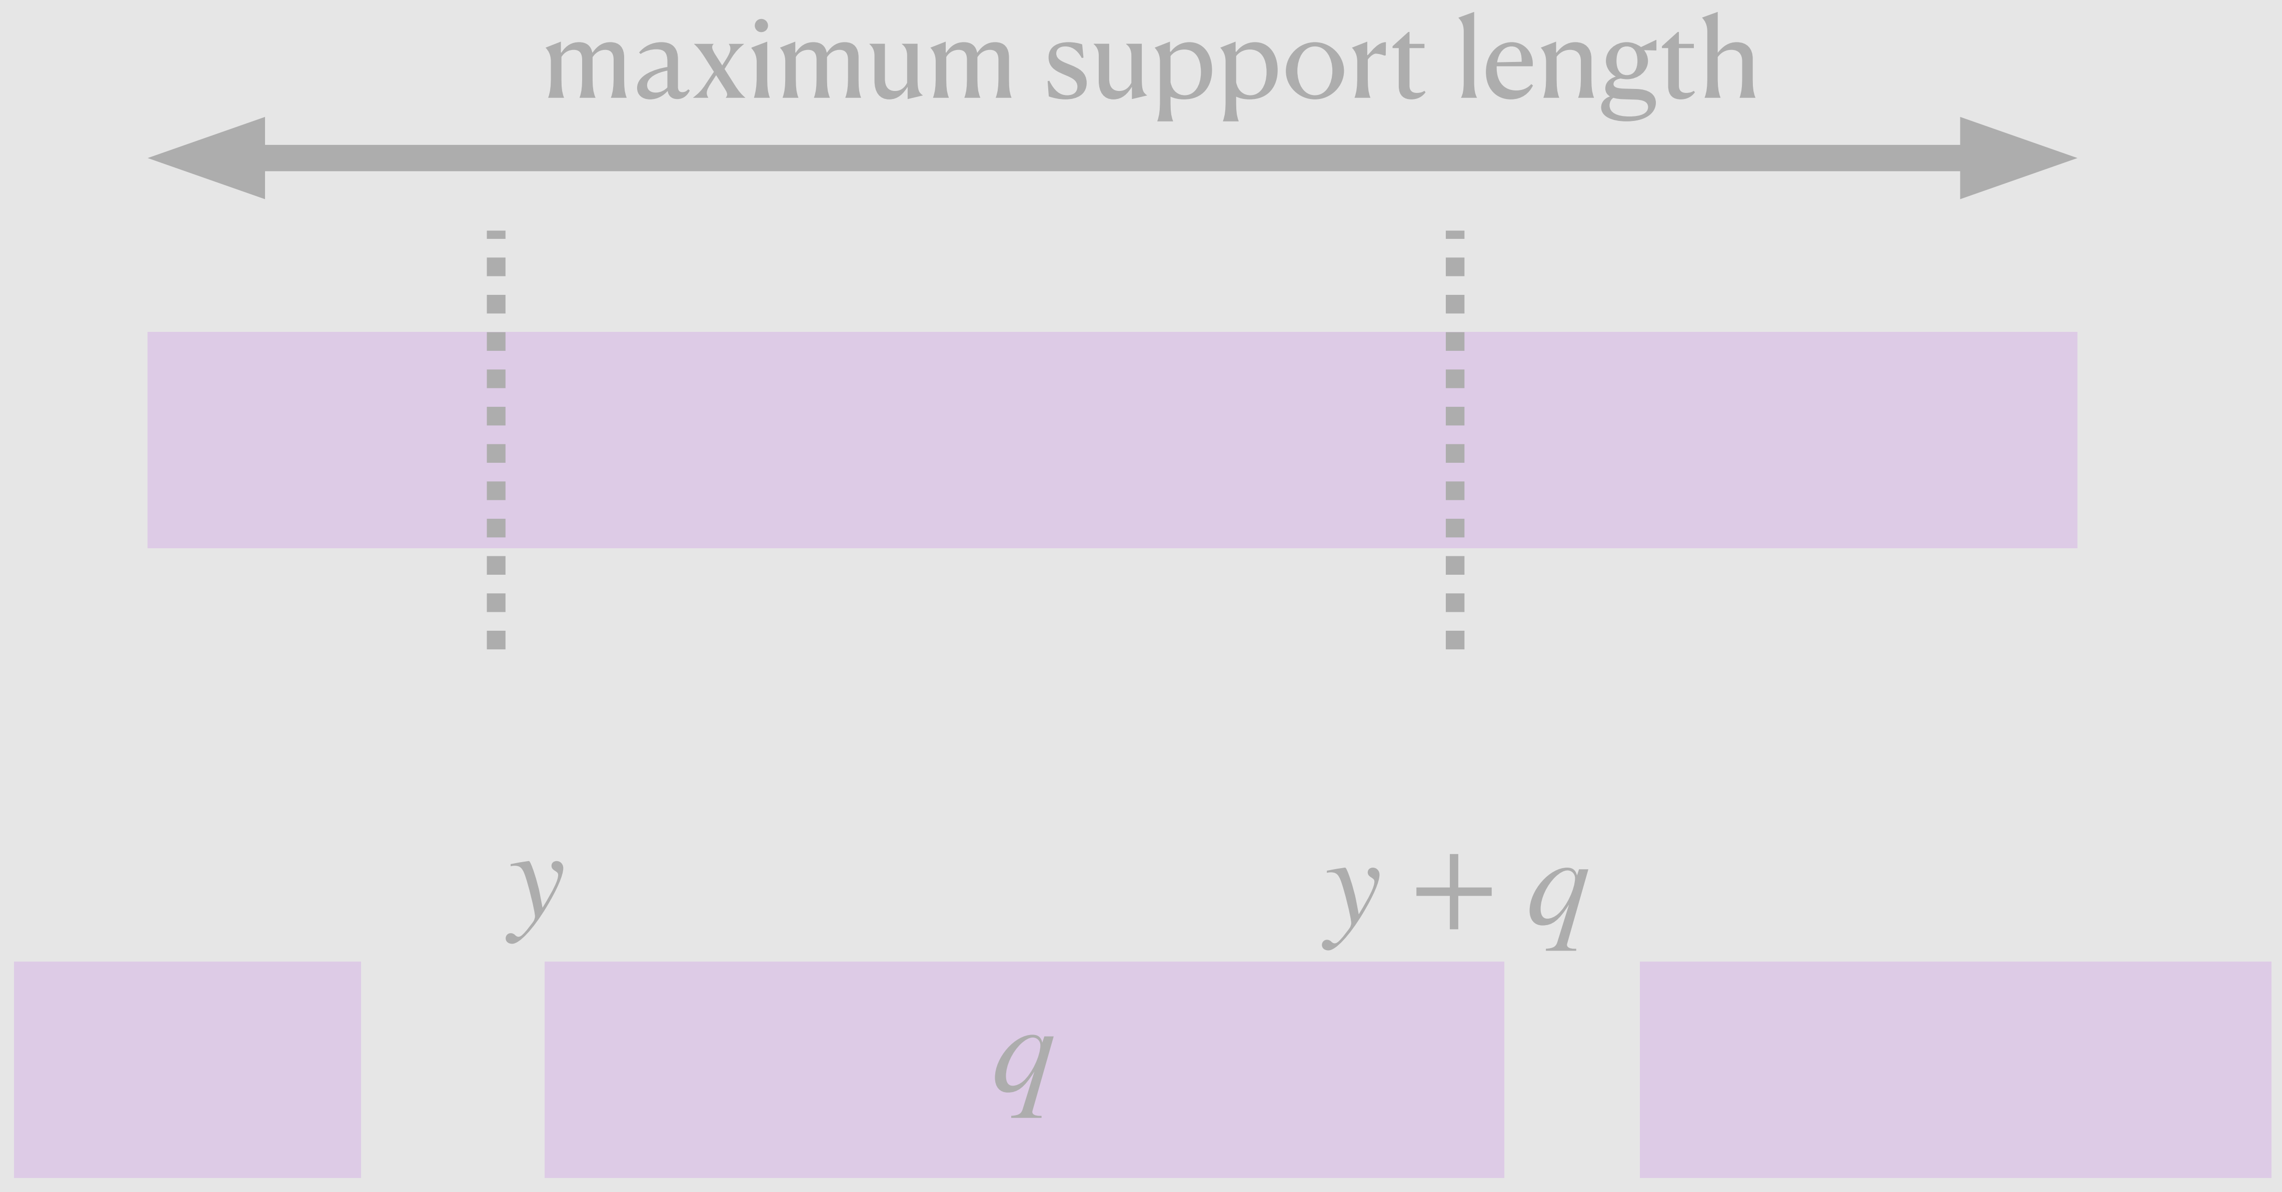

prob. of finishing in interval, given not finished now

$$\max_{q \geq 0} \frac{\mathbb{P}(P_j - y \leq q \mid P_j > y)}{\mathbb{E}[\min\{P_j - y, q\} \mid P_j > y]}$$

expected time investment, given not finished now

Gittins Index Policy

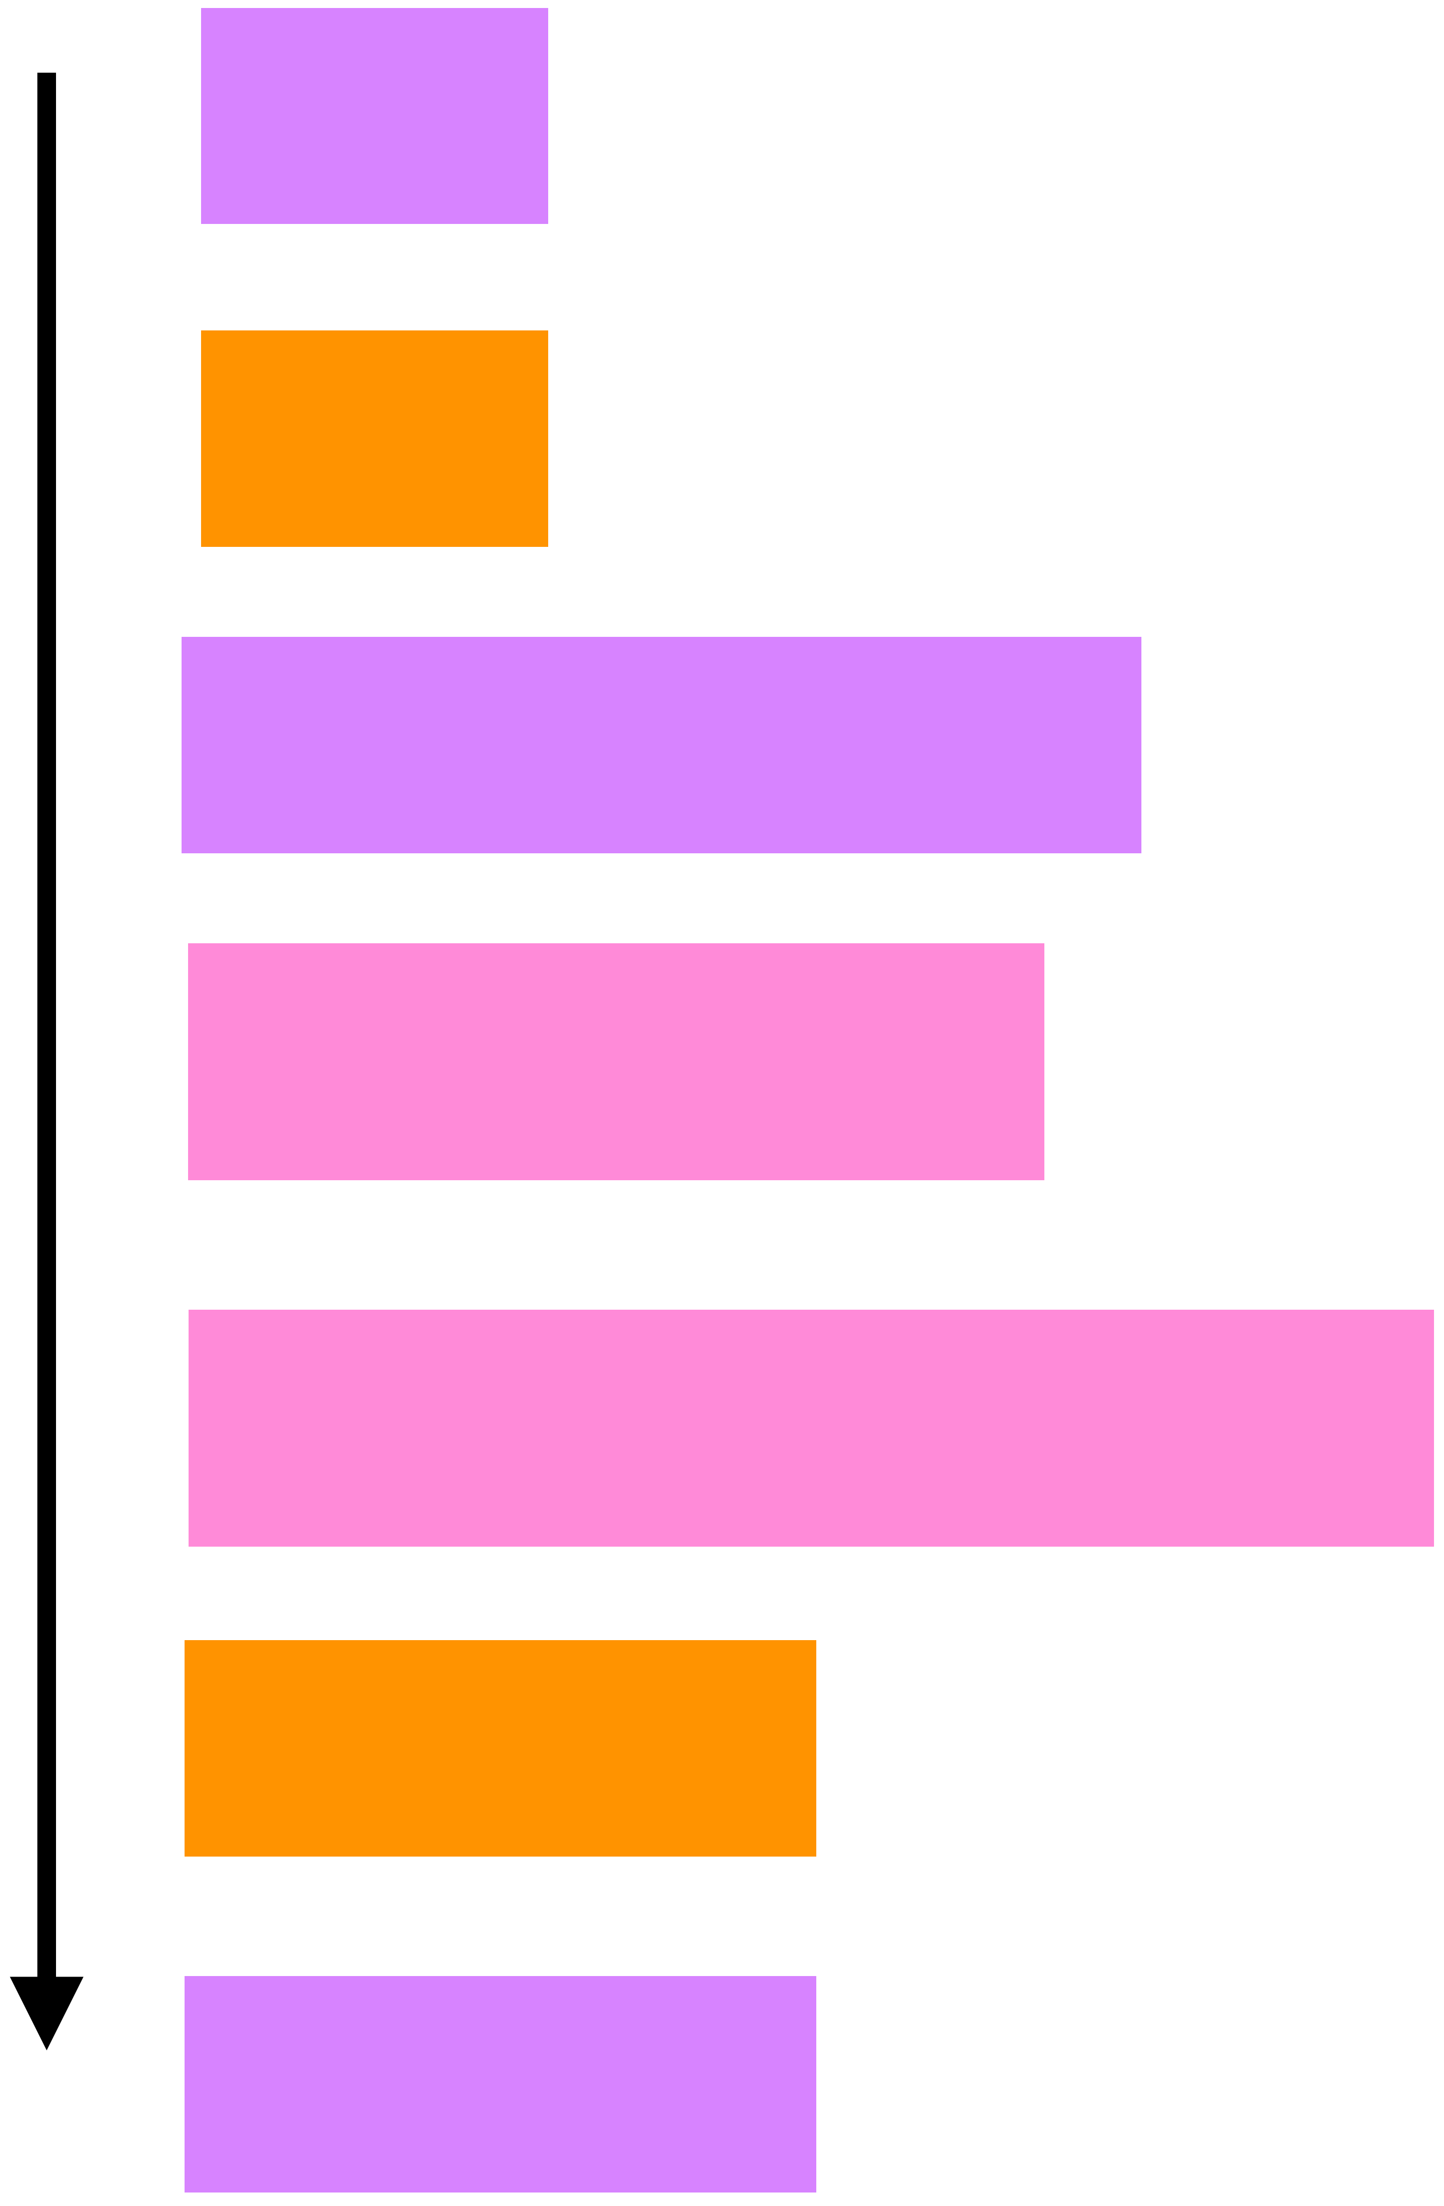

run in decreasing order of rank

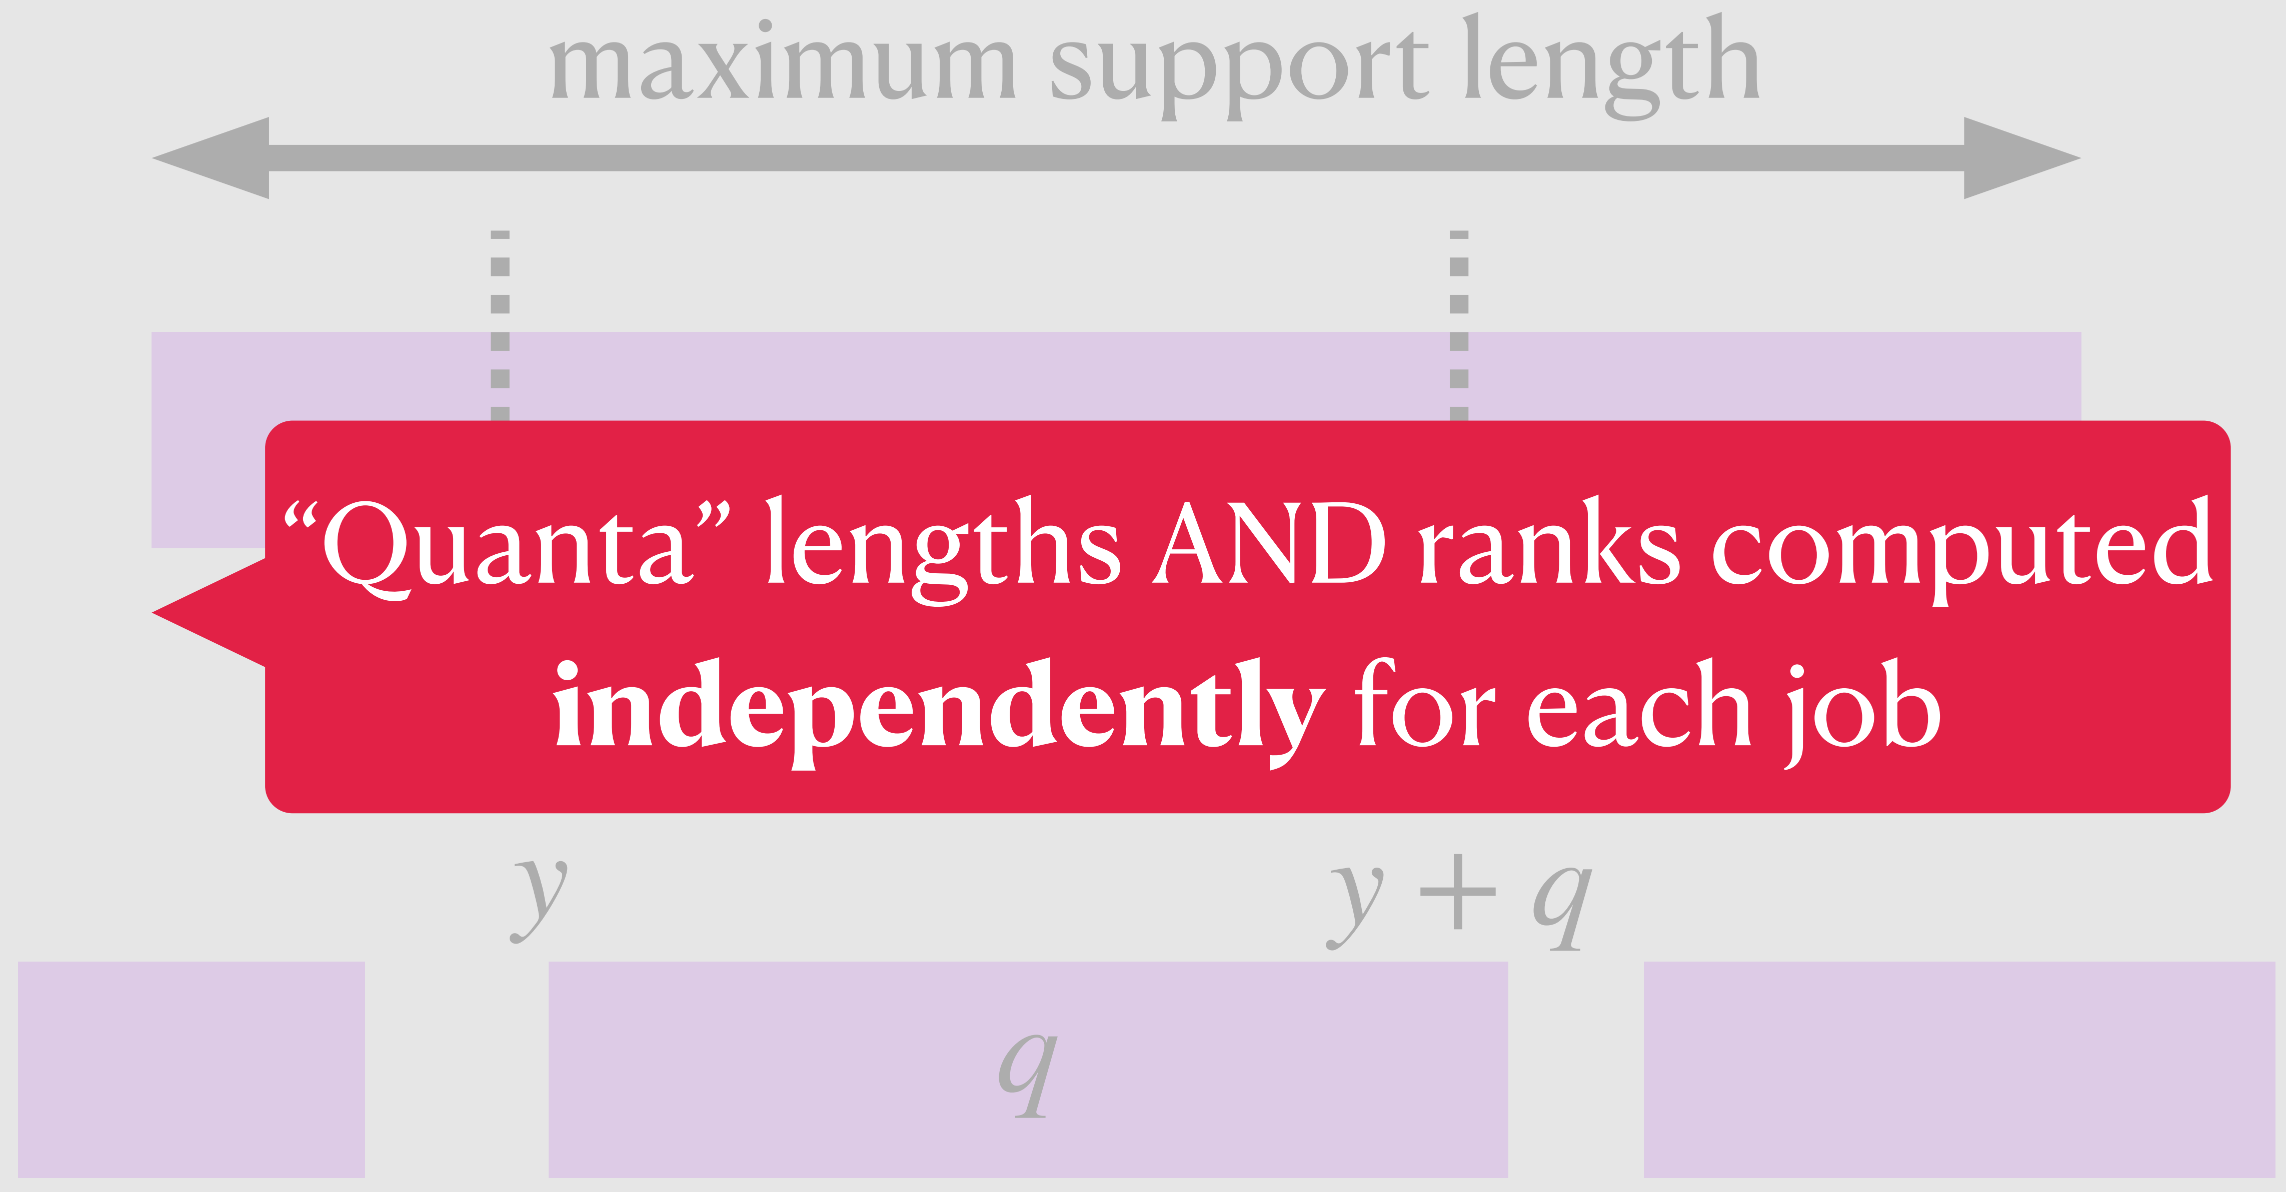

“Quanta” lengths AND ranks computed **independently** for each job

prob. of finishing in interval, given not finished now

$$\max_{q \geq 0} \frac{\mathbb{P}(P_j - y \leq q \mid P_j > y)}{\mathbb{E}[\min\{P_j - y, q\} \mid P_j > y]}$$

expected time investment, given not finished now

Gittins Index Policy

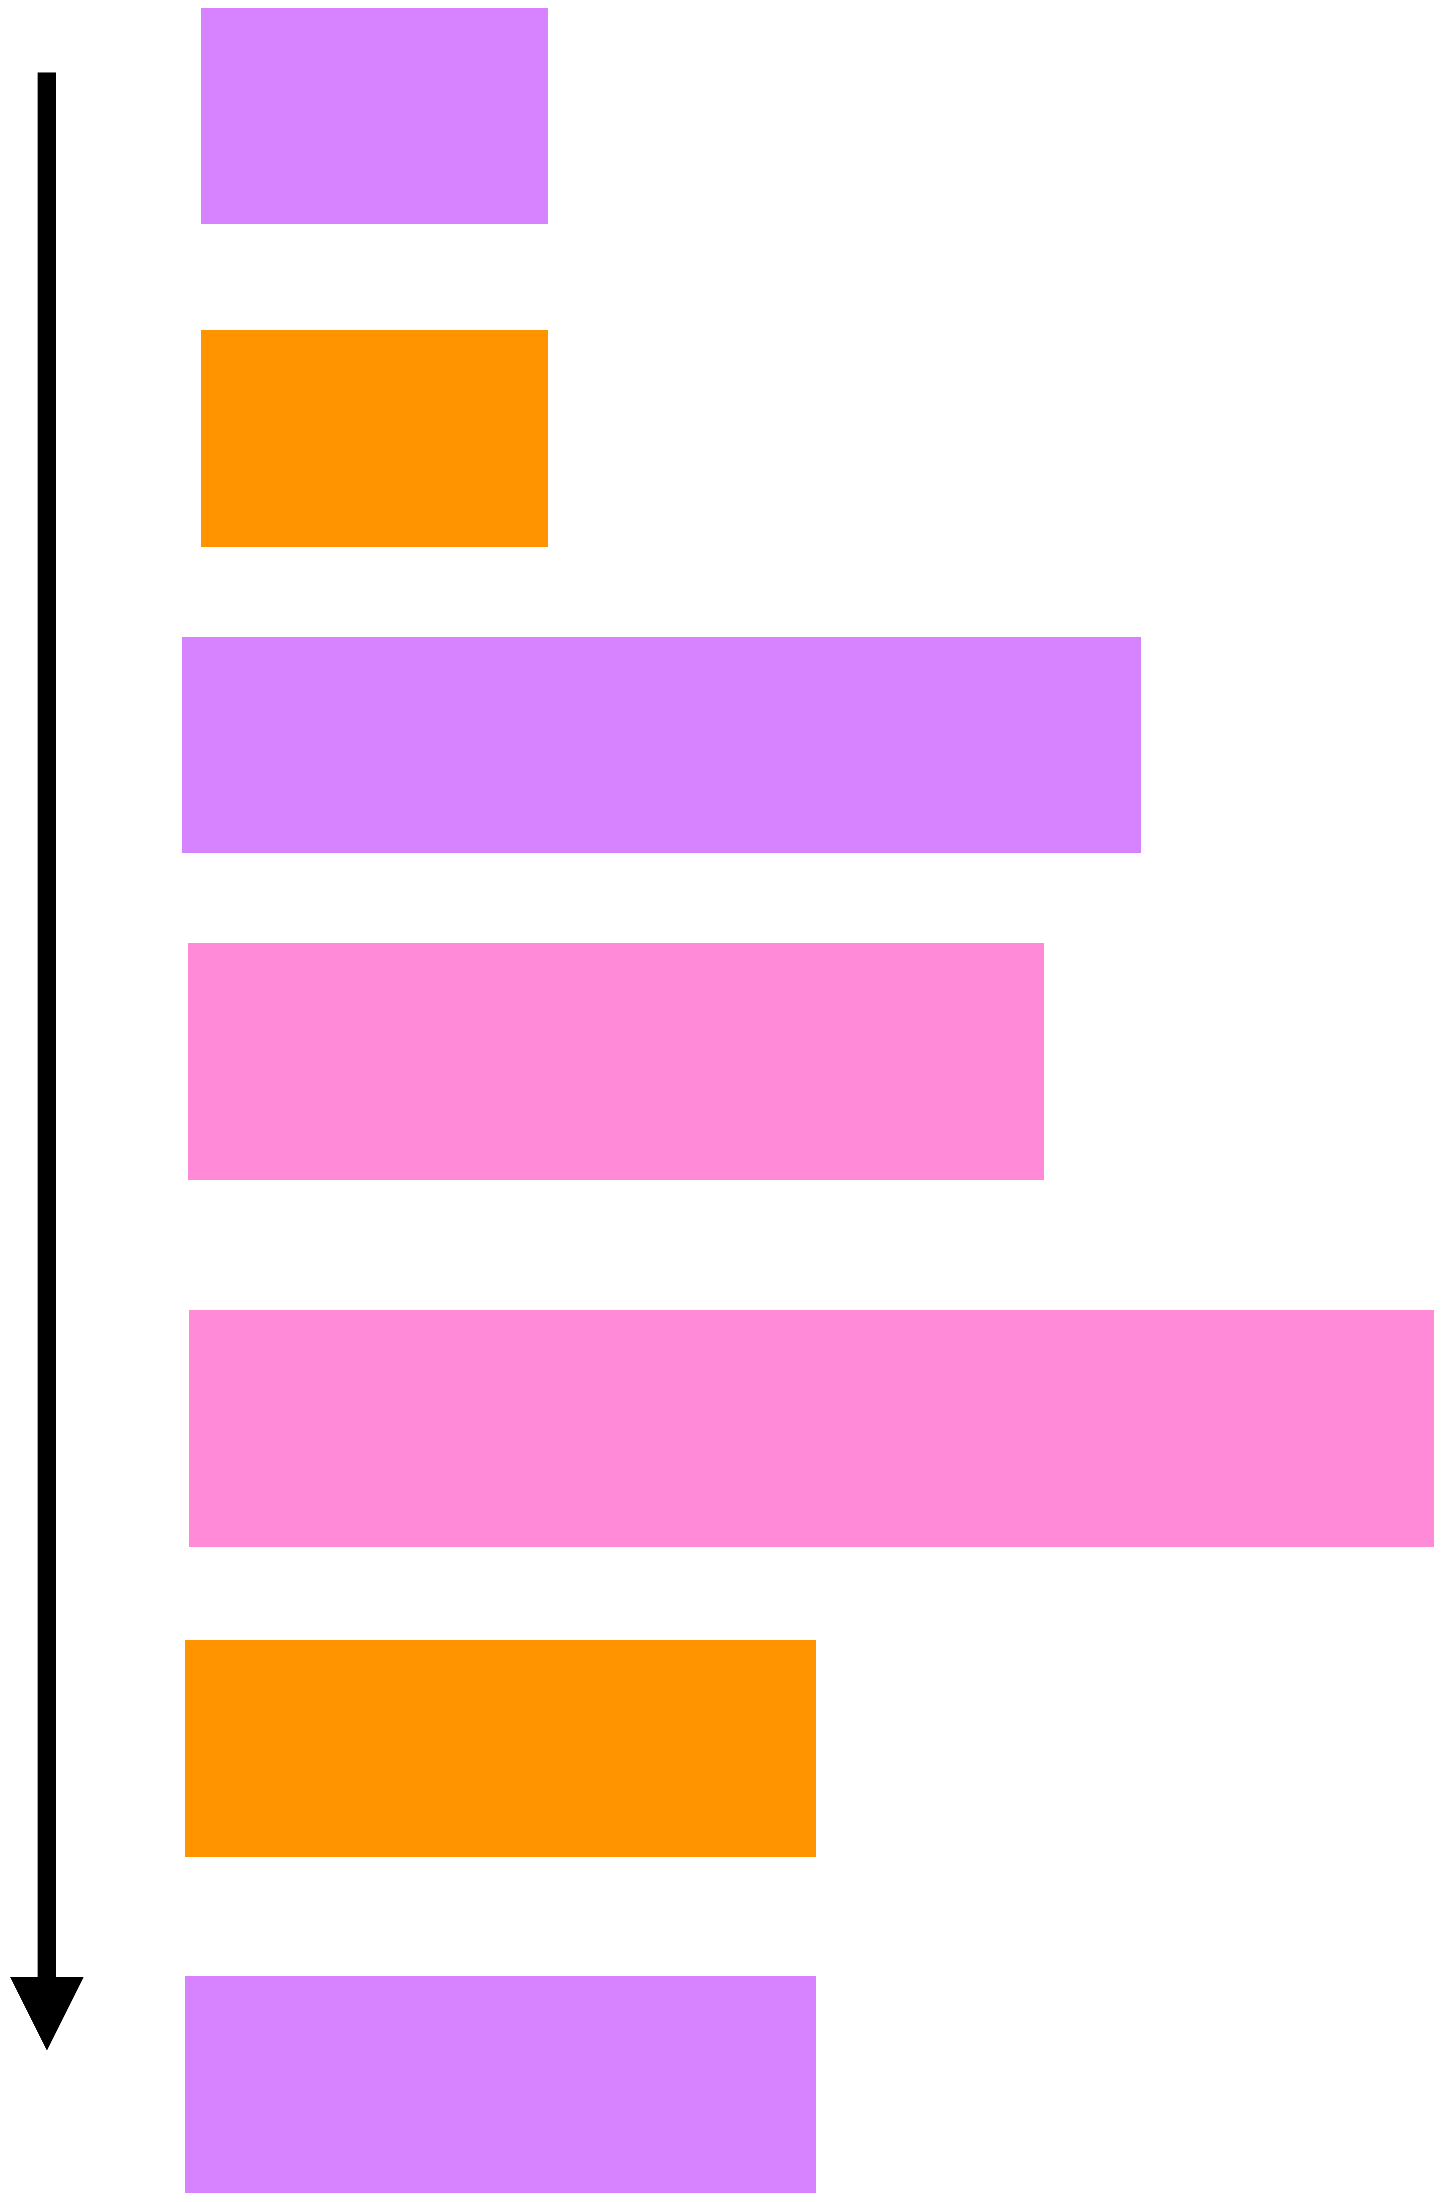

run in
decreasing
order of
rank

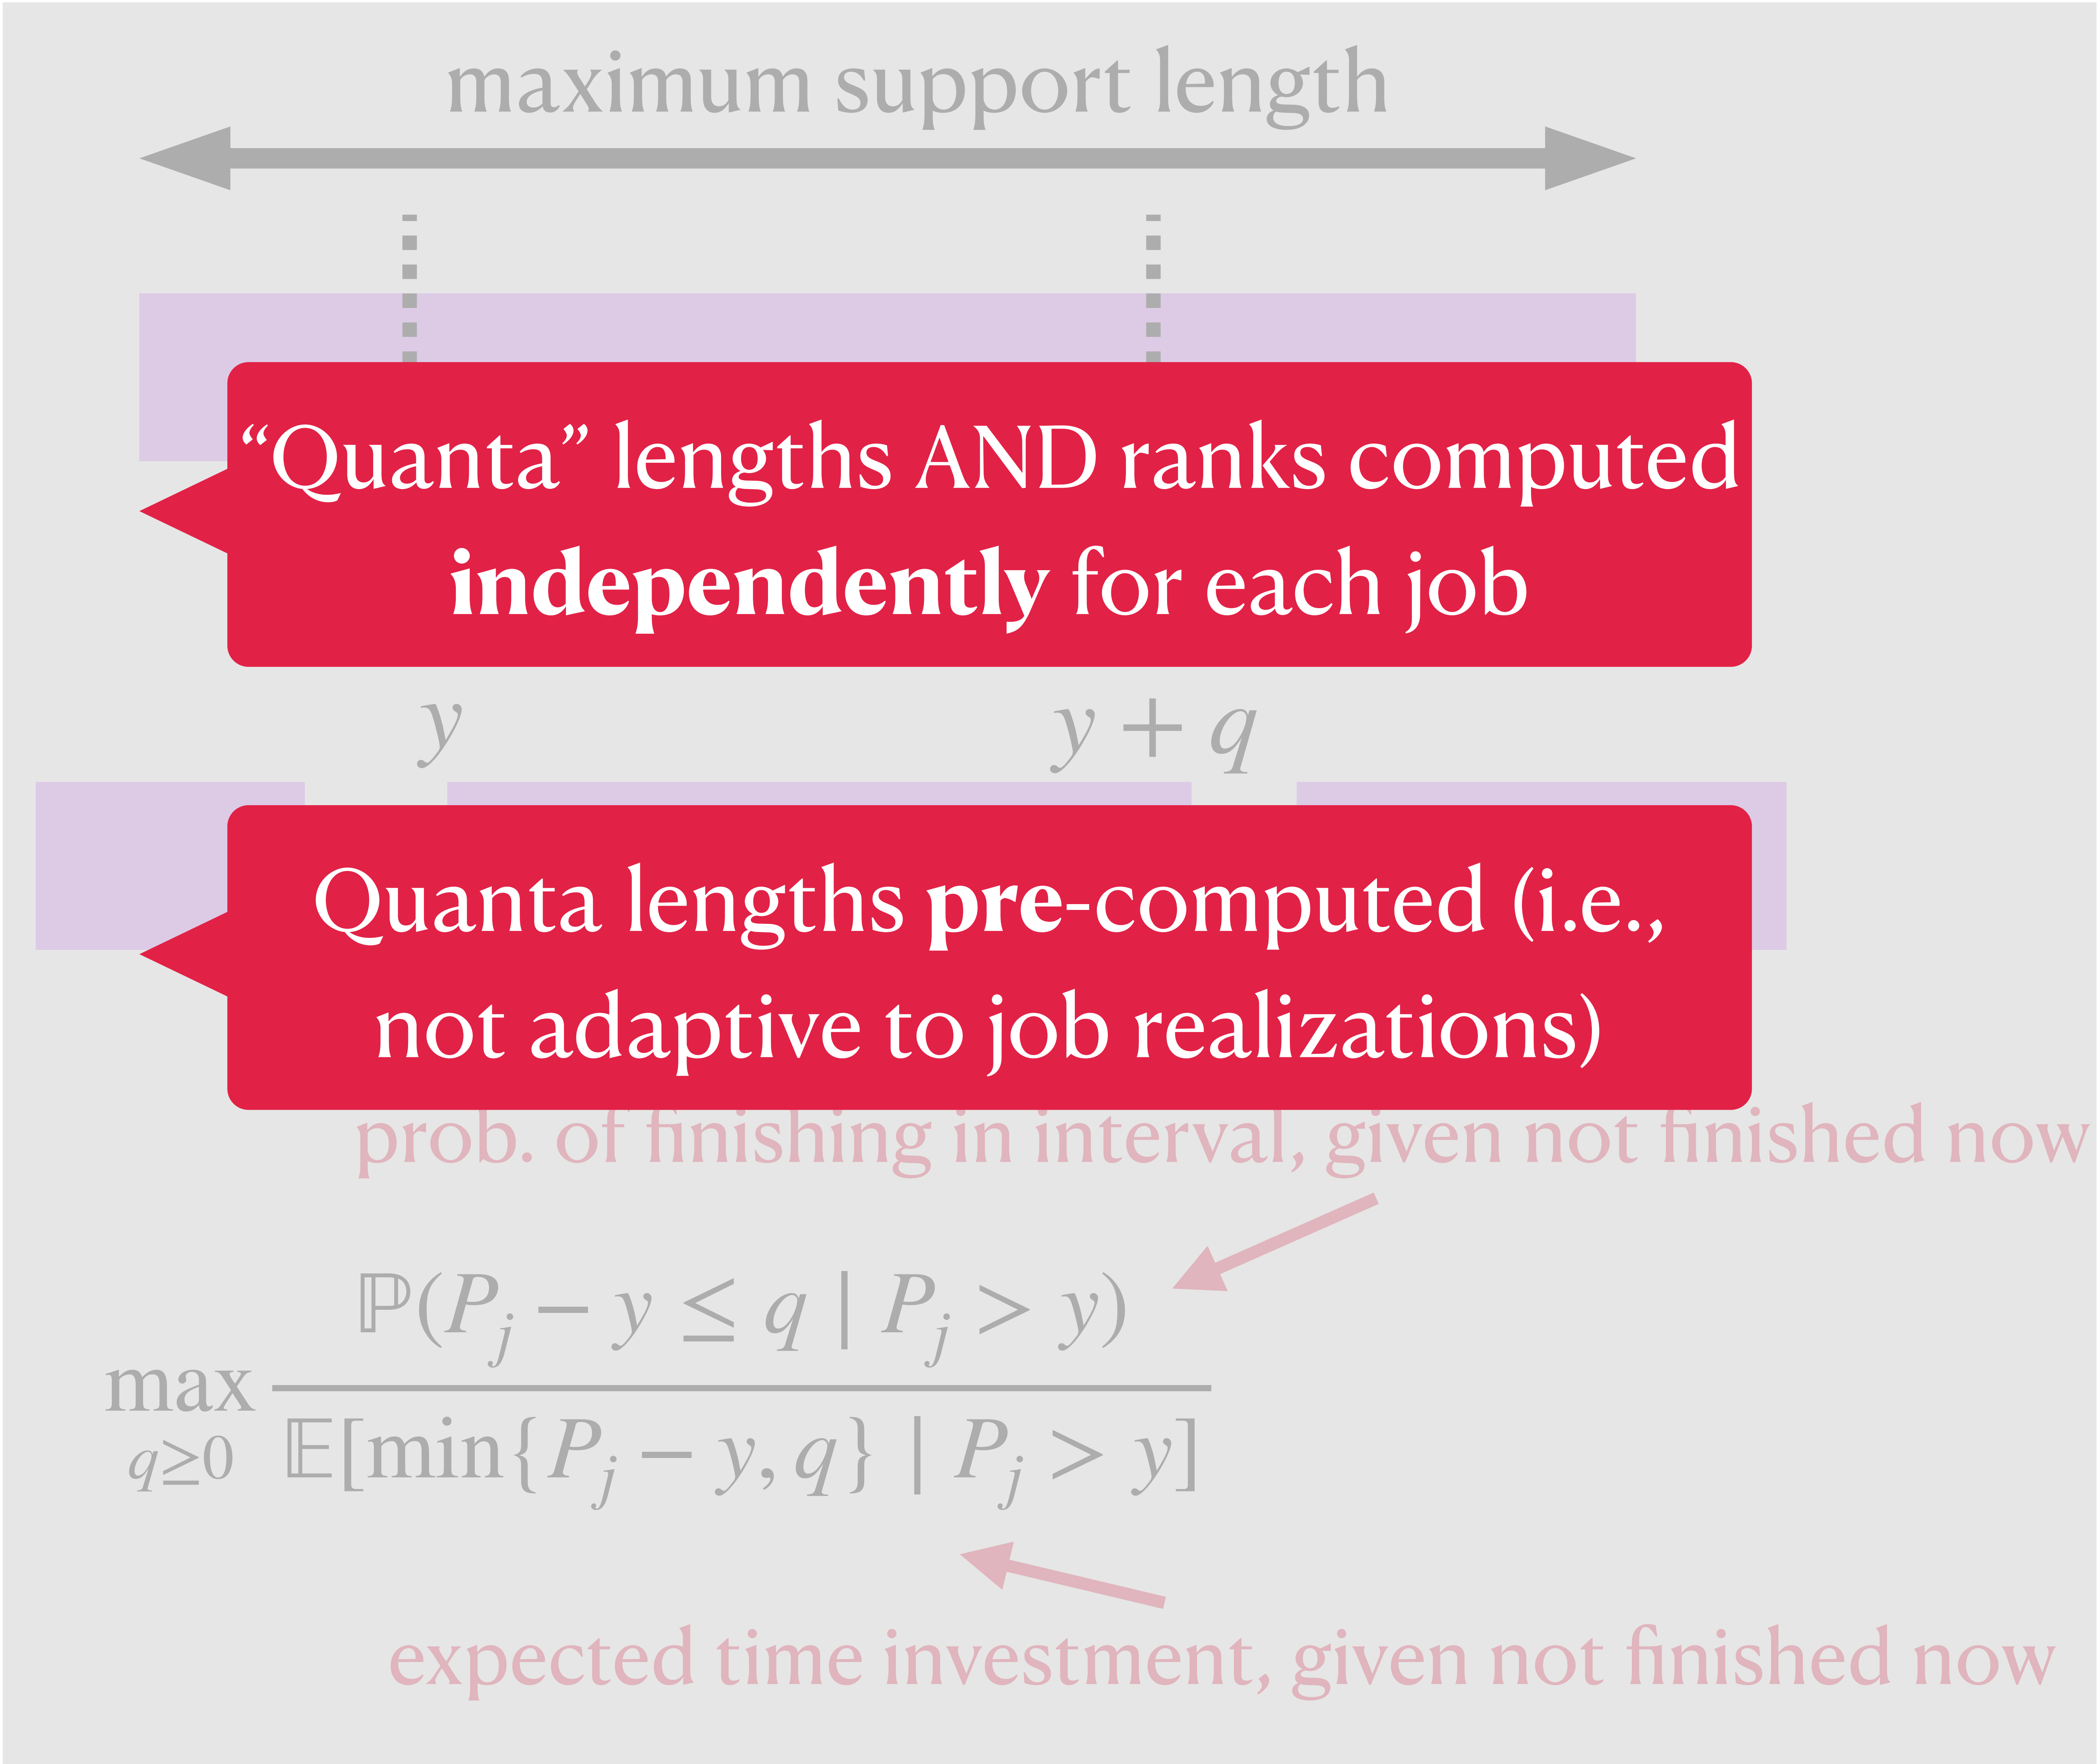

“Quanta” lengths AND ranks computed **independently** for each job

Quanta lengths **pre**-computed (i.e., not adaptive to job realizations)

prob. of finishing in interval, given not finished now

$$\max_{q\geq 0} \frac{\mathbb{P}(P_j - y \leq q \mid P_j > y)}{\mathbb{E}[\min\{P_j - y, q\} \mid P_j > y]}$$

expected time investment, given not finished now

Gittins Index Policy

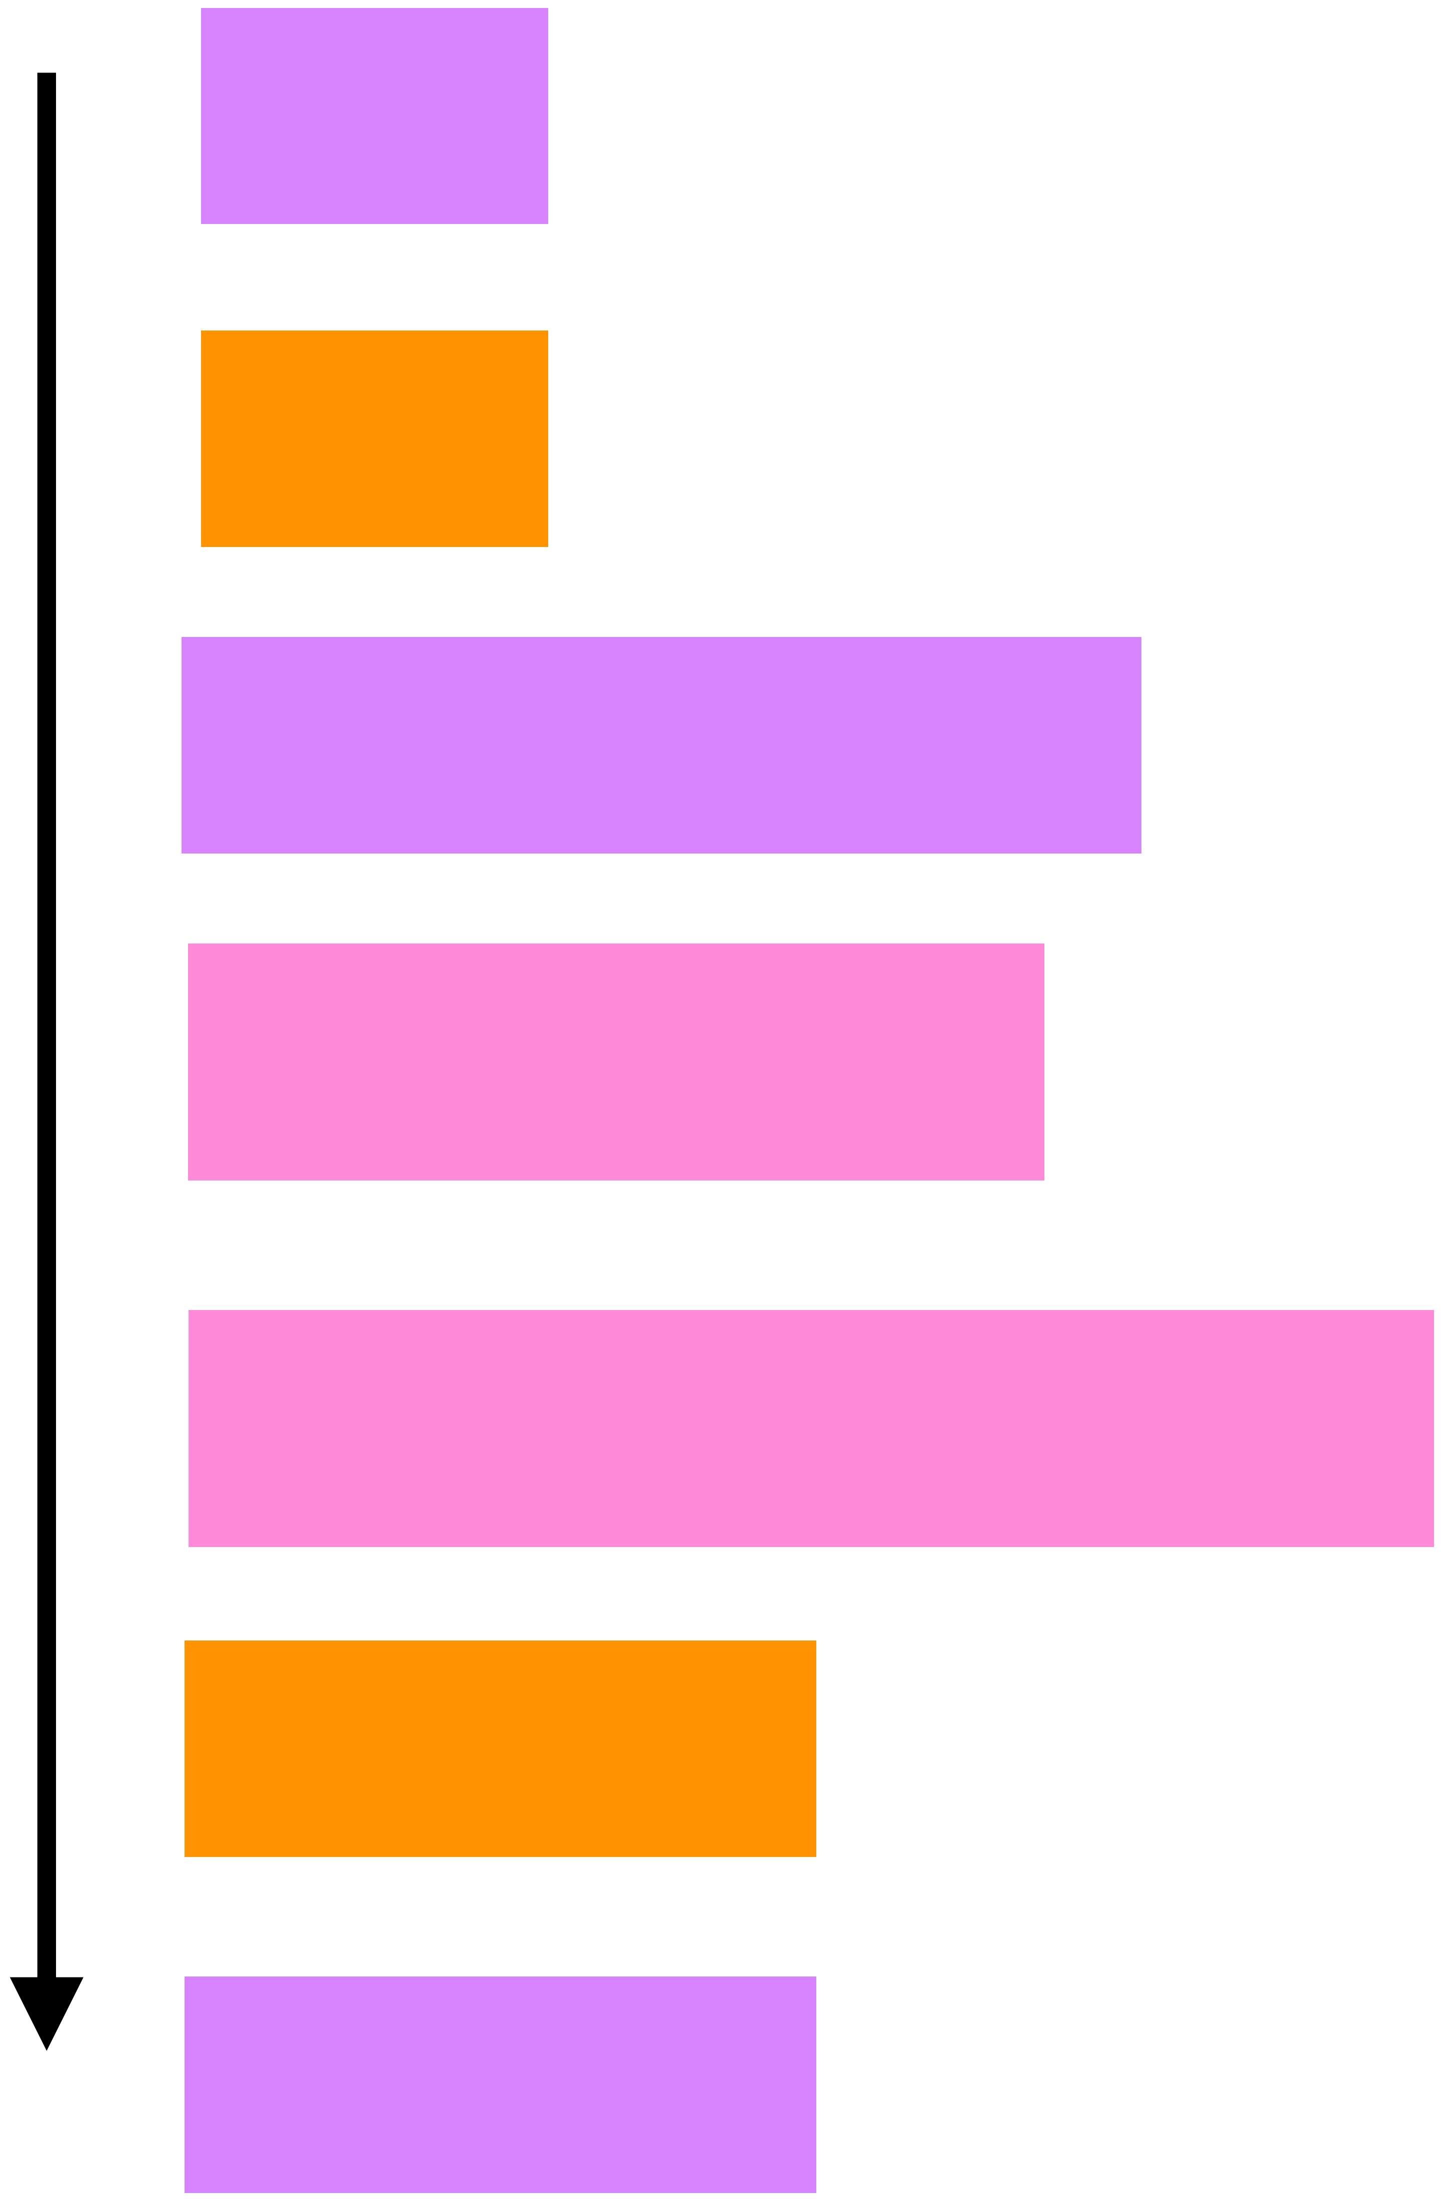

run in decreasing order of rank

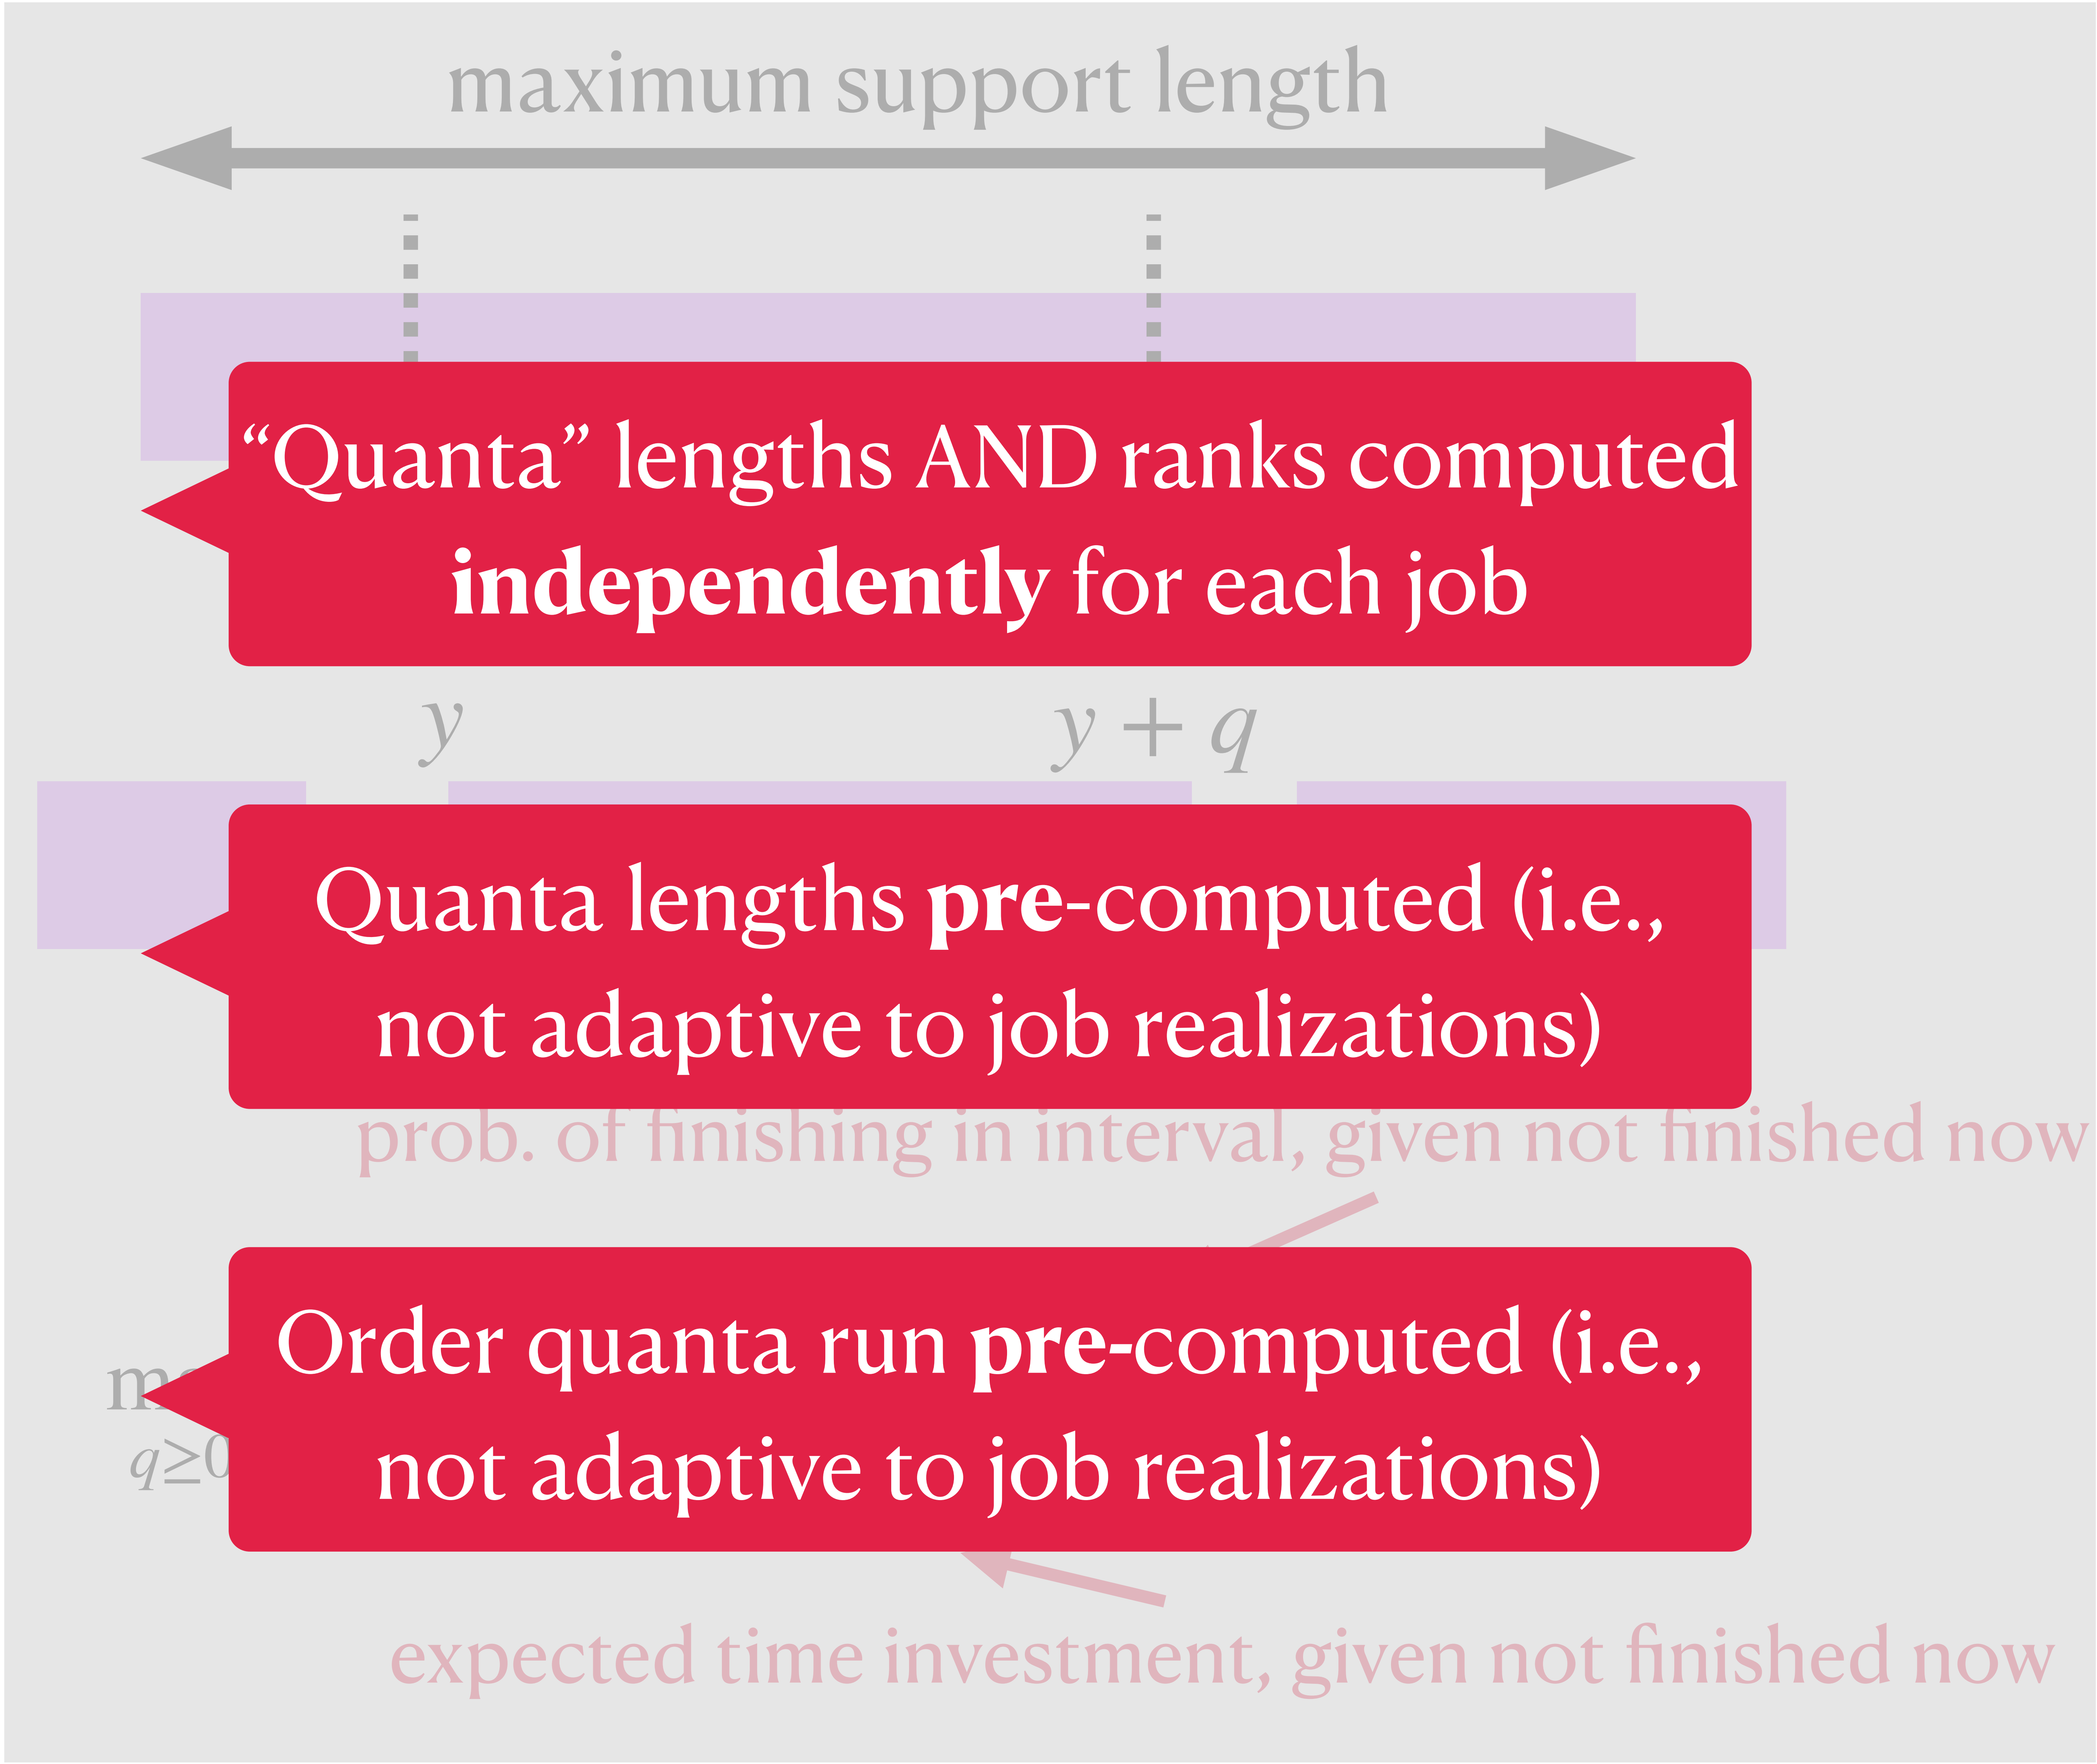

"Quanta" lengths AND ranks computed **independently** for each job

Quanta lengths **pre**-computed (i.e., not adaptive to job realizations)

Order quanta run **pre**-computed (i.e., not adaptive to job realizations)

Gittins is Not Robust

Gittins is Not Robust

Theorem. GIPP has arbitrarily poor performance even when the true distributions $\mathcal{I}^* = \{\mathcal{D}_j^*\}_{j=1}^n$ are arbitrarily "close" to the predicted distributions $\hat{\mathcal{I}} = \{\hat{\mathcal{D}}_j\}_{j=1}^n$

$$\mathrm{GIPP}(\mathcal{I}^*, \hat{\mathcal{I}}) = \Omega(n) \cdot \mathrm{GIPP}(\mathcal{I}^*, \mathcal{I}^*) = \mathrm{OPT}(\mathcal{I}^*)^\dagger$$

Cost of Gittins on true distributions, given predicted distributions (to construct quanta, etc.)

Gittins is Not Robust

Theorem. GIPP has arbitrarily poor performance even when the true distributions $\mathcal{I}^* = \{\mathcal{D}_j^*\}_{j=1}^n$ are arbitrarily “close” to the predicted distributions $\hat{\mathcal{I}} = \{\hat{\mathcal{D}}_j\}_{j=1}^n$

$$\text{GIPP}(\mathcal{I}^*, \hat{\mathcal{I}}) = \Omega(n) \cdot \text{GIPP}(\mathcal{I}^*, \mathcal{I}^*) = \text{OPT}(\mathcal{I}^*)$$ [†]

Cost of Gittins on true distributions, given predicted distributions (to construct quanta, etc.)

[†] Restricting to instances where $\text{GIPP}(\mathcal{I}^*, \hat{\mathcal{I}})$ completes all jobs of $\mathcal{I}^*$

Gittins is Not Robust

Theorem. GIPP has arbitrarily poor performance even when the true distributions $\mathcal{J}^* = \{\mathcal{D}_j^*\}_{j=1}^n$ are arbitrarily “close” to the predicted distributions $\hat{\mathcal{J}} = \{\hat{\mathcal{D}}_j\}_{j=1}^n$

$$\mathrm{GIPP}(\mathcal{J}^*, \hat{\mathcal{J}}) = \Omega(n) \cdot \mathrm{GIPP}(\mathcal{J}^*, \mathcal{J}^*) = \mathrm{OPT}(\mathcal{J}^*)^\dagger$$

Cost of Gittins on true distributions, given predicted distributions (to construct quanta, etc.)

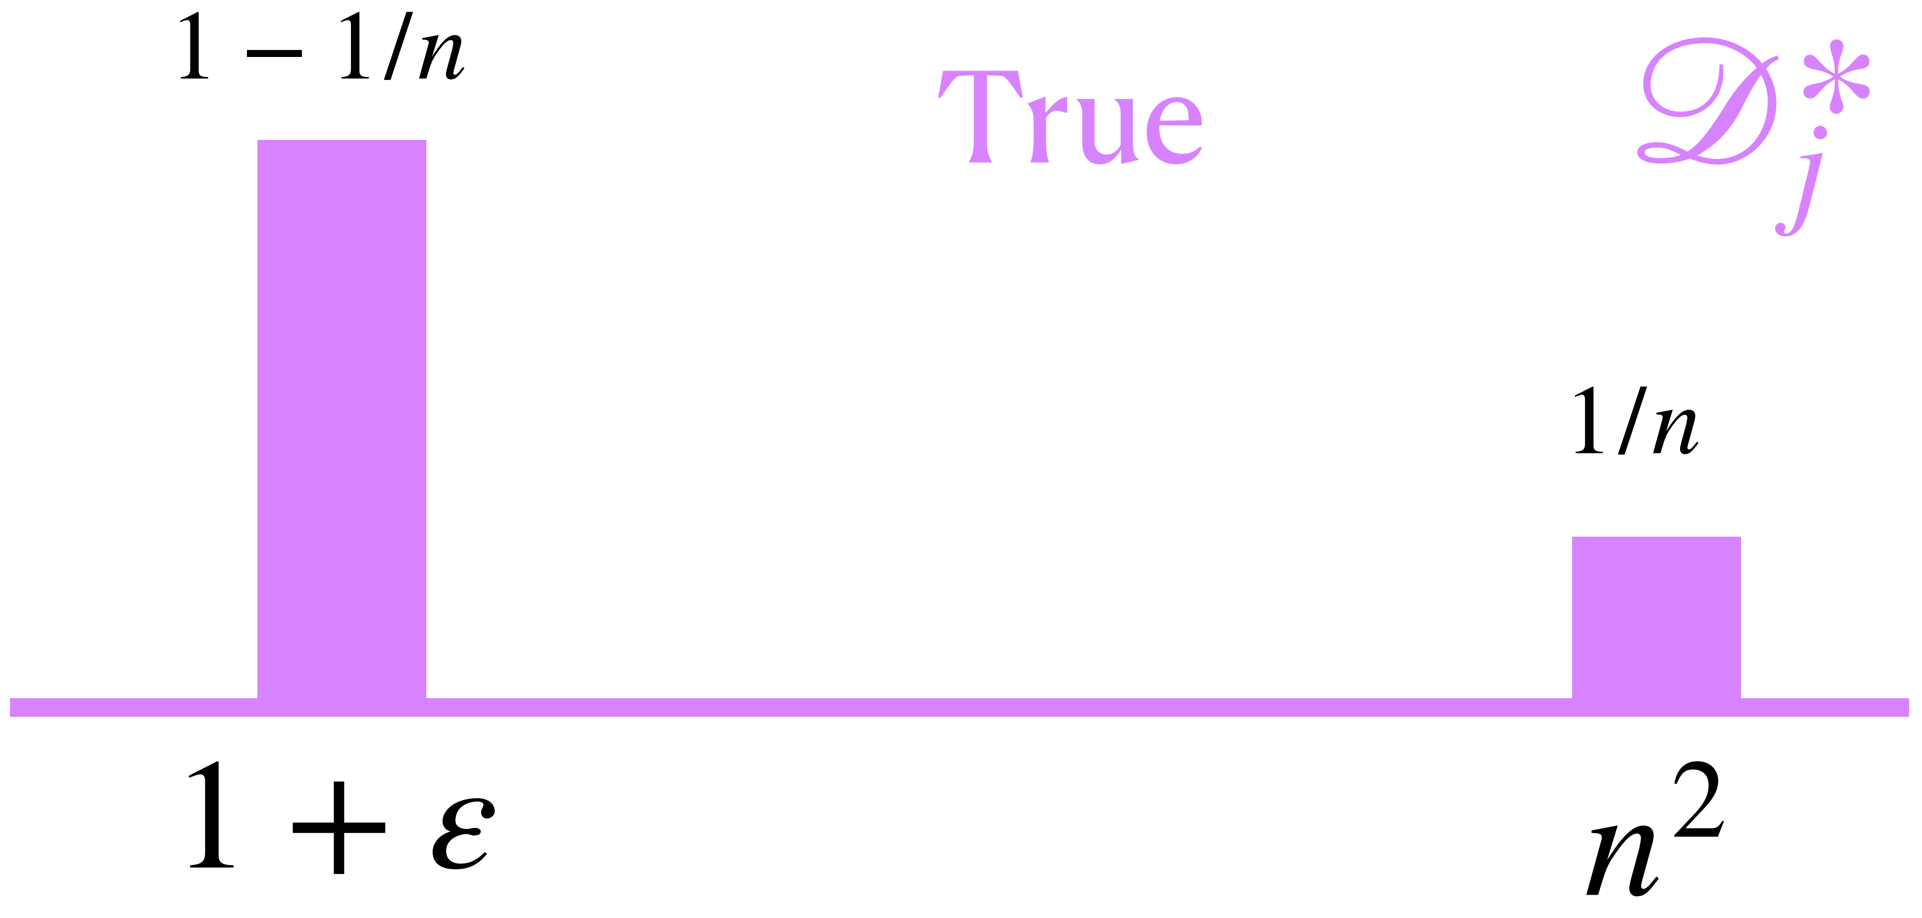

[†] Restricting to instances where $\mathrm{GIPP}(\mathcal{J}^*, \hat{\mathcal{J}})$ completes all jobs of $\mathcal{J}^*$

Gittins is Not Robust

Theorem. GIPP has arbitrarily poor performance even when the true distributions $\mathcal{J}^* = \{\mathcal{D}_j^*\}_{j=1}^n$ are arbitrarily “close” to the predicted distributions $\hat{\mathcal{J}} = \{\hat{\mathcal{D}}_j\}_{j=1}^n$

$$\mathrm{GIPP}(\mathcal{J}^*, \hat{\mathcal{J}}) = \Omega(n) \cdot \mathrm{GIPP}(\mathcal{J}^*, \mathcal{J}^*) = \mathrm{OPT}(\mathcal{J}^*)^\dagger$$

Cost of Gittins on true distributions, given predicted distributions (to construct quanta, etc.)

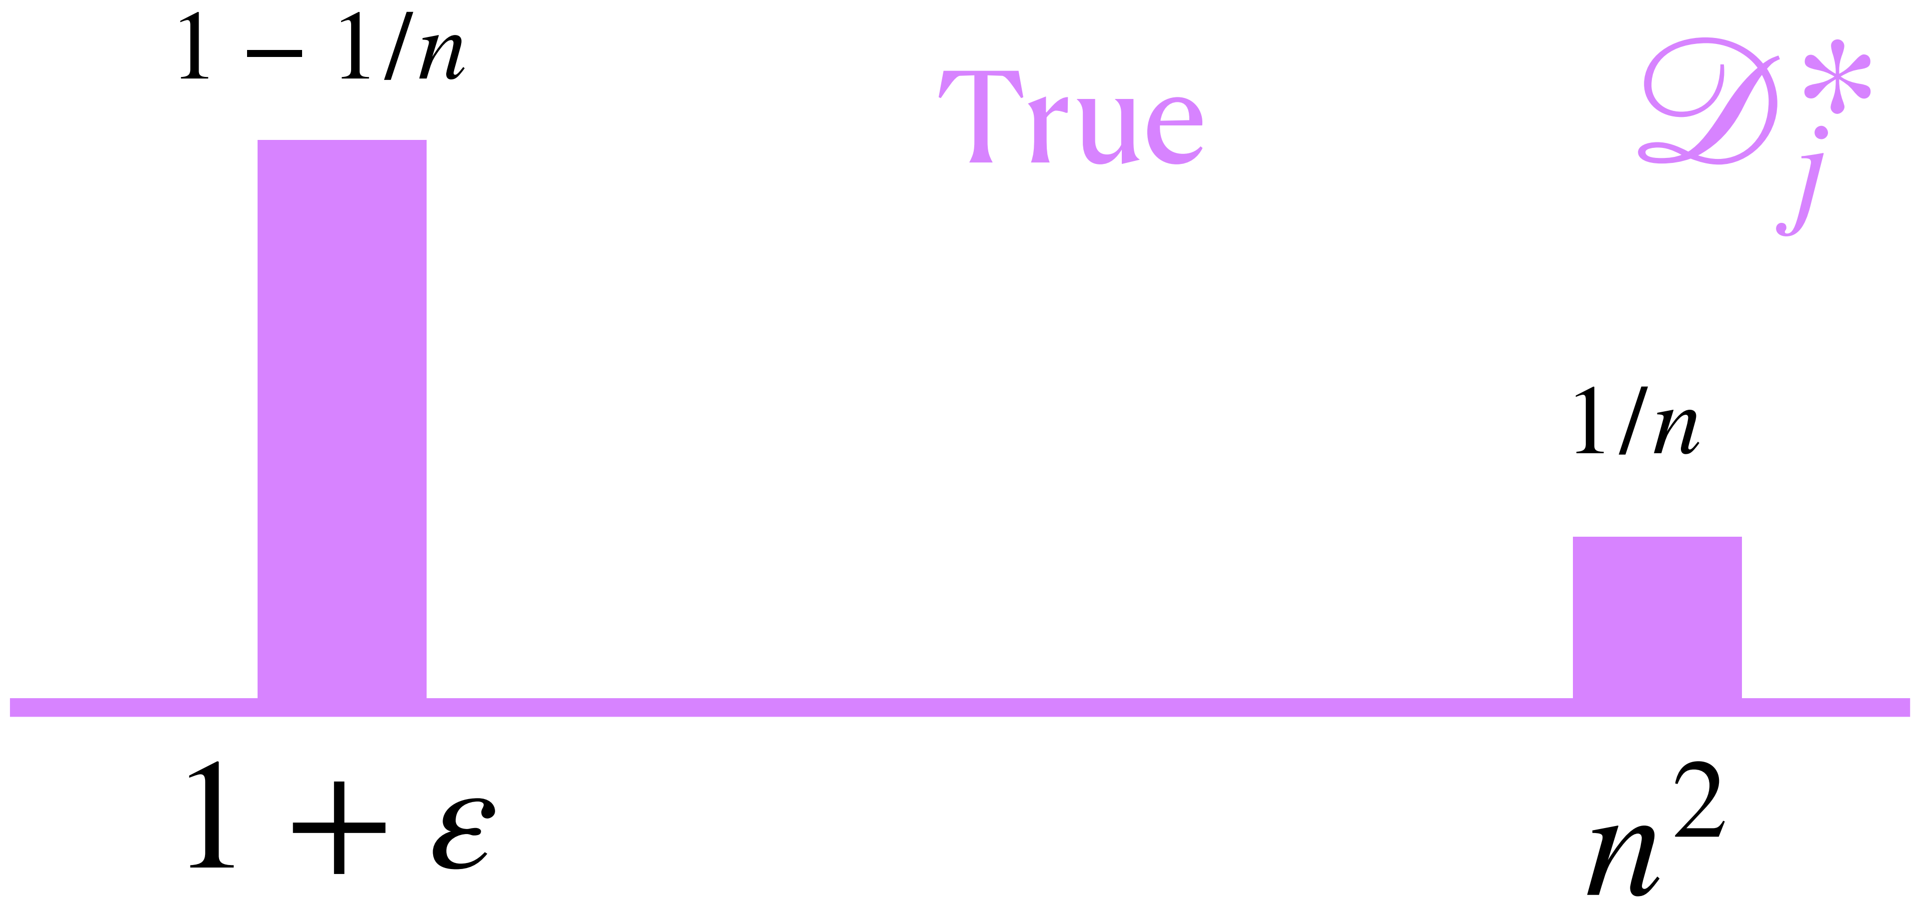

[†] Restricting to instances where $\mathrm{GIPP}(\mathcal{J}^*, \hat{\mathcal{J}})$ completes all jobs of $\mathcal{J}^*$

Gittins is Not Robust

$\mathrm{GIPP}(\mathcal{I}^*, \hat{\mathcal{I}}) = \Omega(n) \cdot \mathrm{OPT}(\mathcal{I}^*)$

n i.i.d. distributions

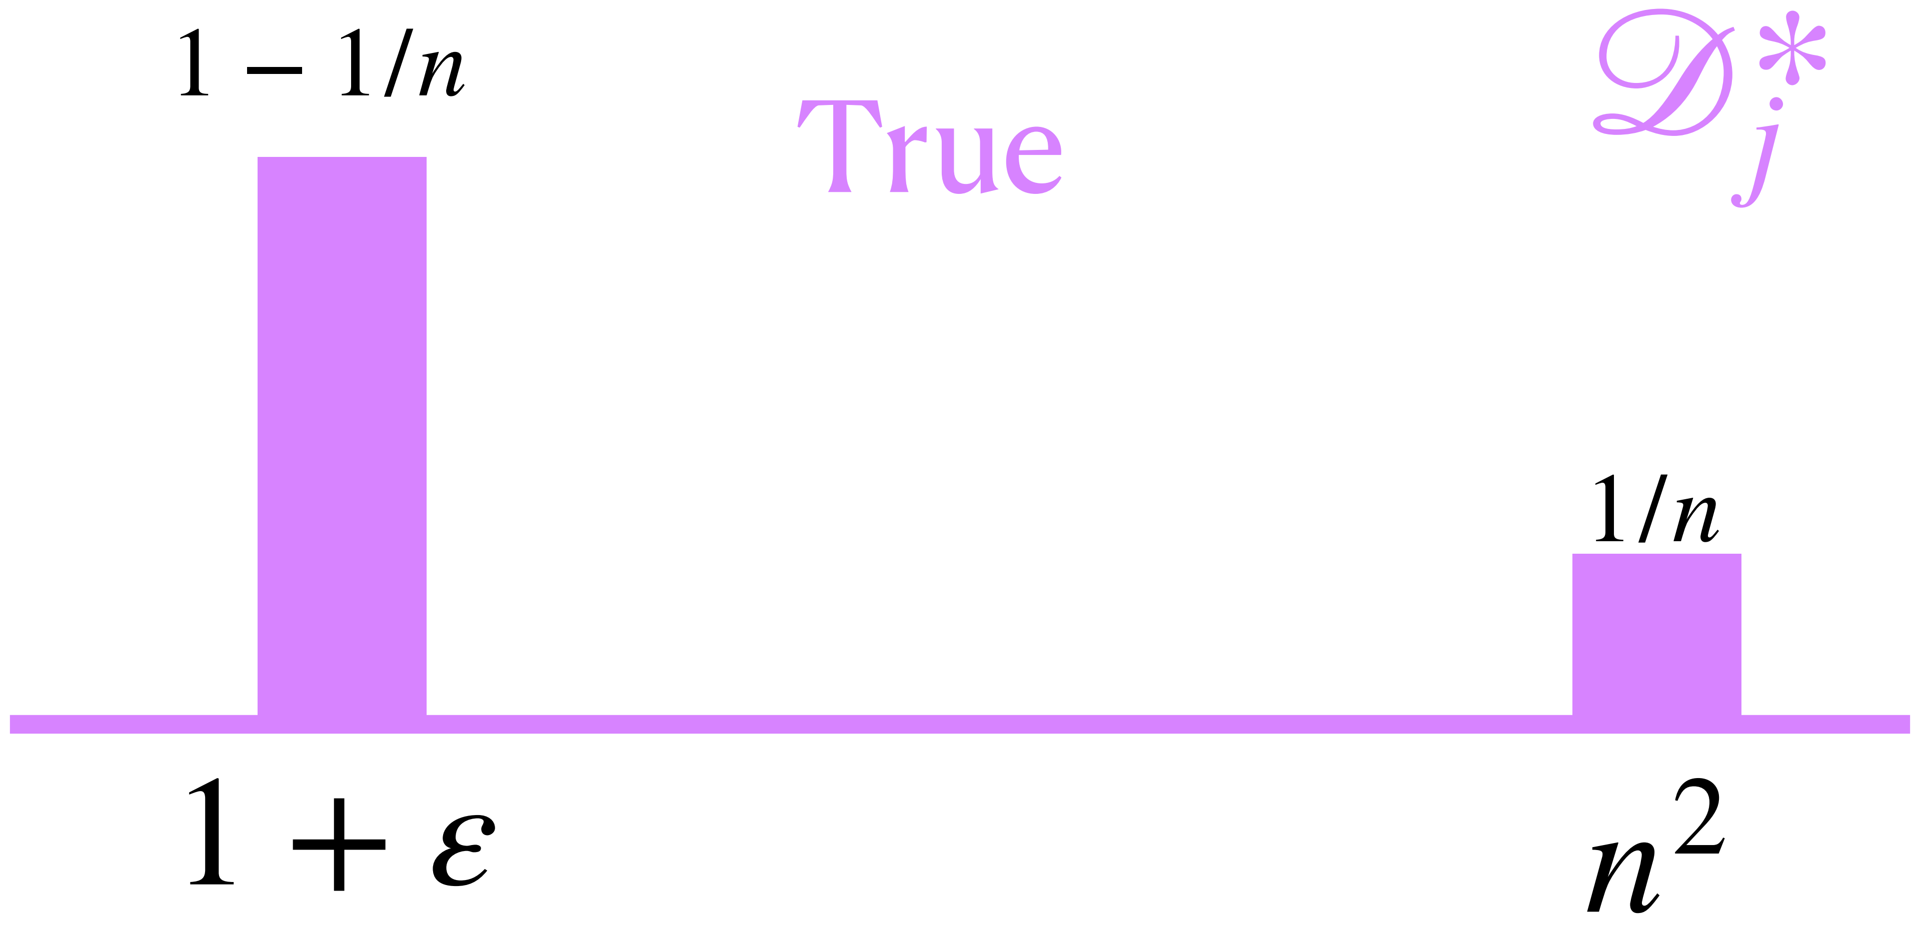

Gittins is Not Robust

$\mathrm{GIPP}(\mathcal{J}^*, \hat{\mathcal{J}}) = \Omega(n) \cdot \mathrm{OPT}(\mathcal{J}^*)$

n i.i.d. distributions

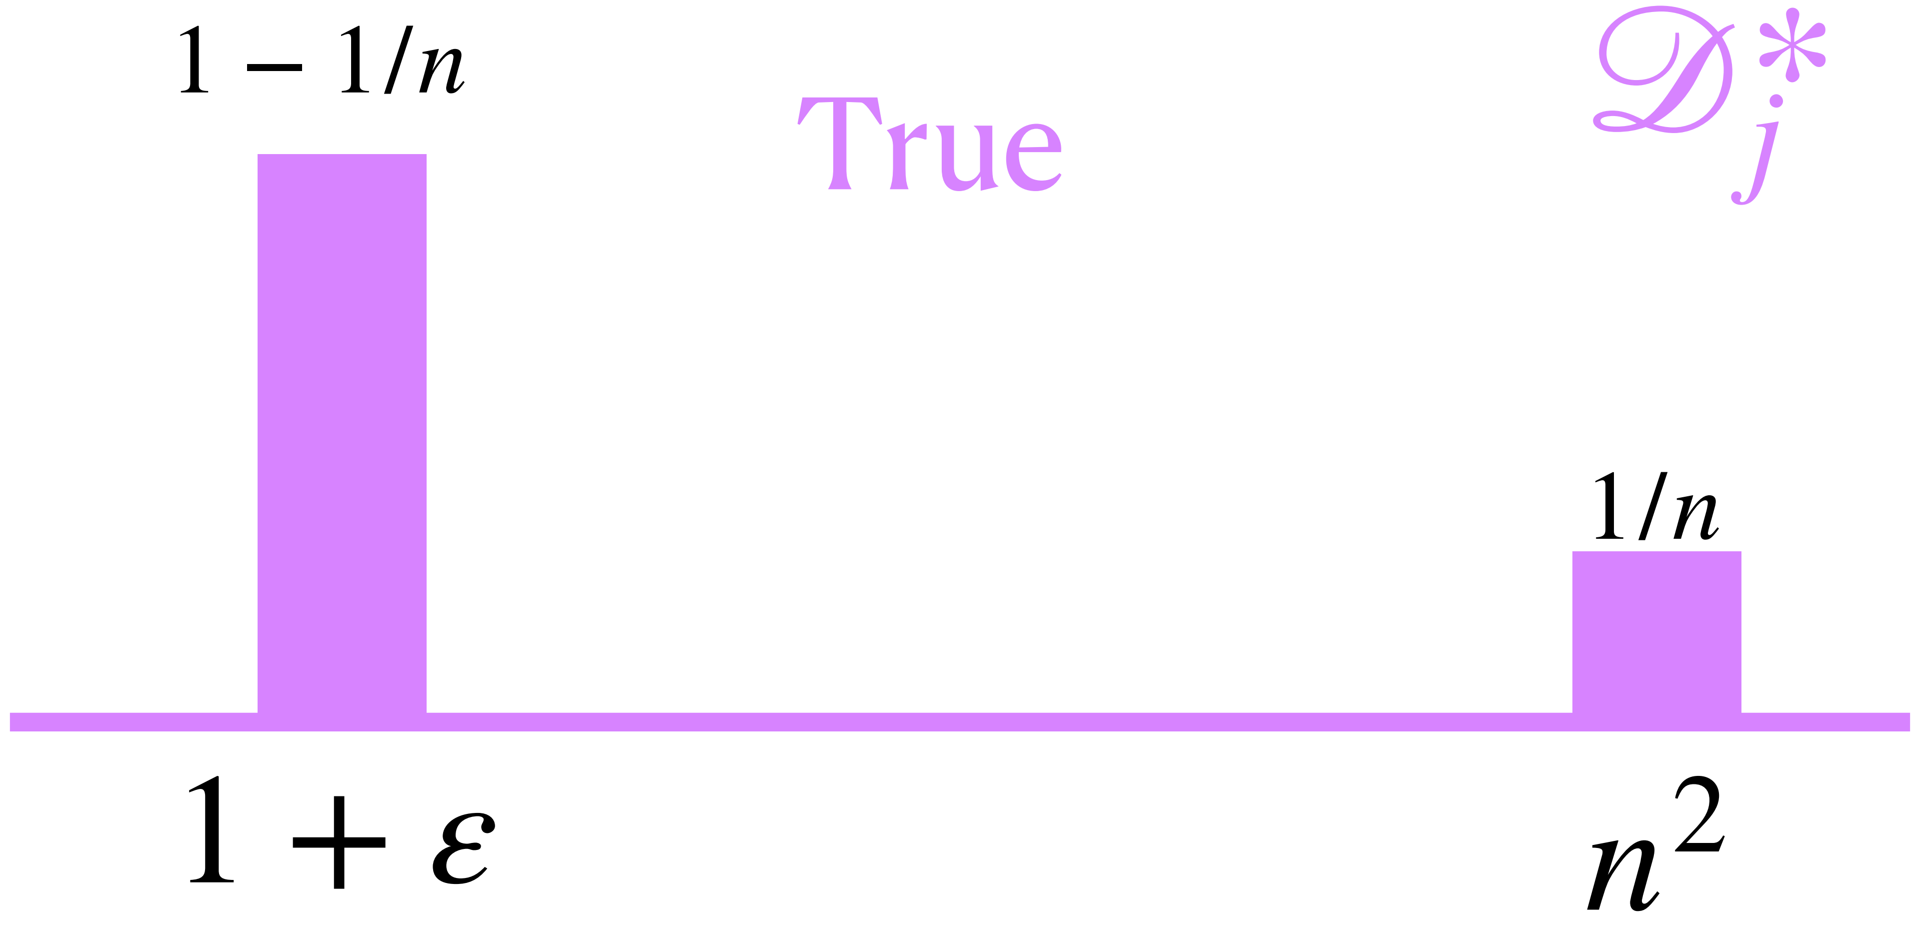

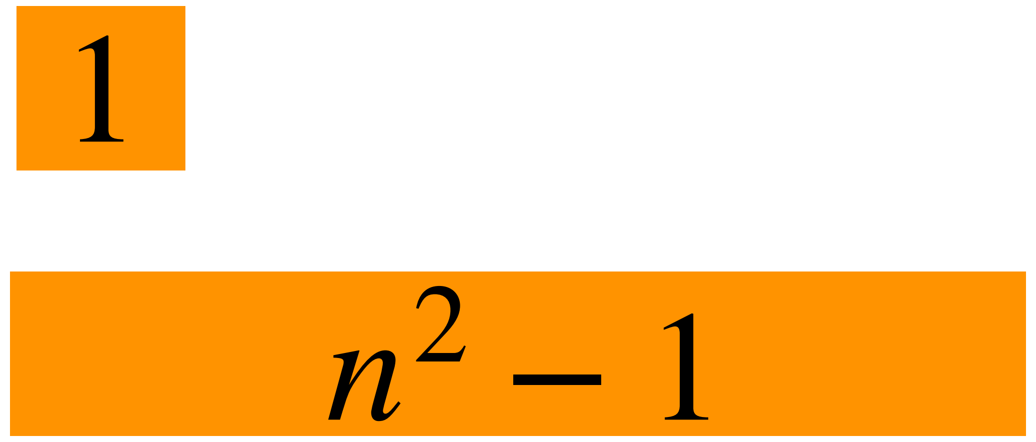

Gittins is Not Robust

$\mathrm{GIPP}(\mathscr{I}^*, \hat{\mathscr{I}}) = \Omega(n) \cdot \mathrm{OPT}(\mathscr{I}^*)$

n i.i.d. distributions

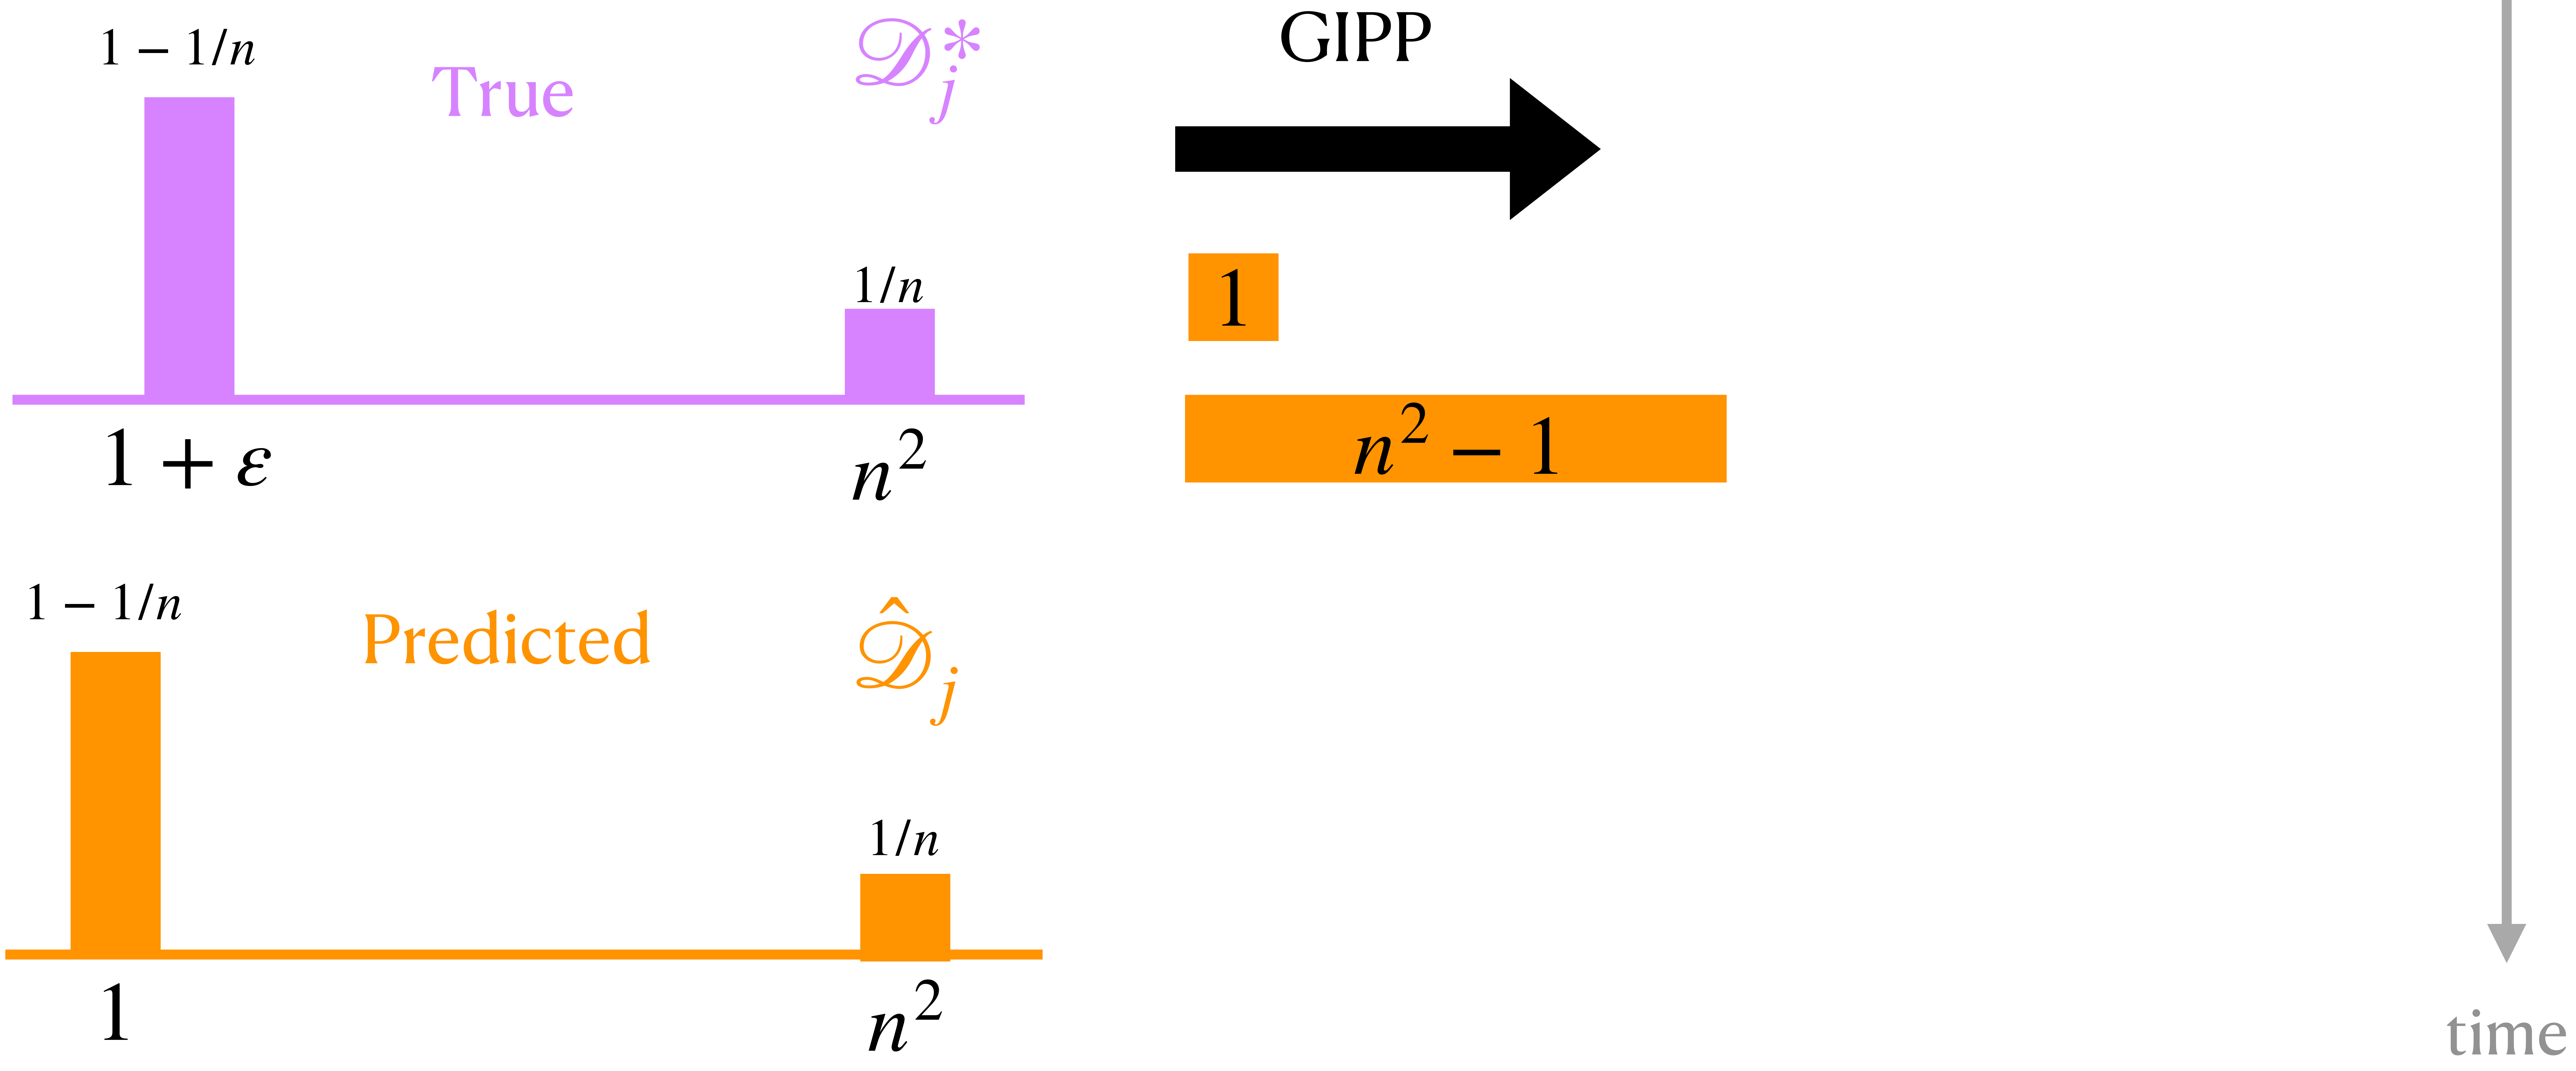

Gittins is Not Robust

$$\mathrm{GIPP}(\mathscr{I}^*, \hat{\mathscr{I}}) = \Omega(n) \cdot \mathrm{OPT}(\mathscr{I}^*)$$

n i.i.d. distributions

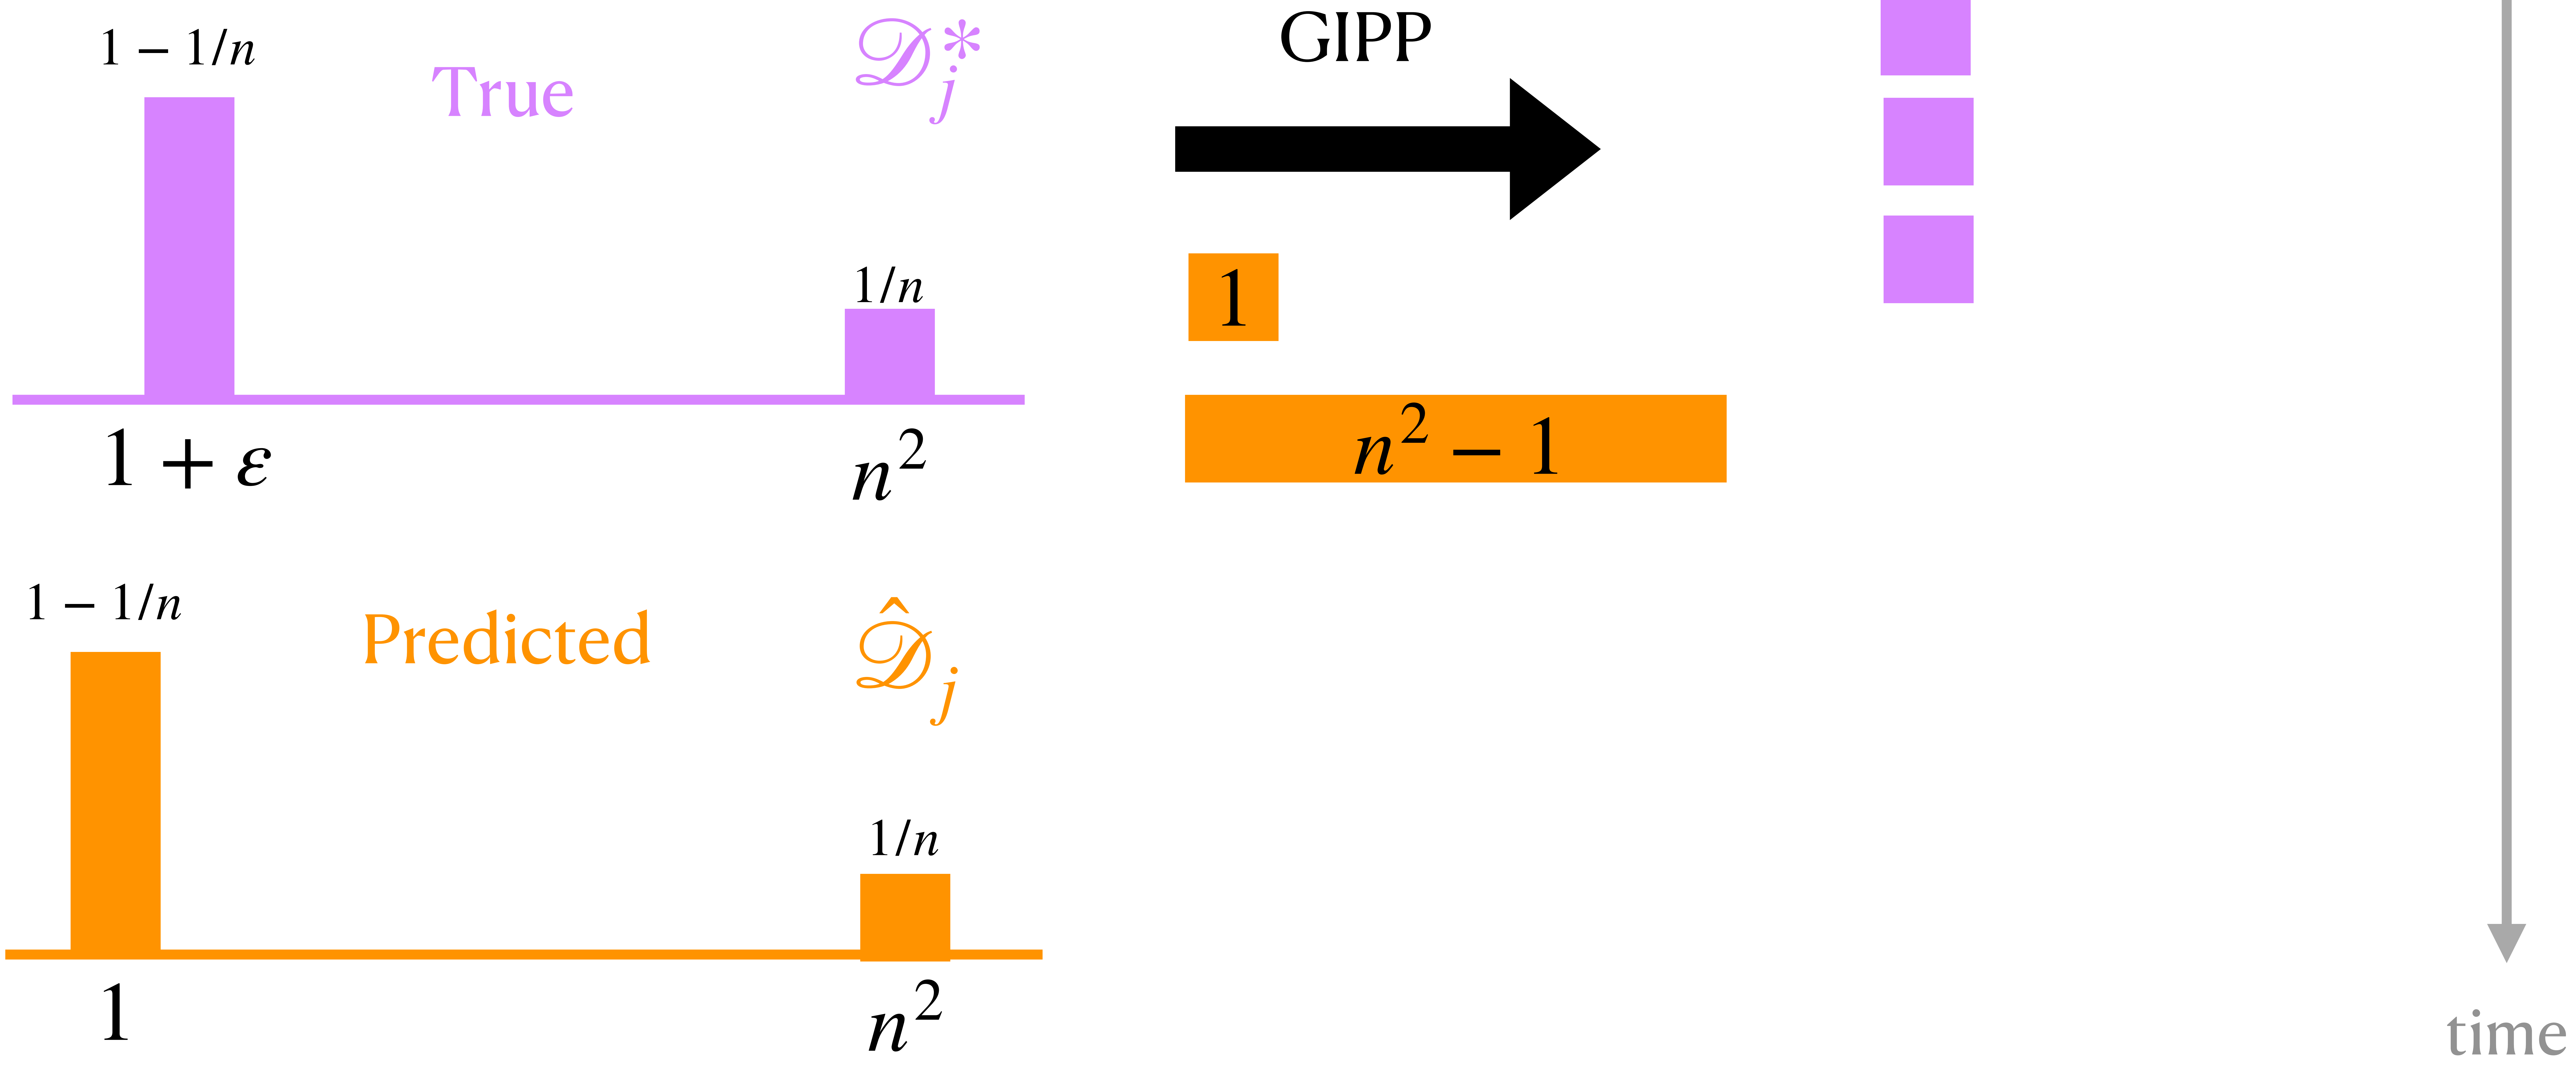

Gittins is Not Robust

$\mathrm{GIPP}(\mathcal{I}^*, \hat{\mathcal{I}}) = \Omega(n) \cdot \mathrm{OPT}(\mathcal{I}^*)$

n i.i.d. distributions

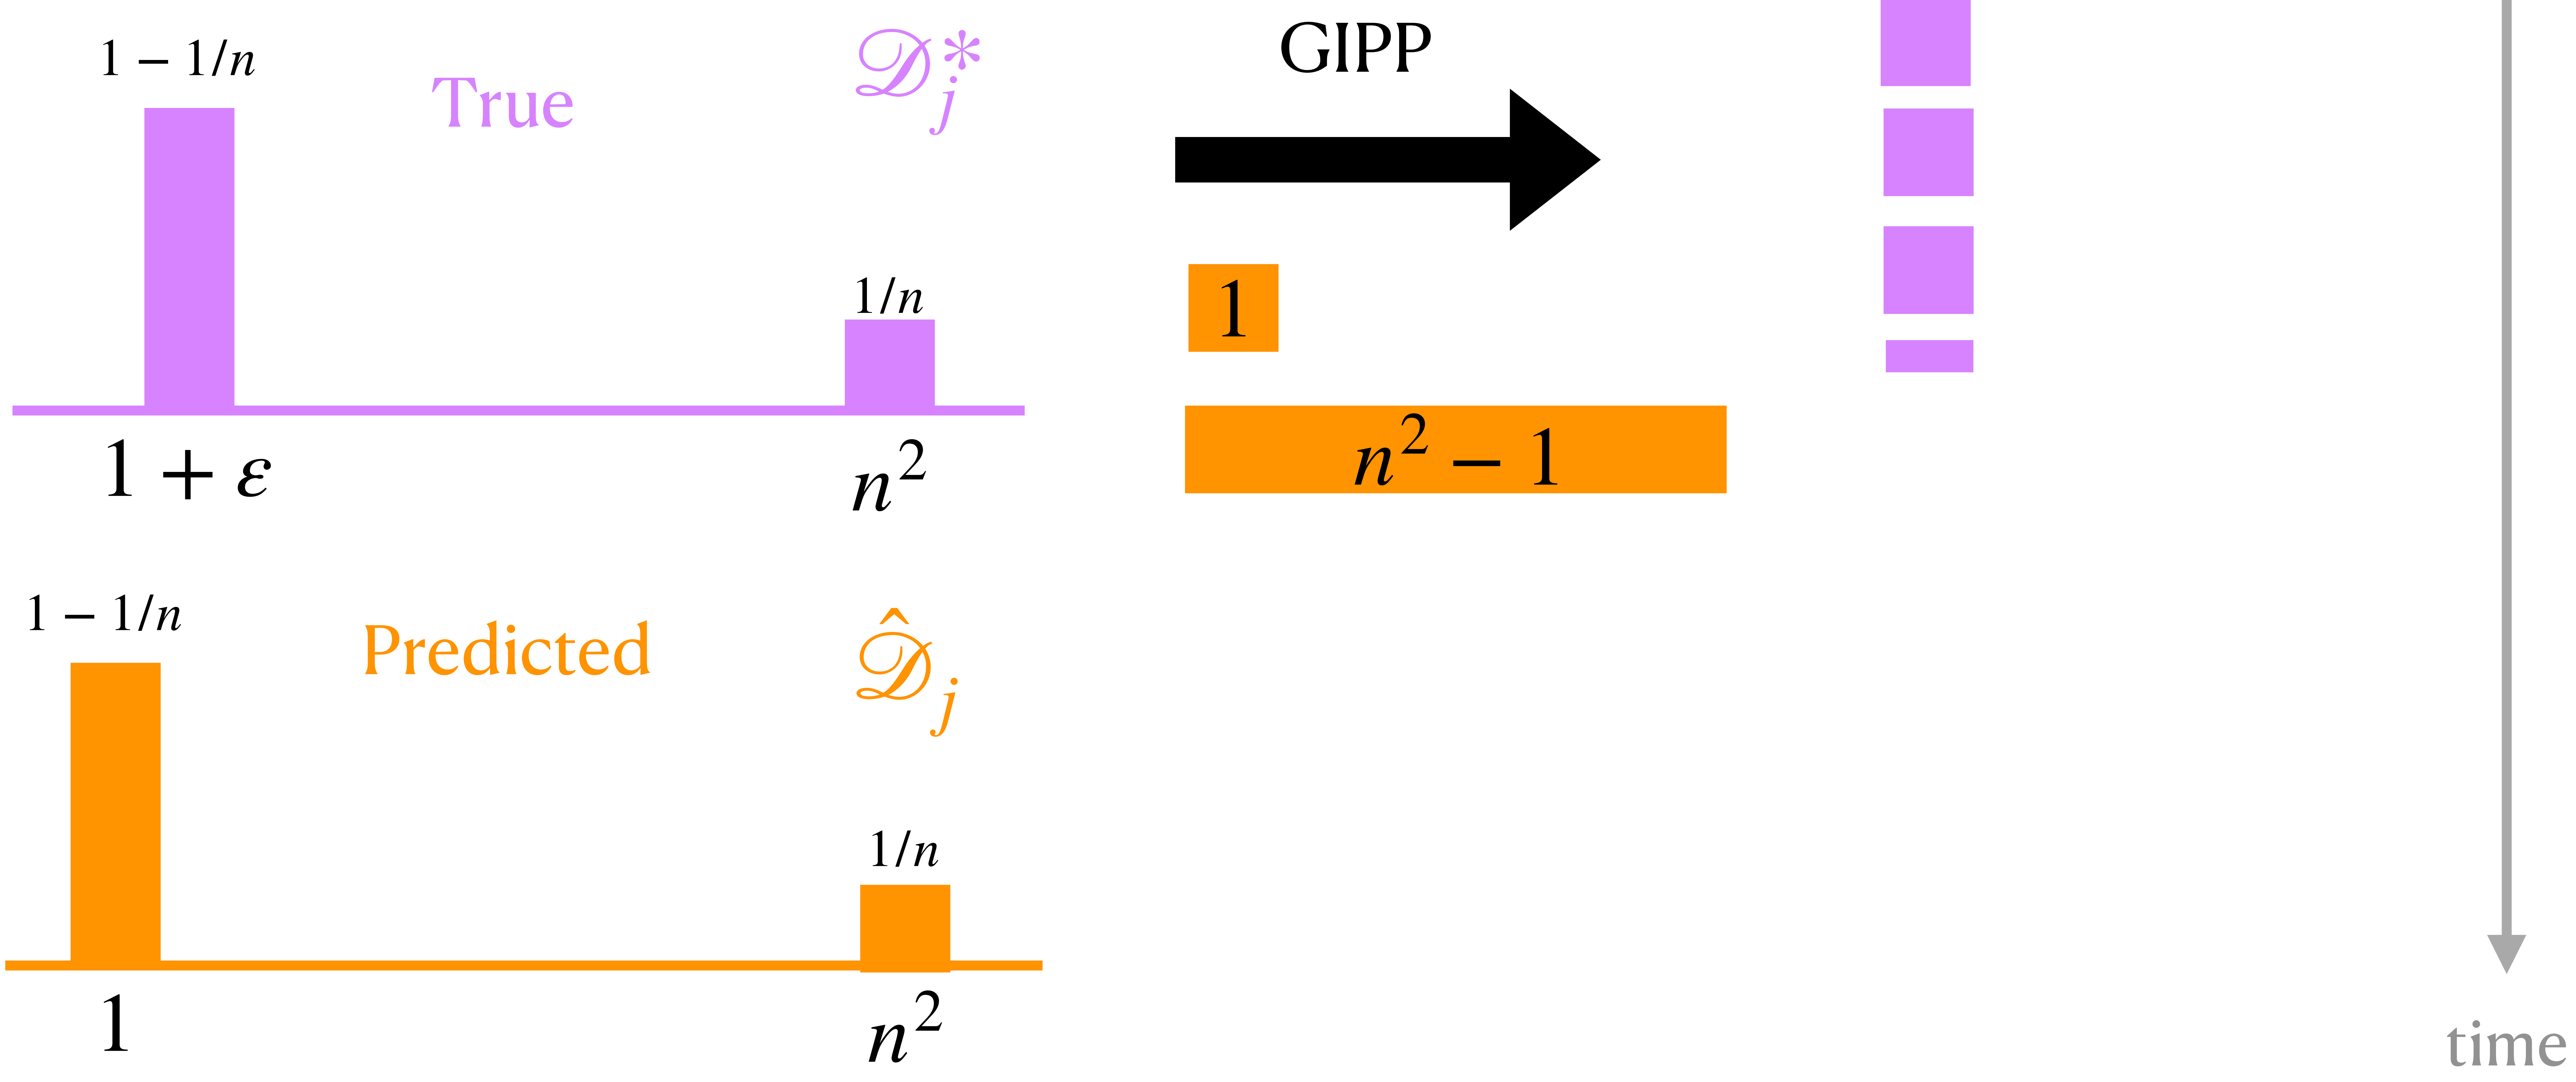

$1 - 1/n$

True

$\mathcal{D}_j^*$

$1/n$

$1 + \varepsilon$

n^2

$1 - 1/n$

Predicted

$\hat{\mathcal{D}}_j$

$1/n$

1

n^2

GIPP

1

$n^2 - 1$

time

Gittins is Not Robust

$\mathrm{GIPP}(\mathscr{J}^*, \hat{\mathscr{J}}) = \Omega(n) \cdot \mathrm{OPT}(\mathscr{J}^*)$

n i.i.d. distributions

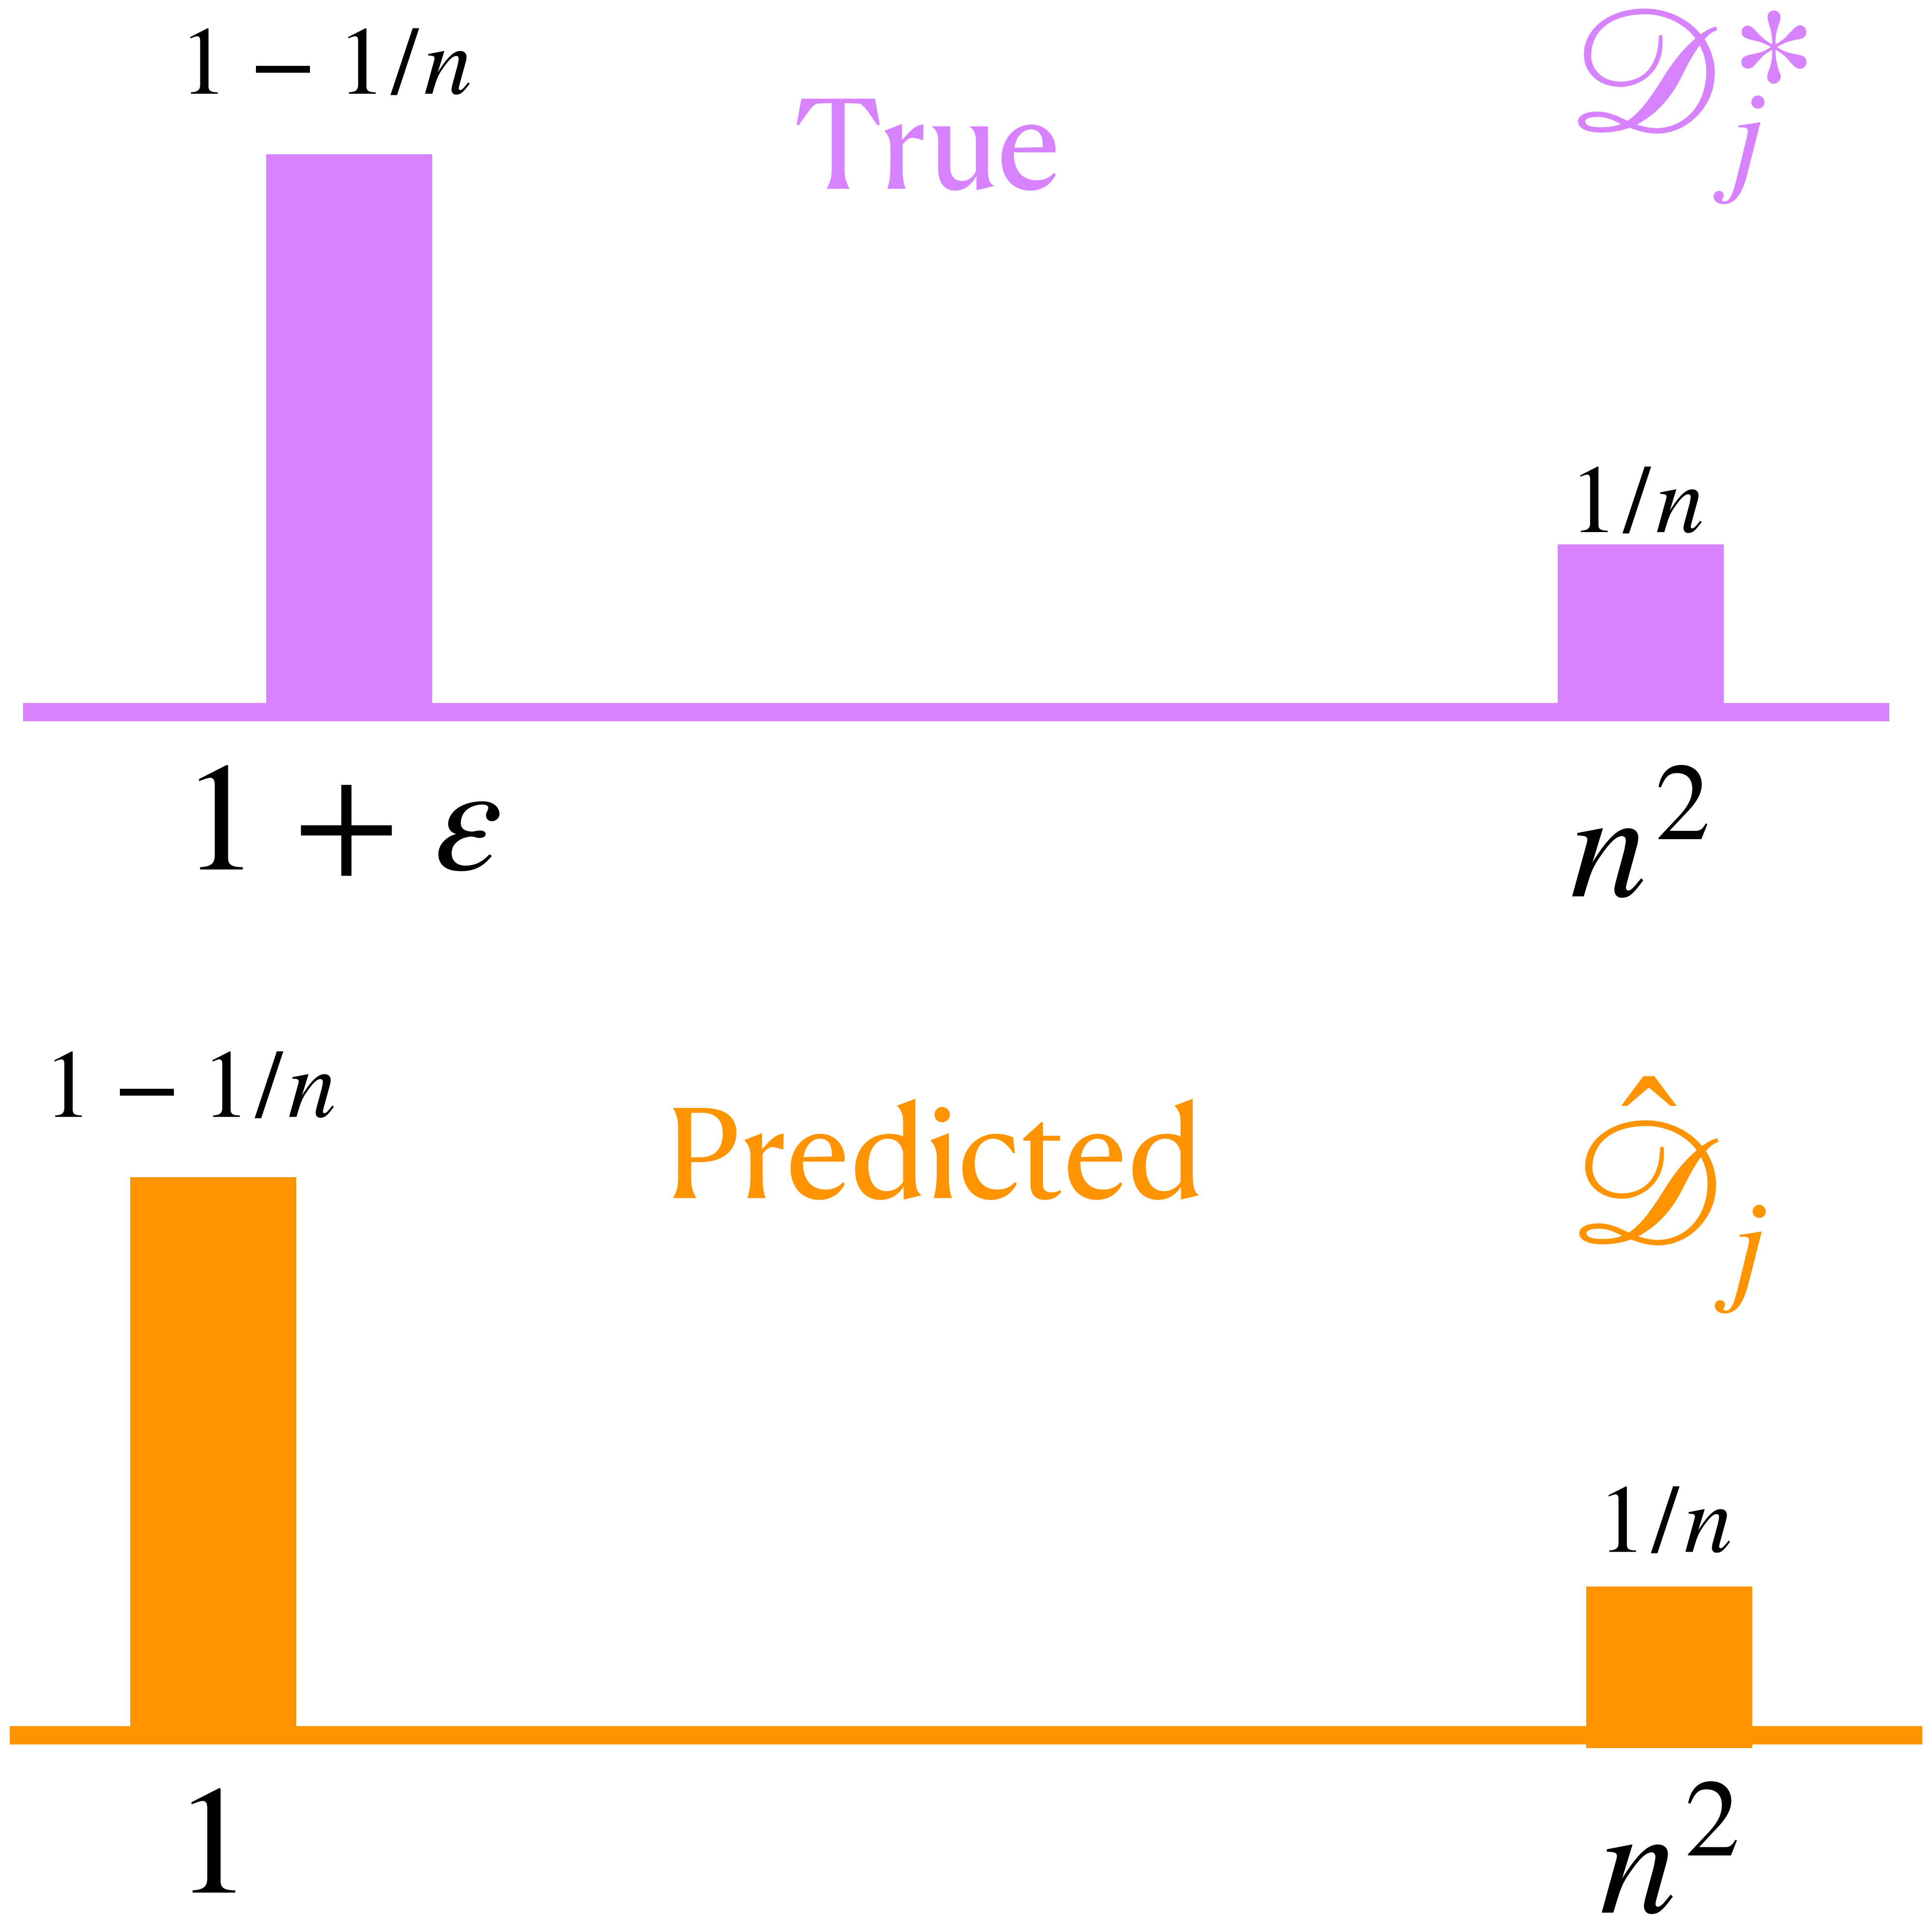

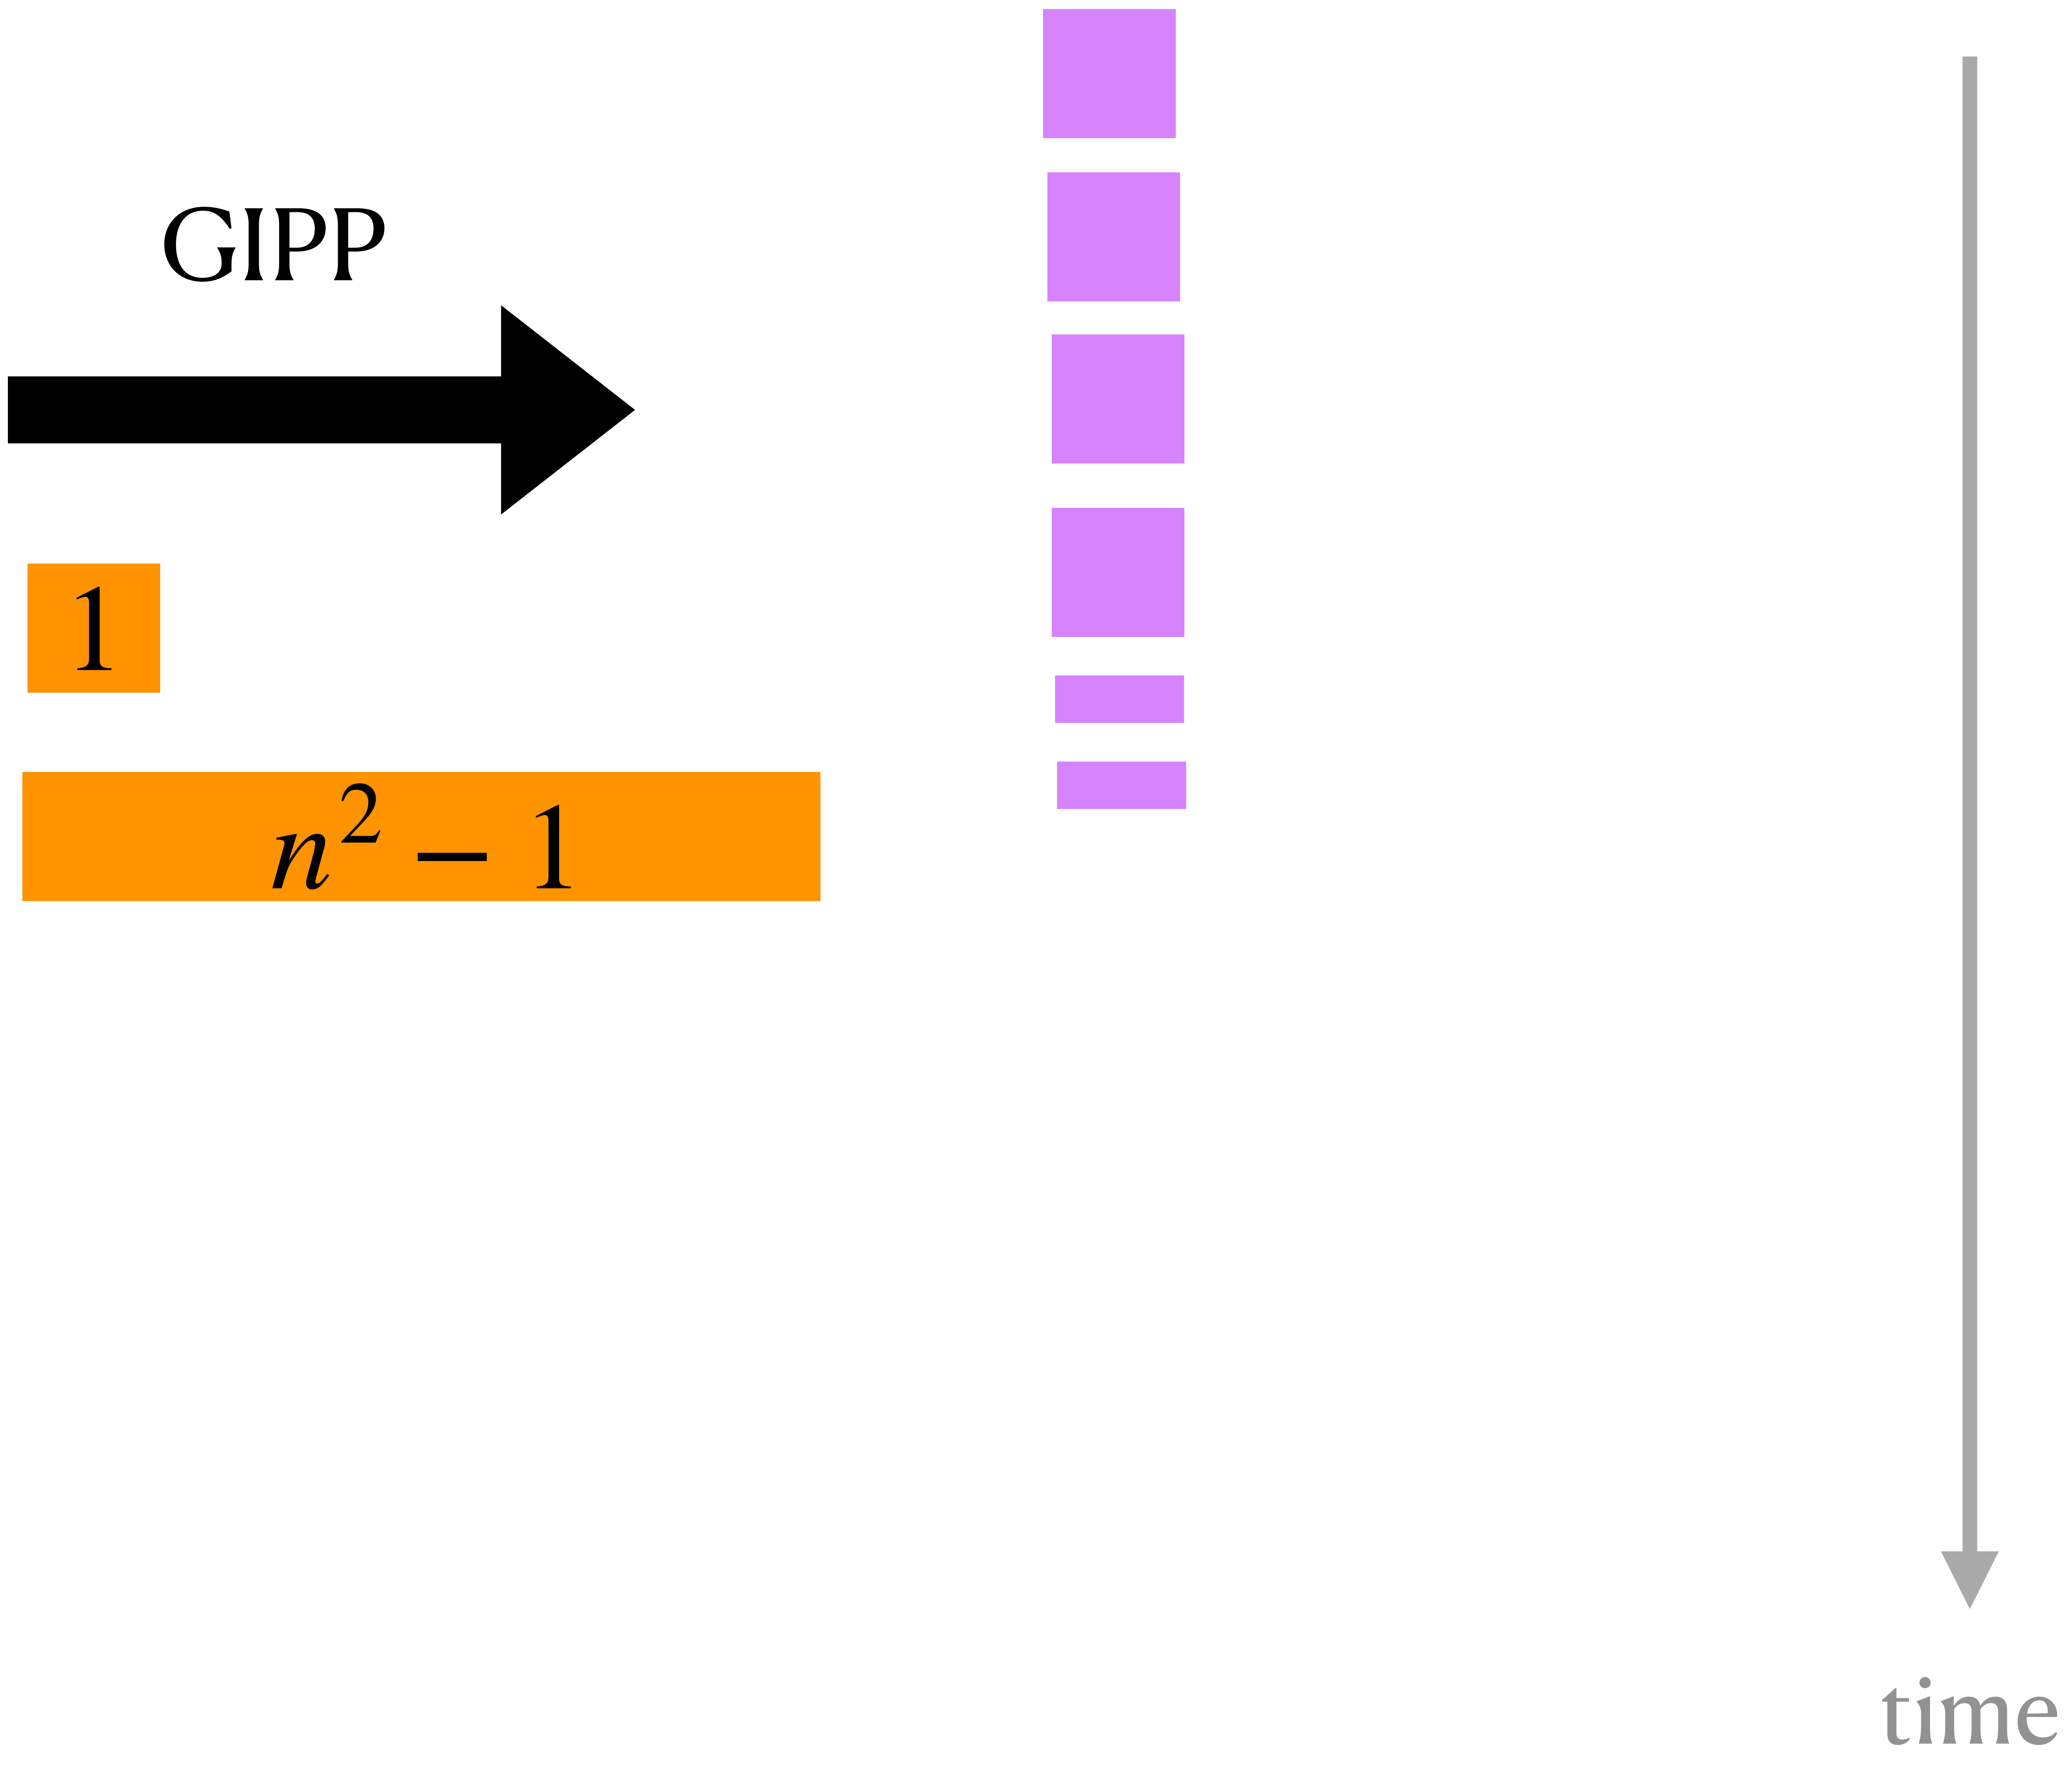

Gittins is Not Robust

$\text{GIPP}(\mathscr{J}^*, \hat{\mathscr{J}}) = \Omega(n) \cdot \text{OPT}(\mathscr{J}^*)$

n i.i.d. distributions

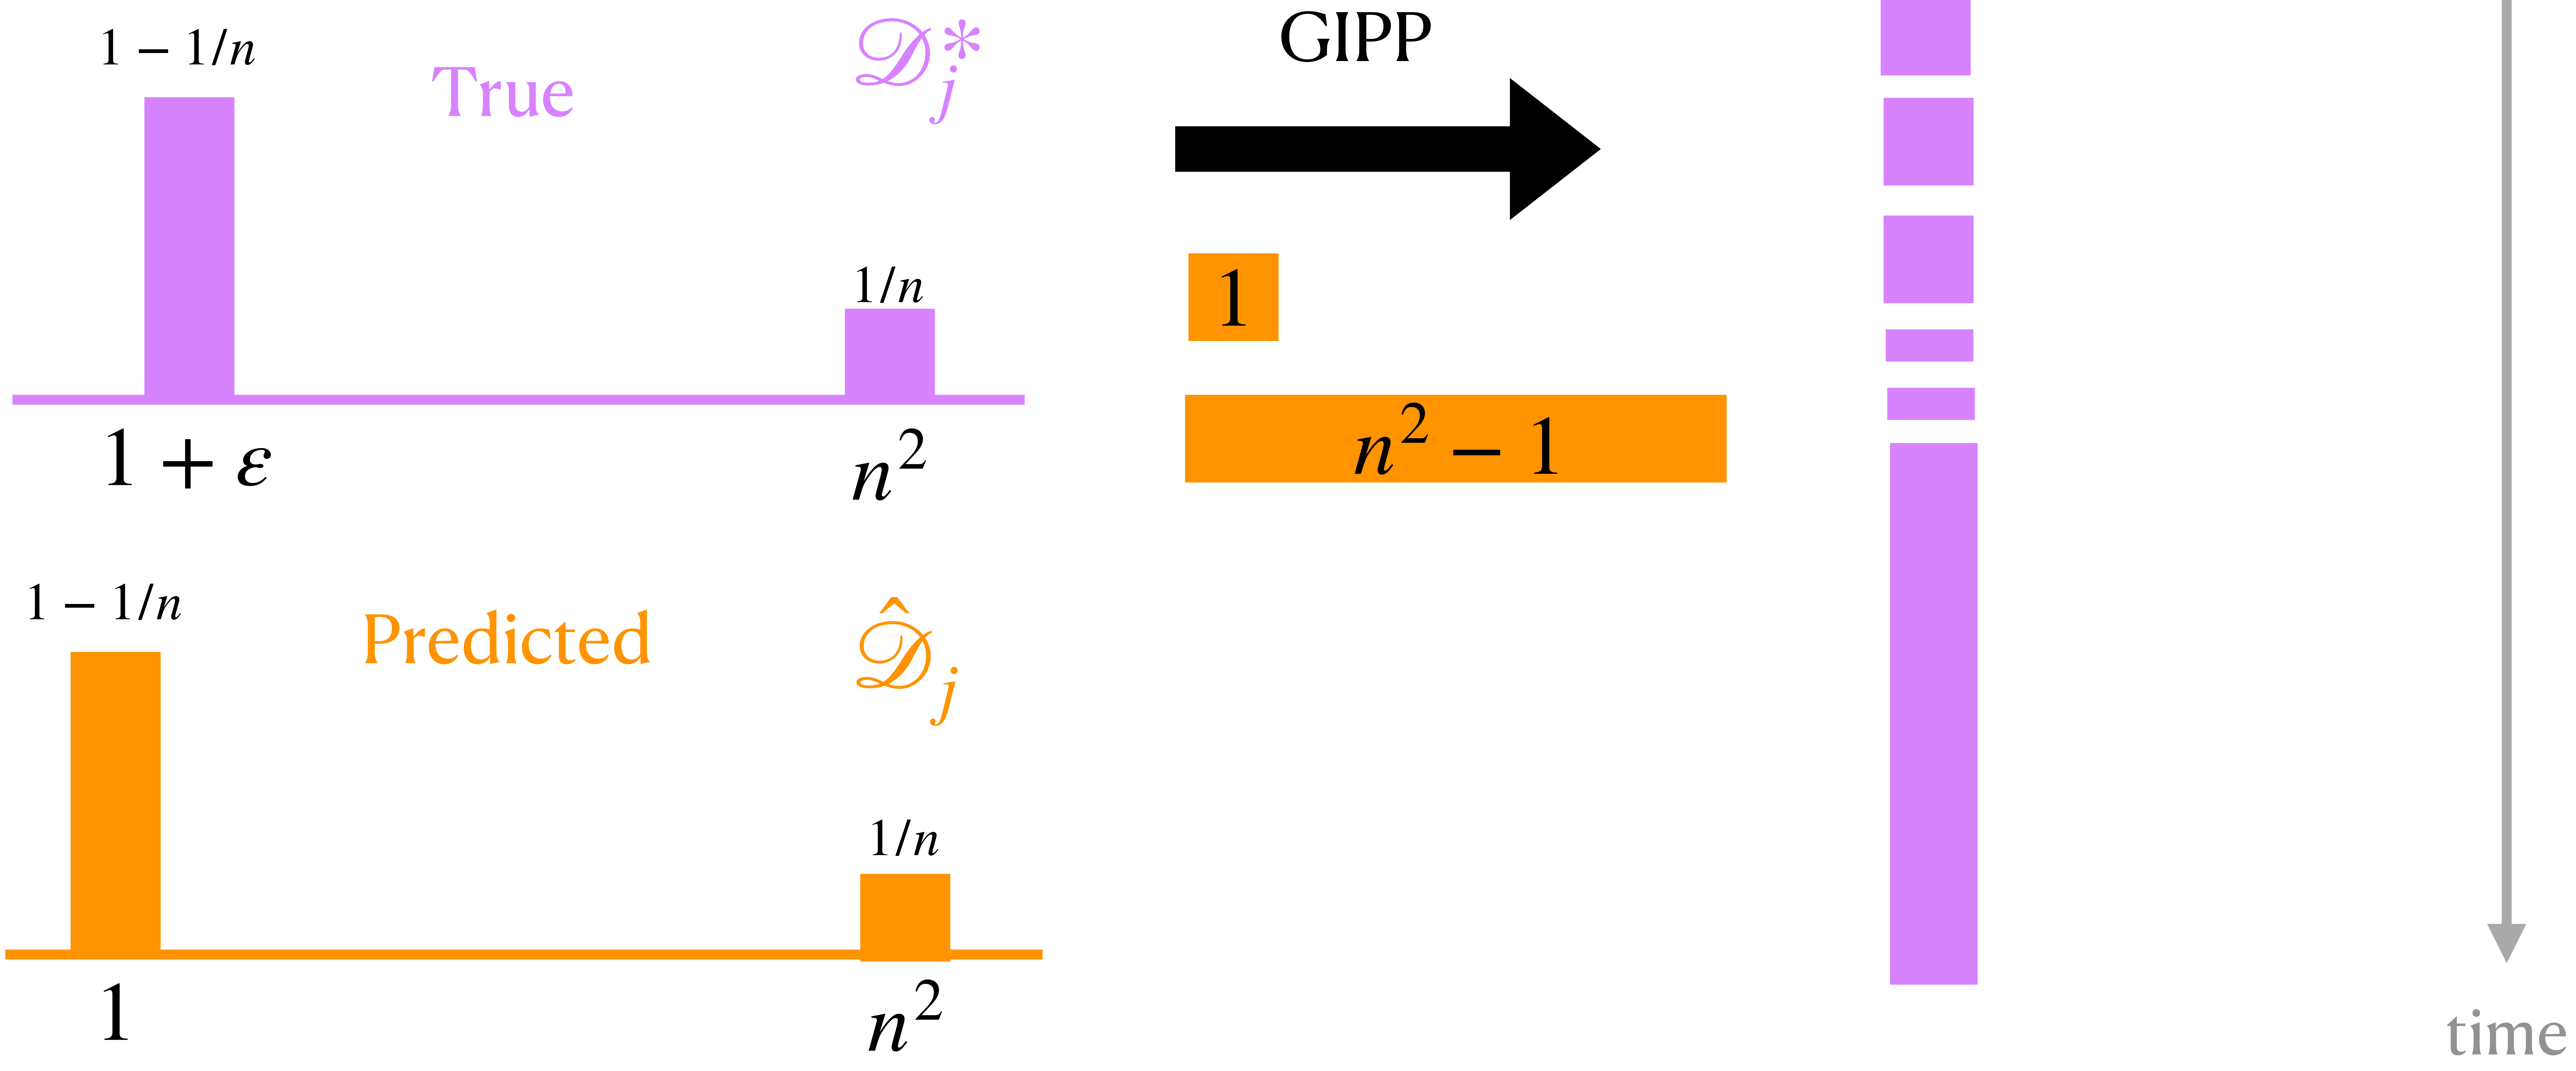

Gittins is Not Robust

$\mathrm{GIPP}(\mathcal{J}^*, \hat{\mathcal{J}}) = \Omega(n) \cdot \mathrm{OPT}(\mathcal{J}^*)$

n i.i.d. distributions

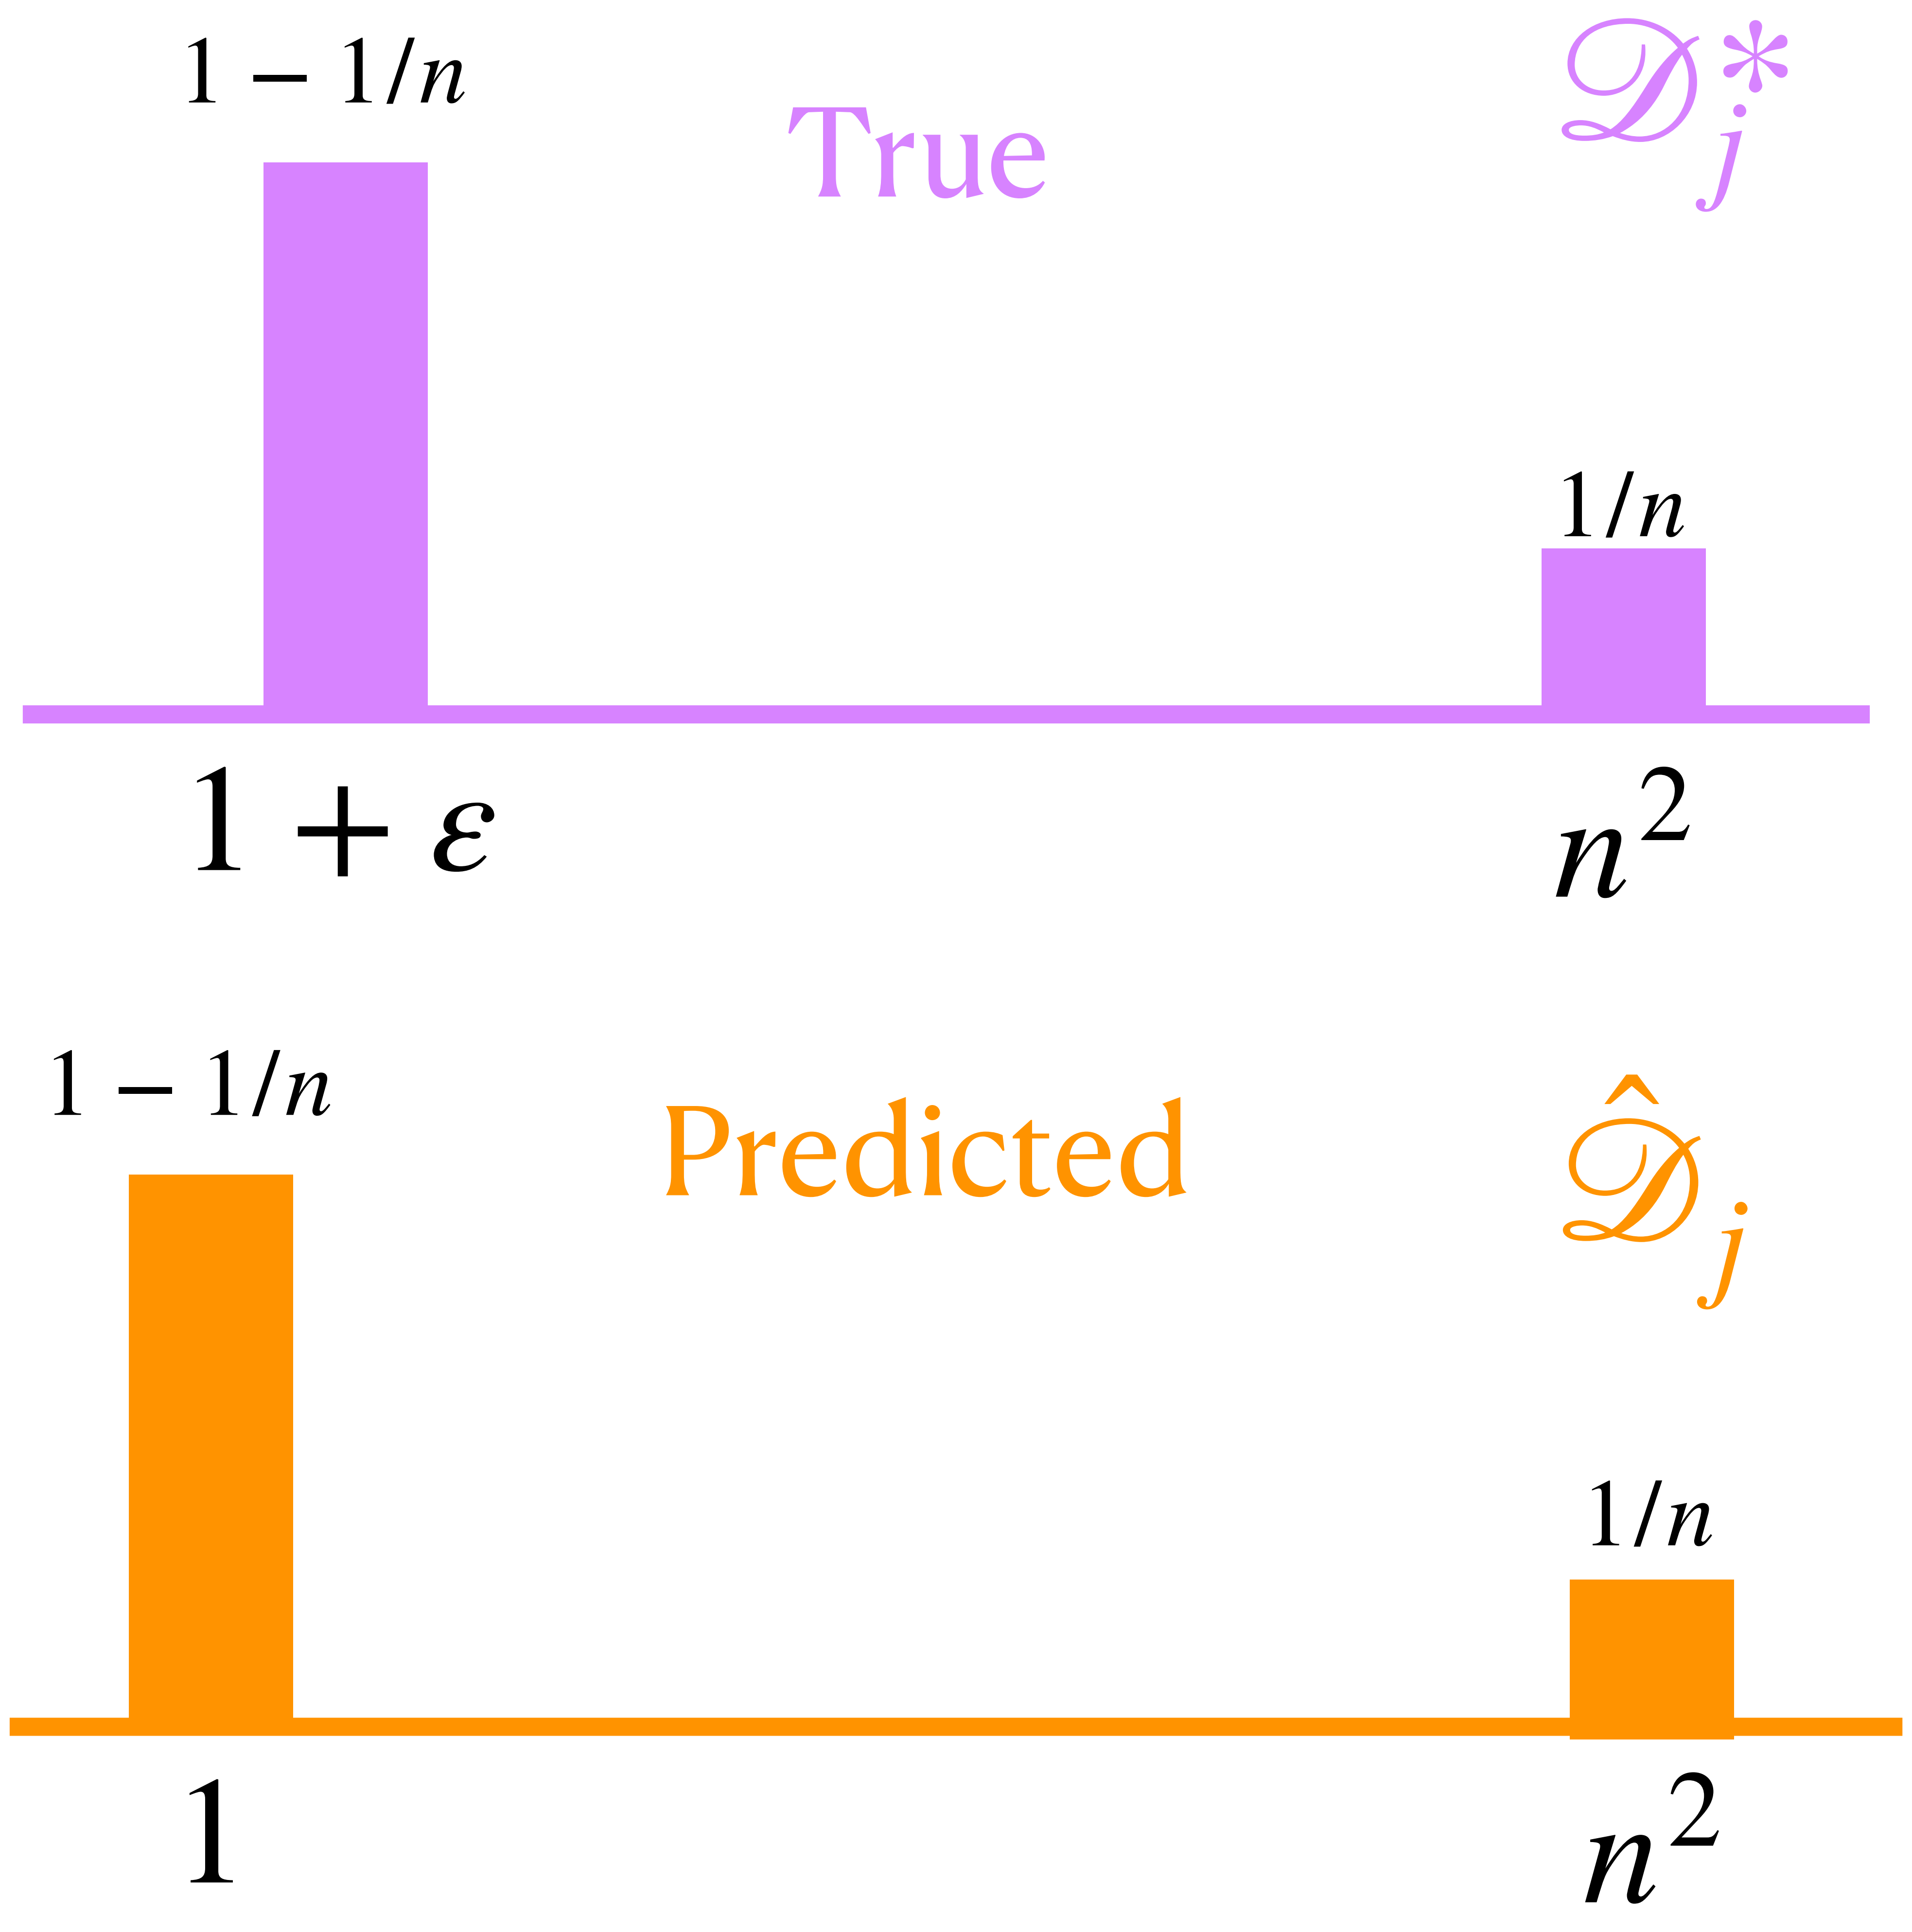

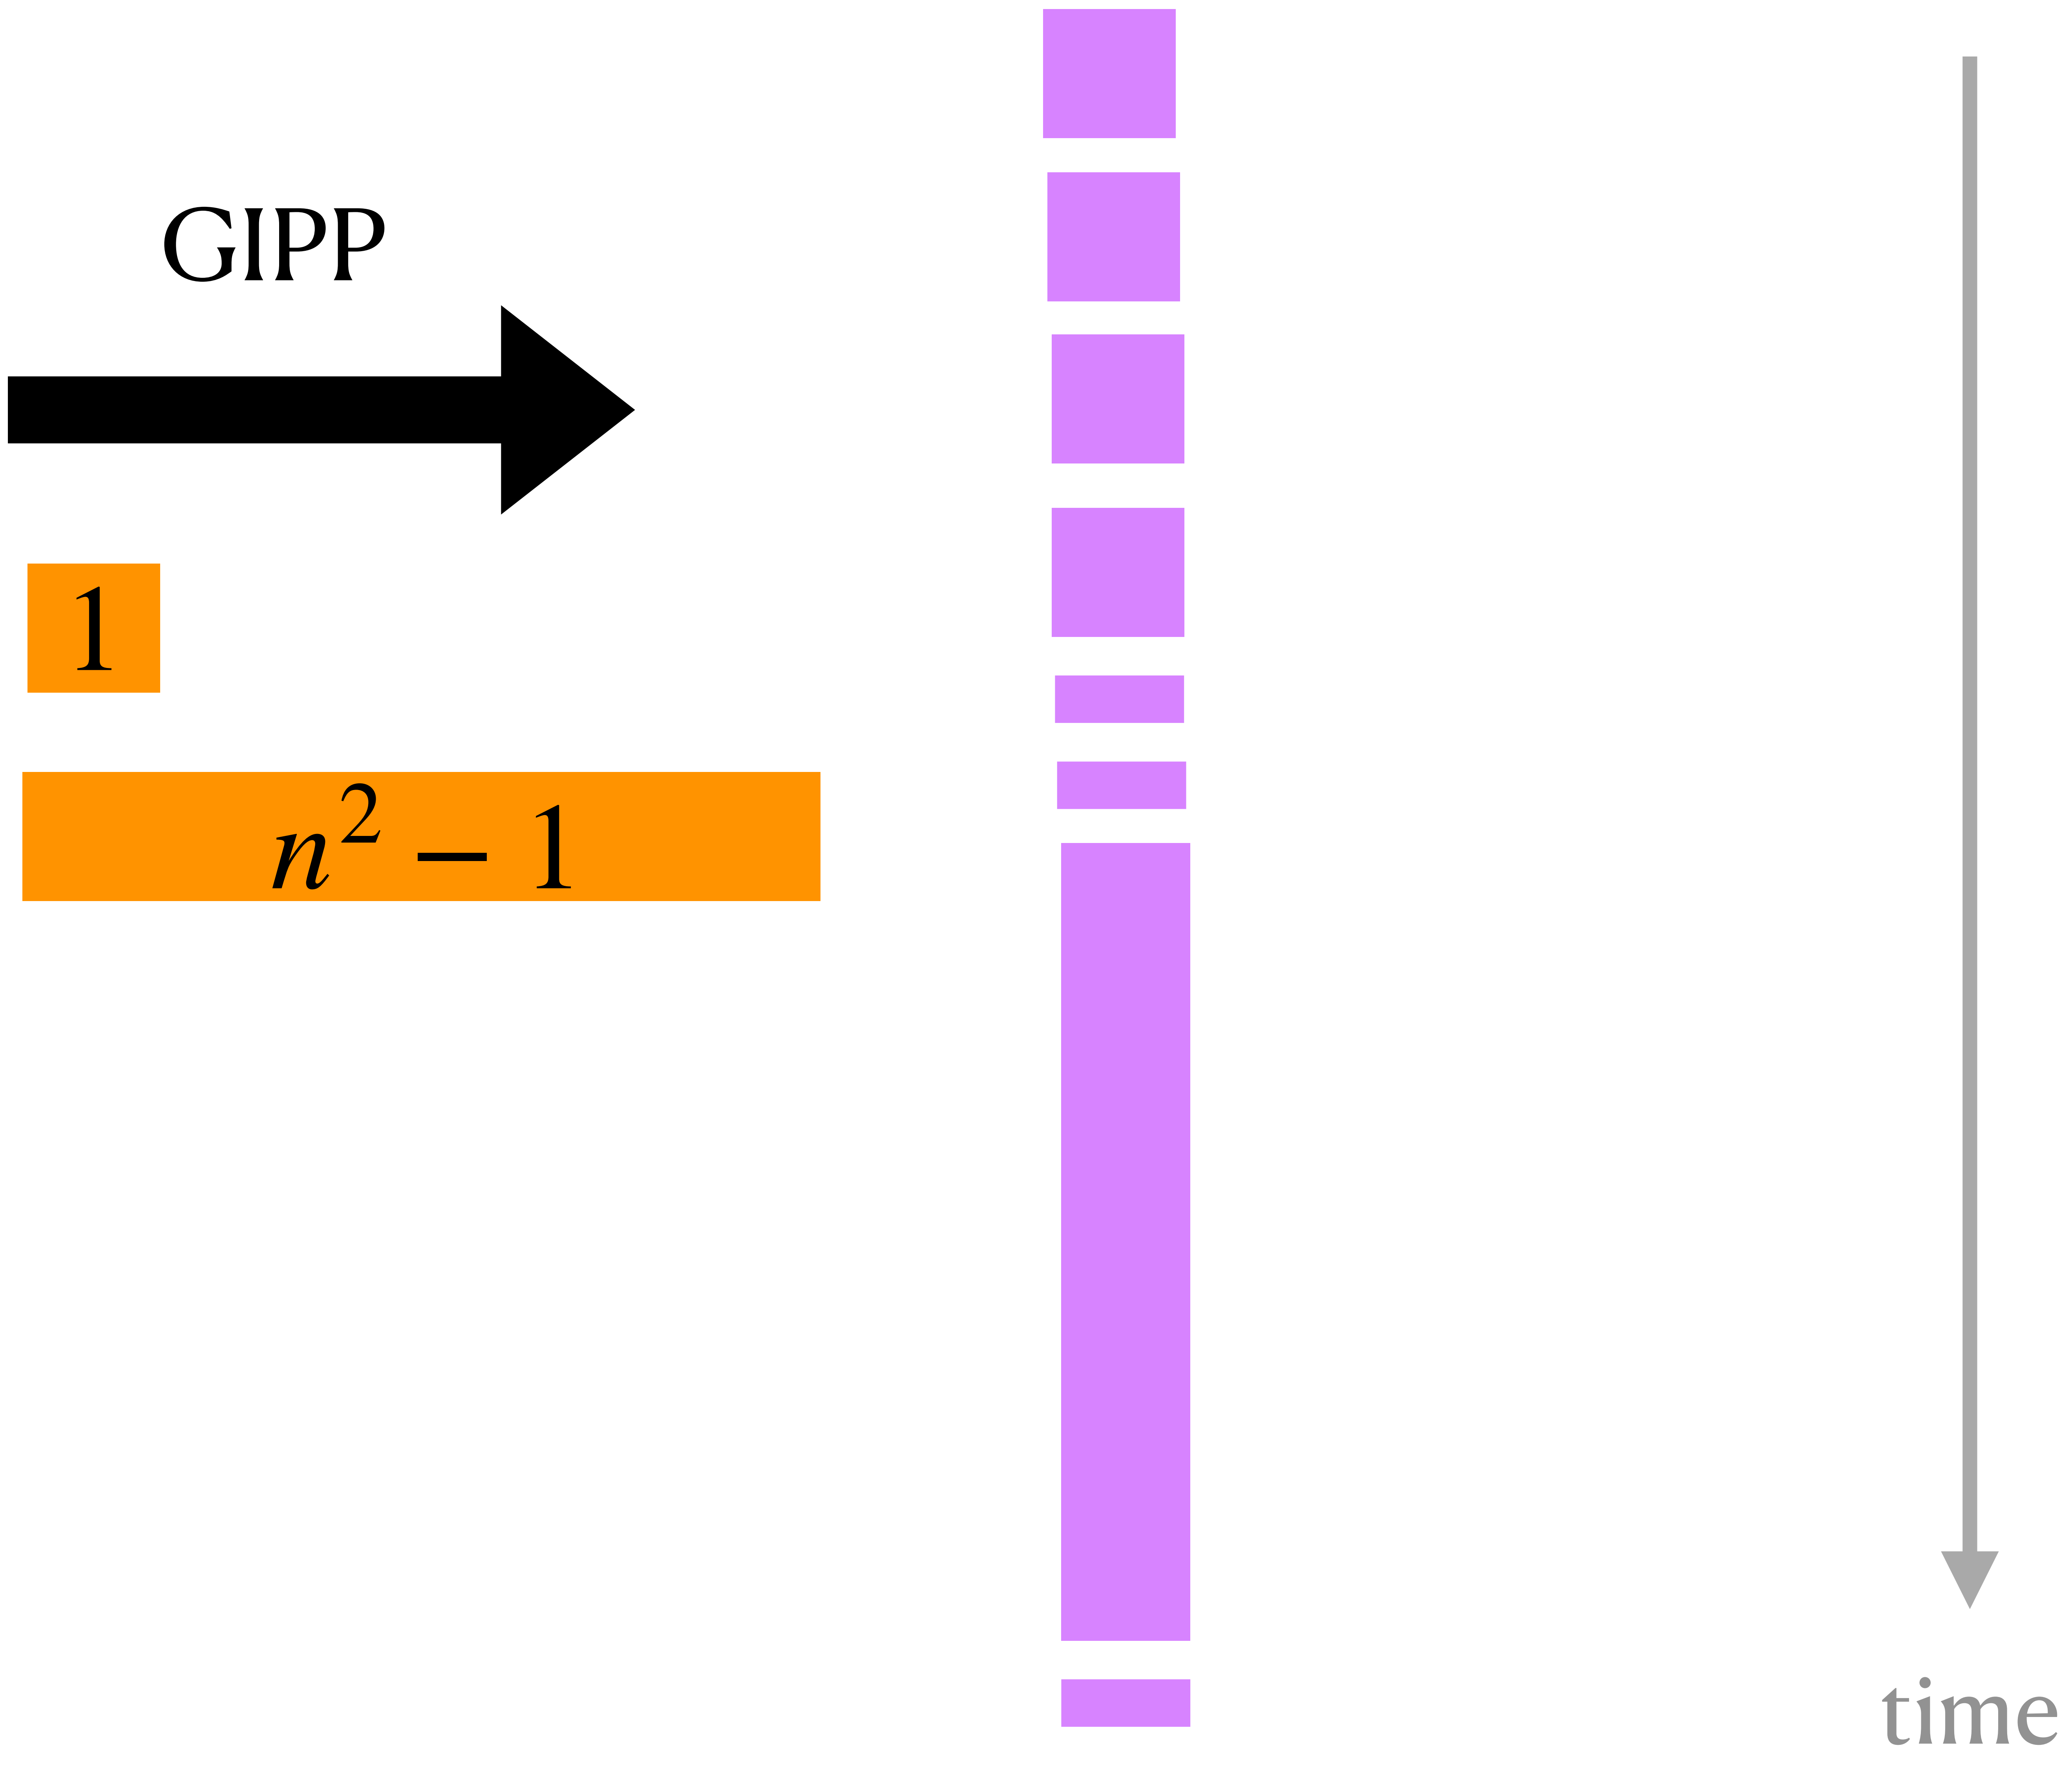

Gittins is Not Robust

$\mathrm{GIPP}(\mathscr{J}^*, \hat{\mathscr{J}}) = \Omega(n) \cdot \mathrm{OPT}(\mathscr{J}^*)$

n i.i.d. distributions

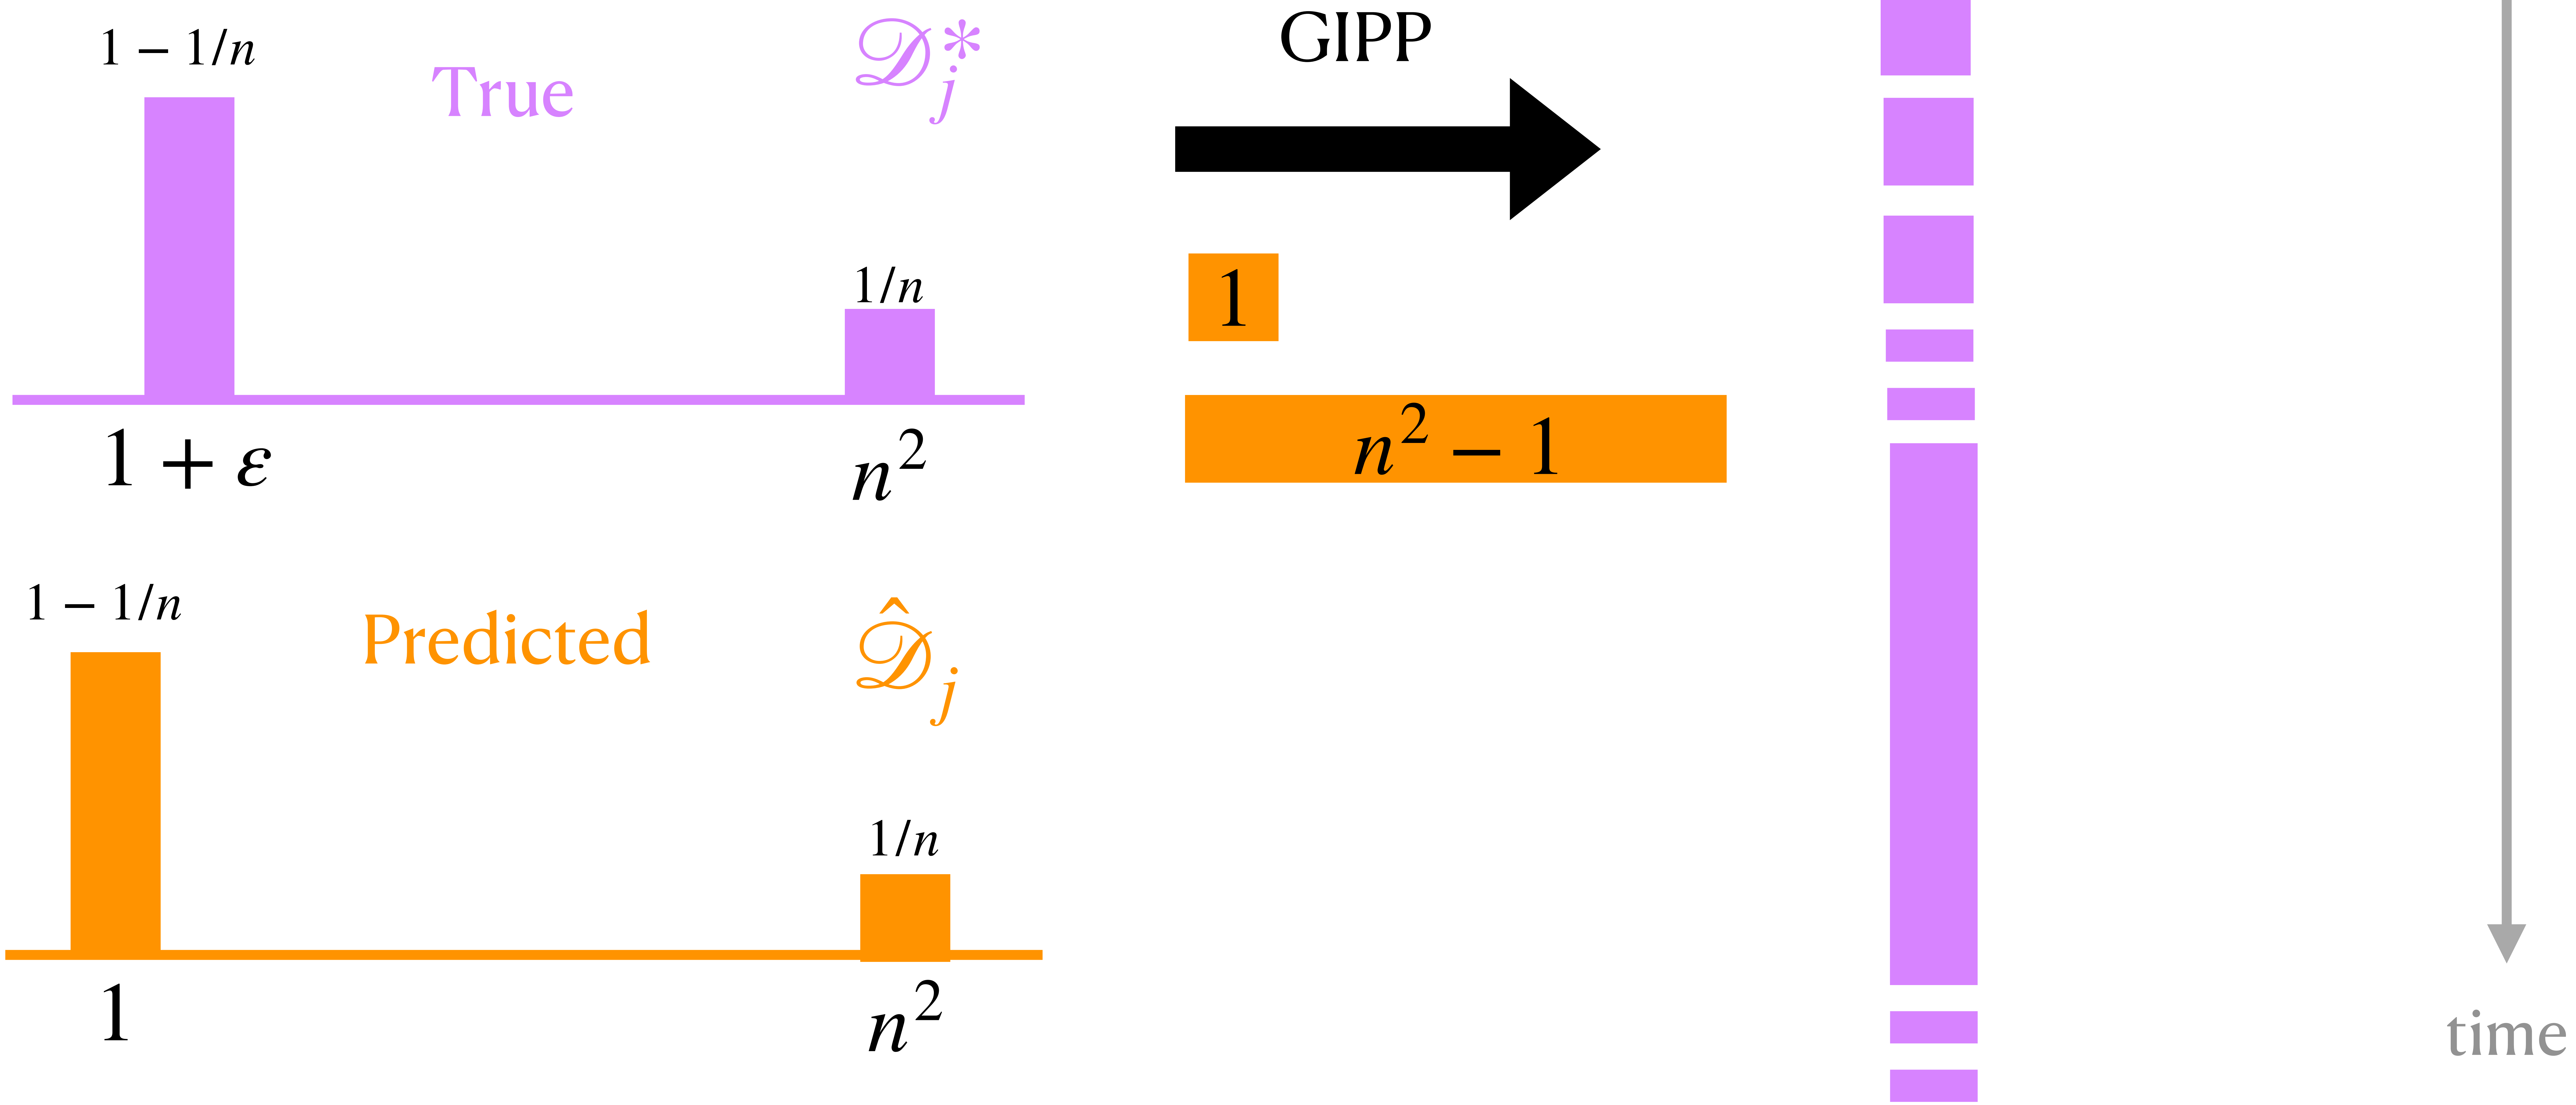

Gittins is Not Robust

$\mathrm{GIPP}(\mathscr{I}^*, \hat{\mathscr{I}}) = \Omega(n) \cdot \mathrm{OPT}(\mathscr{I}^*)$

n i.i.d. distributions

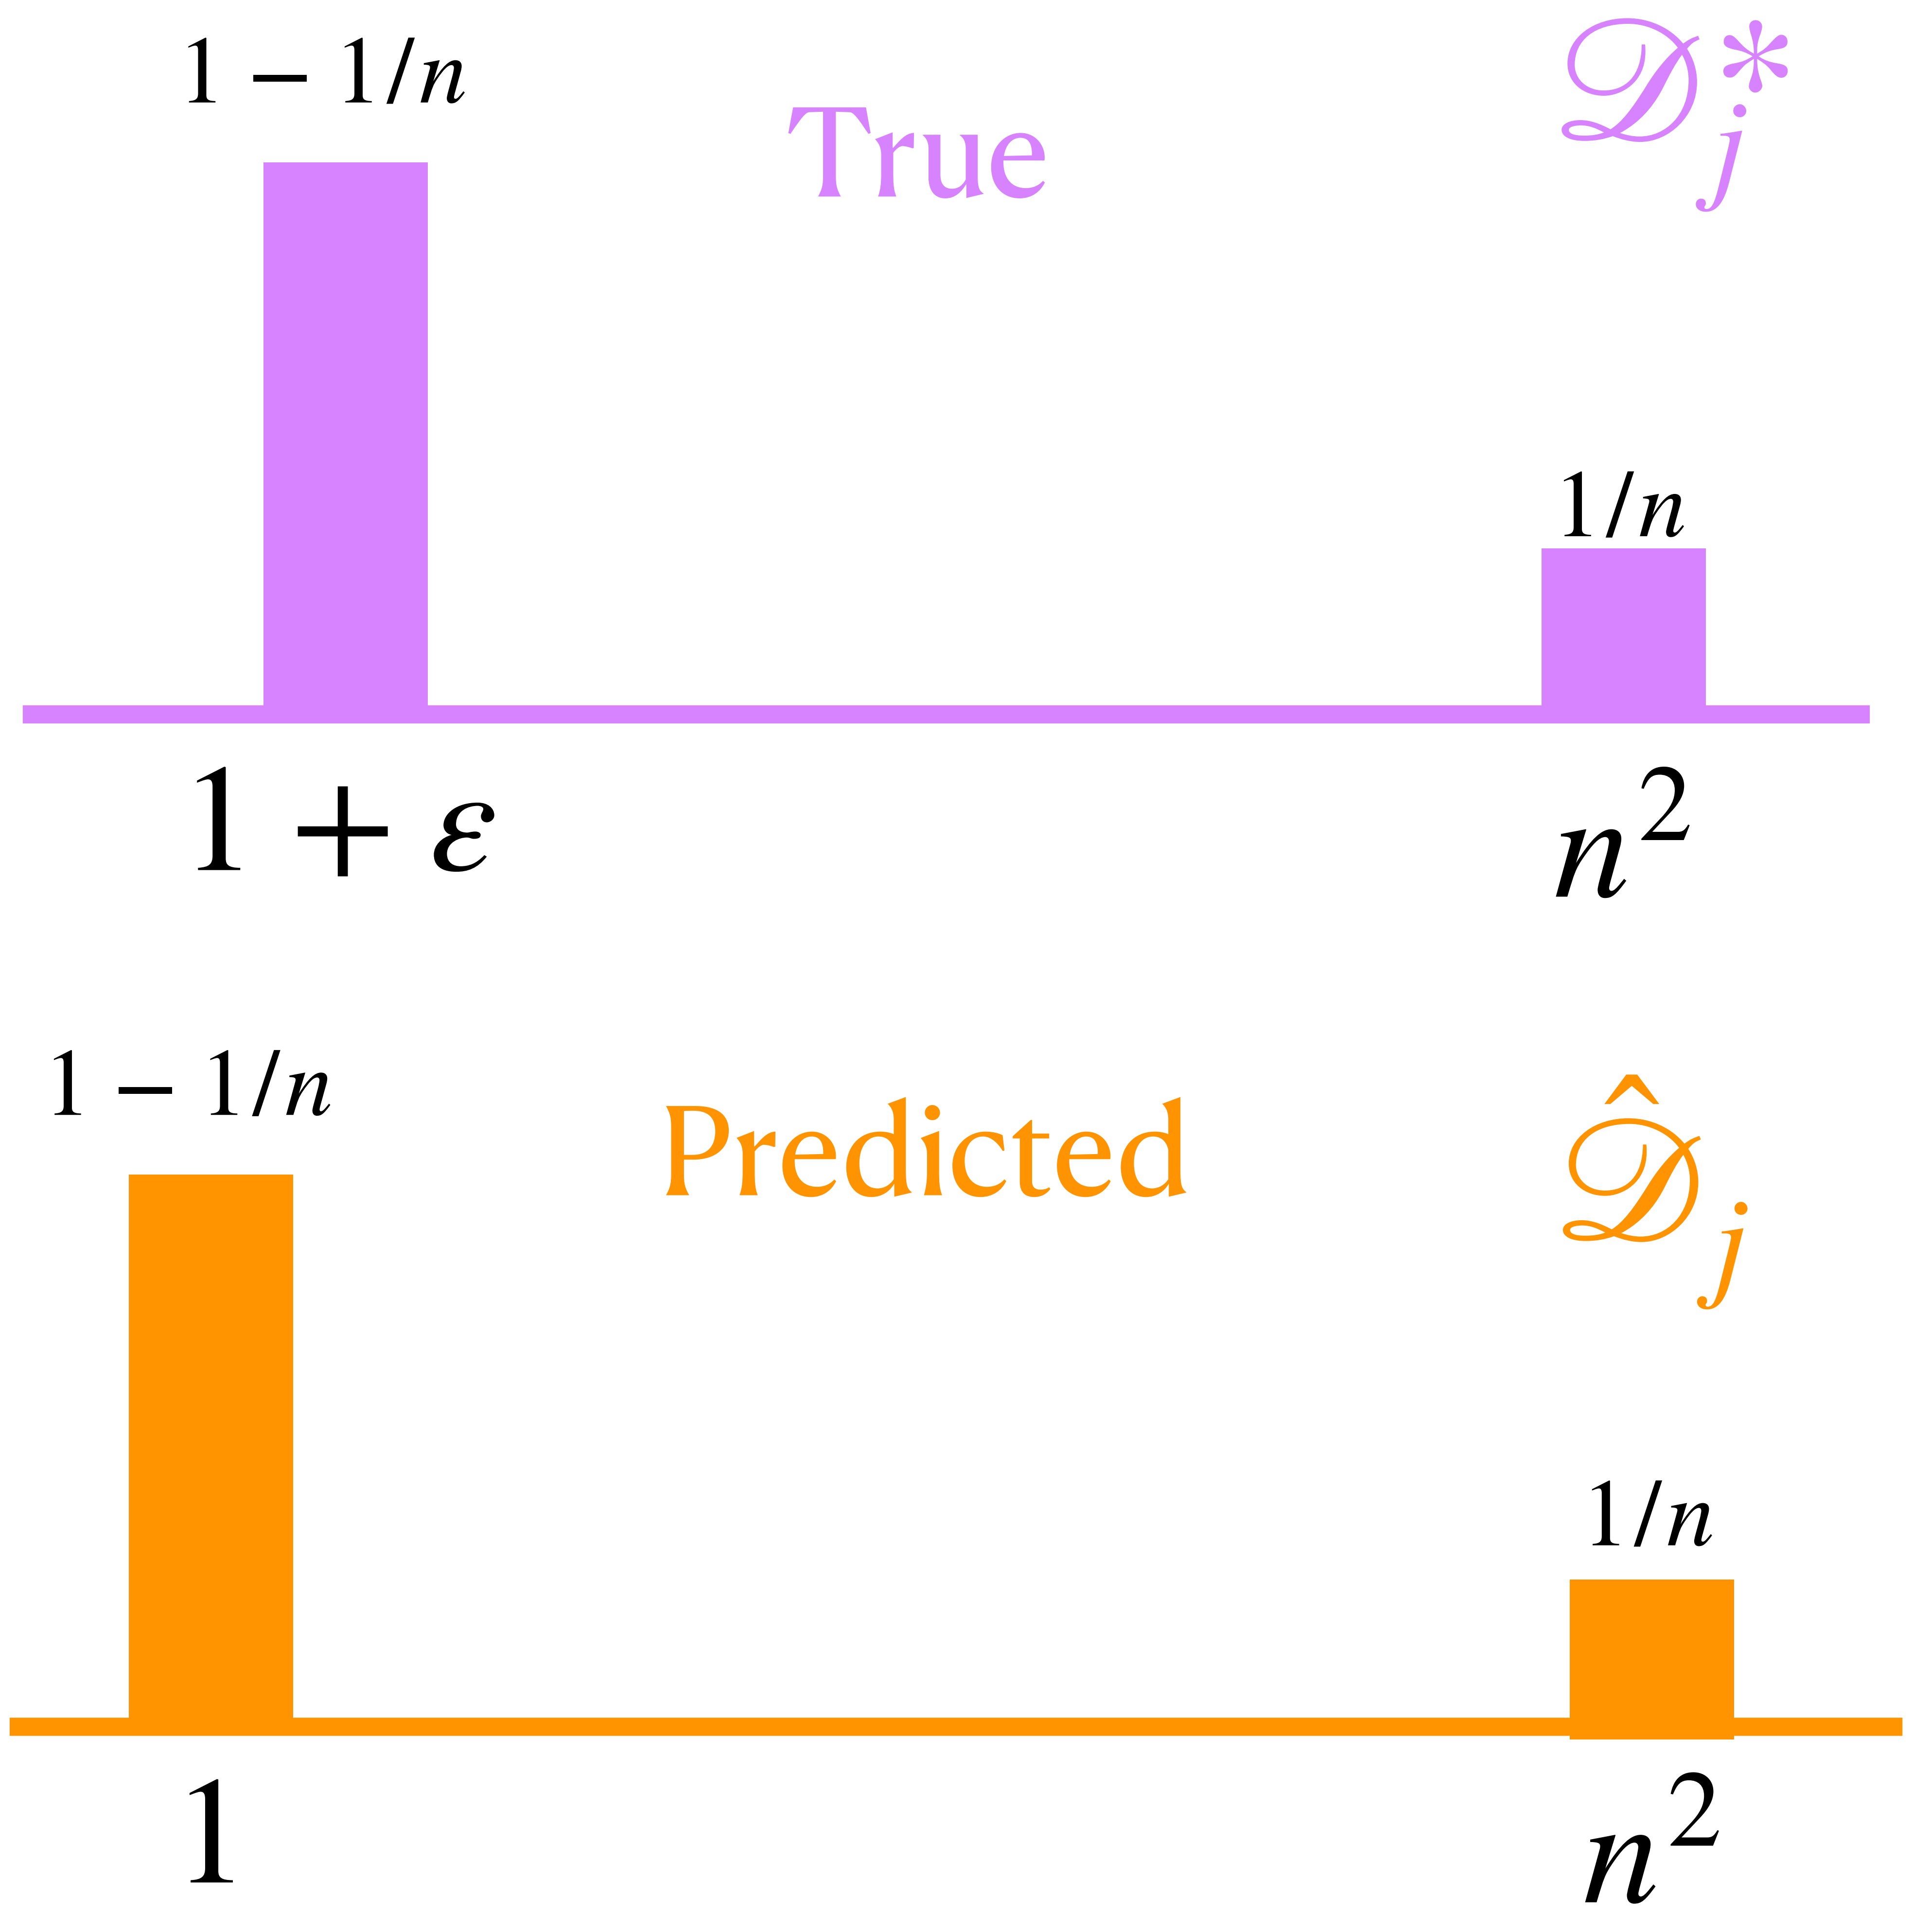

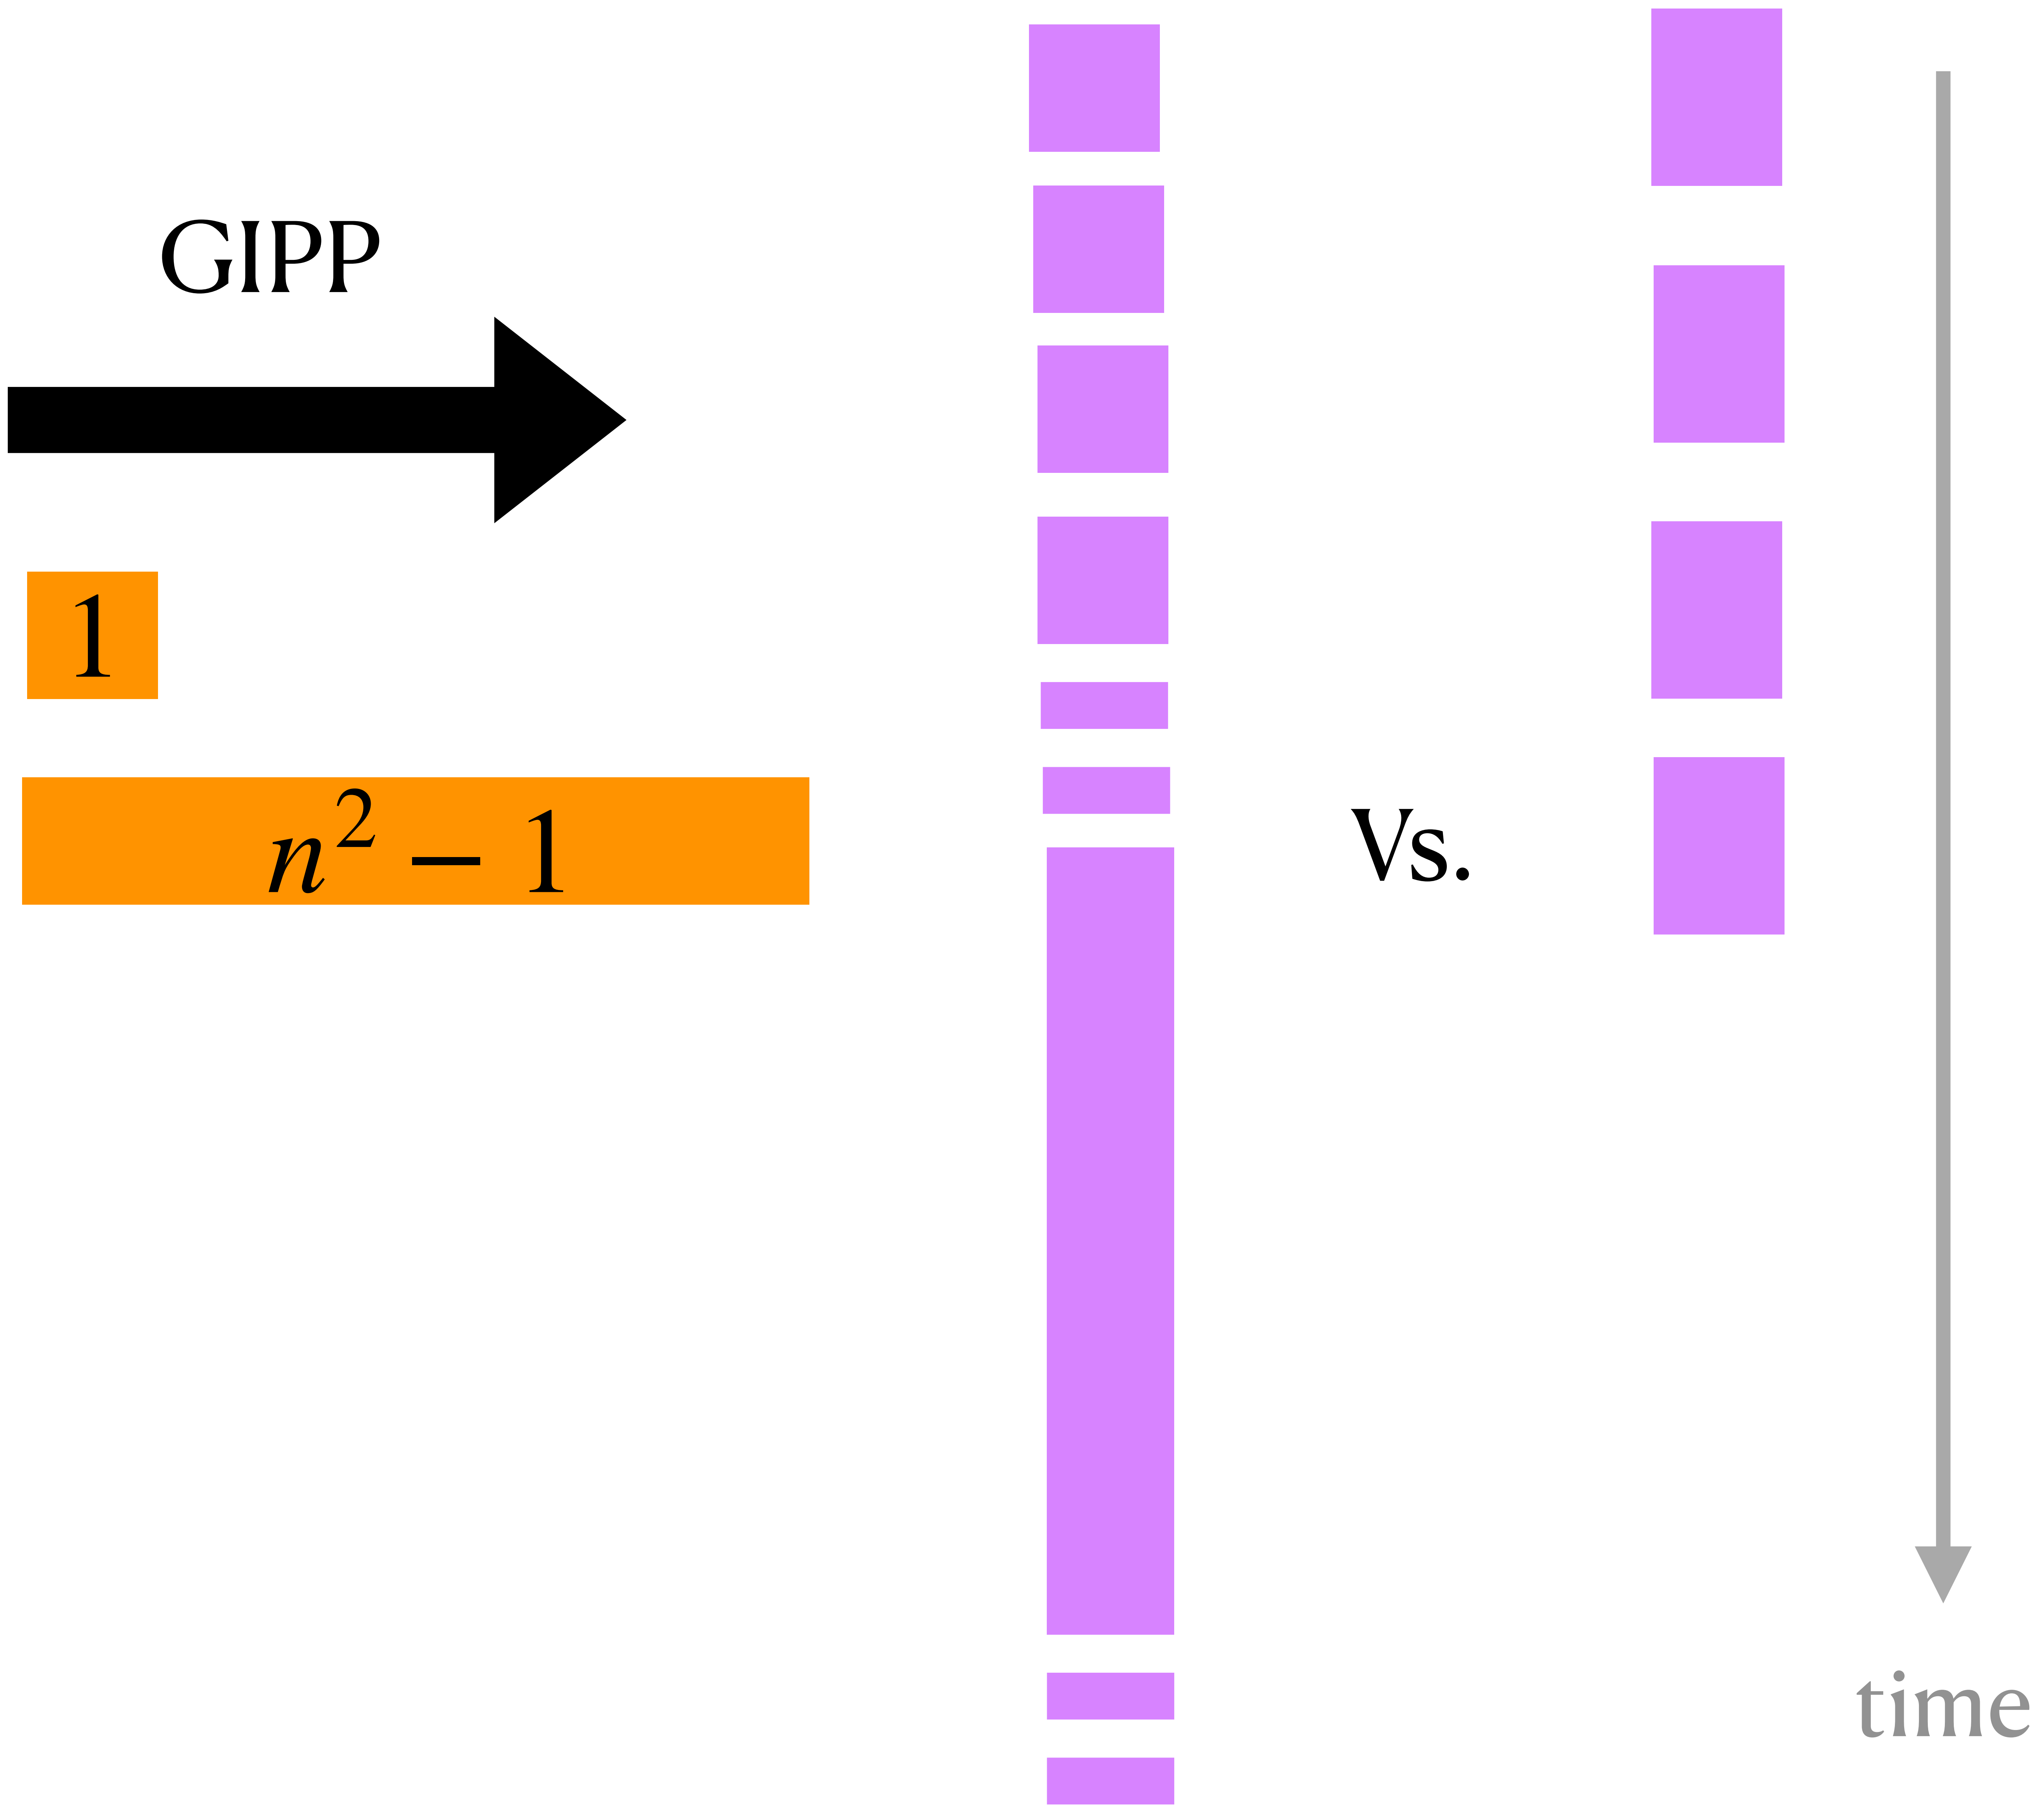

Gittins is Not Robust

$$\mathrm{GIPP}(\mathcal{I}^*, \hat{\mathcal{I}}) = \Omega(n) \cdot \mathrm{OPT}(\mathcal{I}^*)$$

n i.i.d. distributions

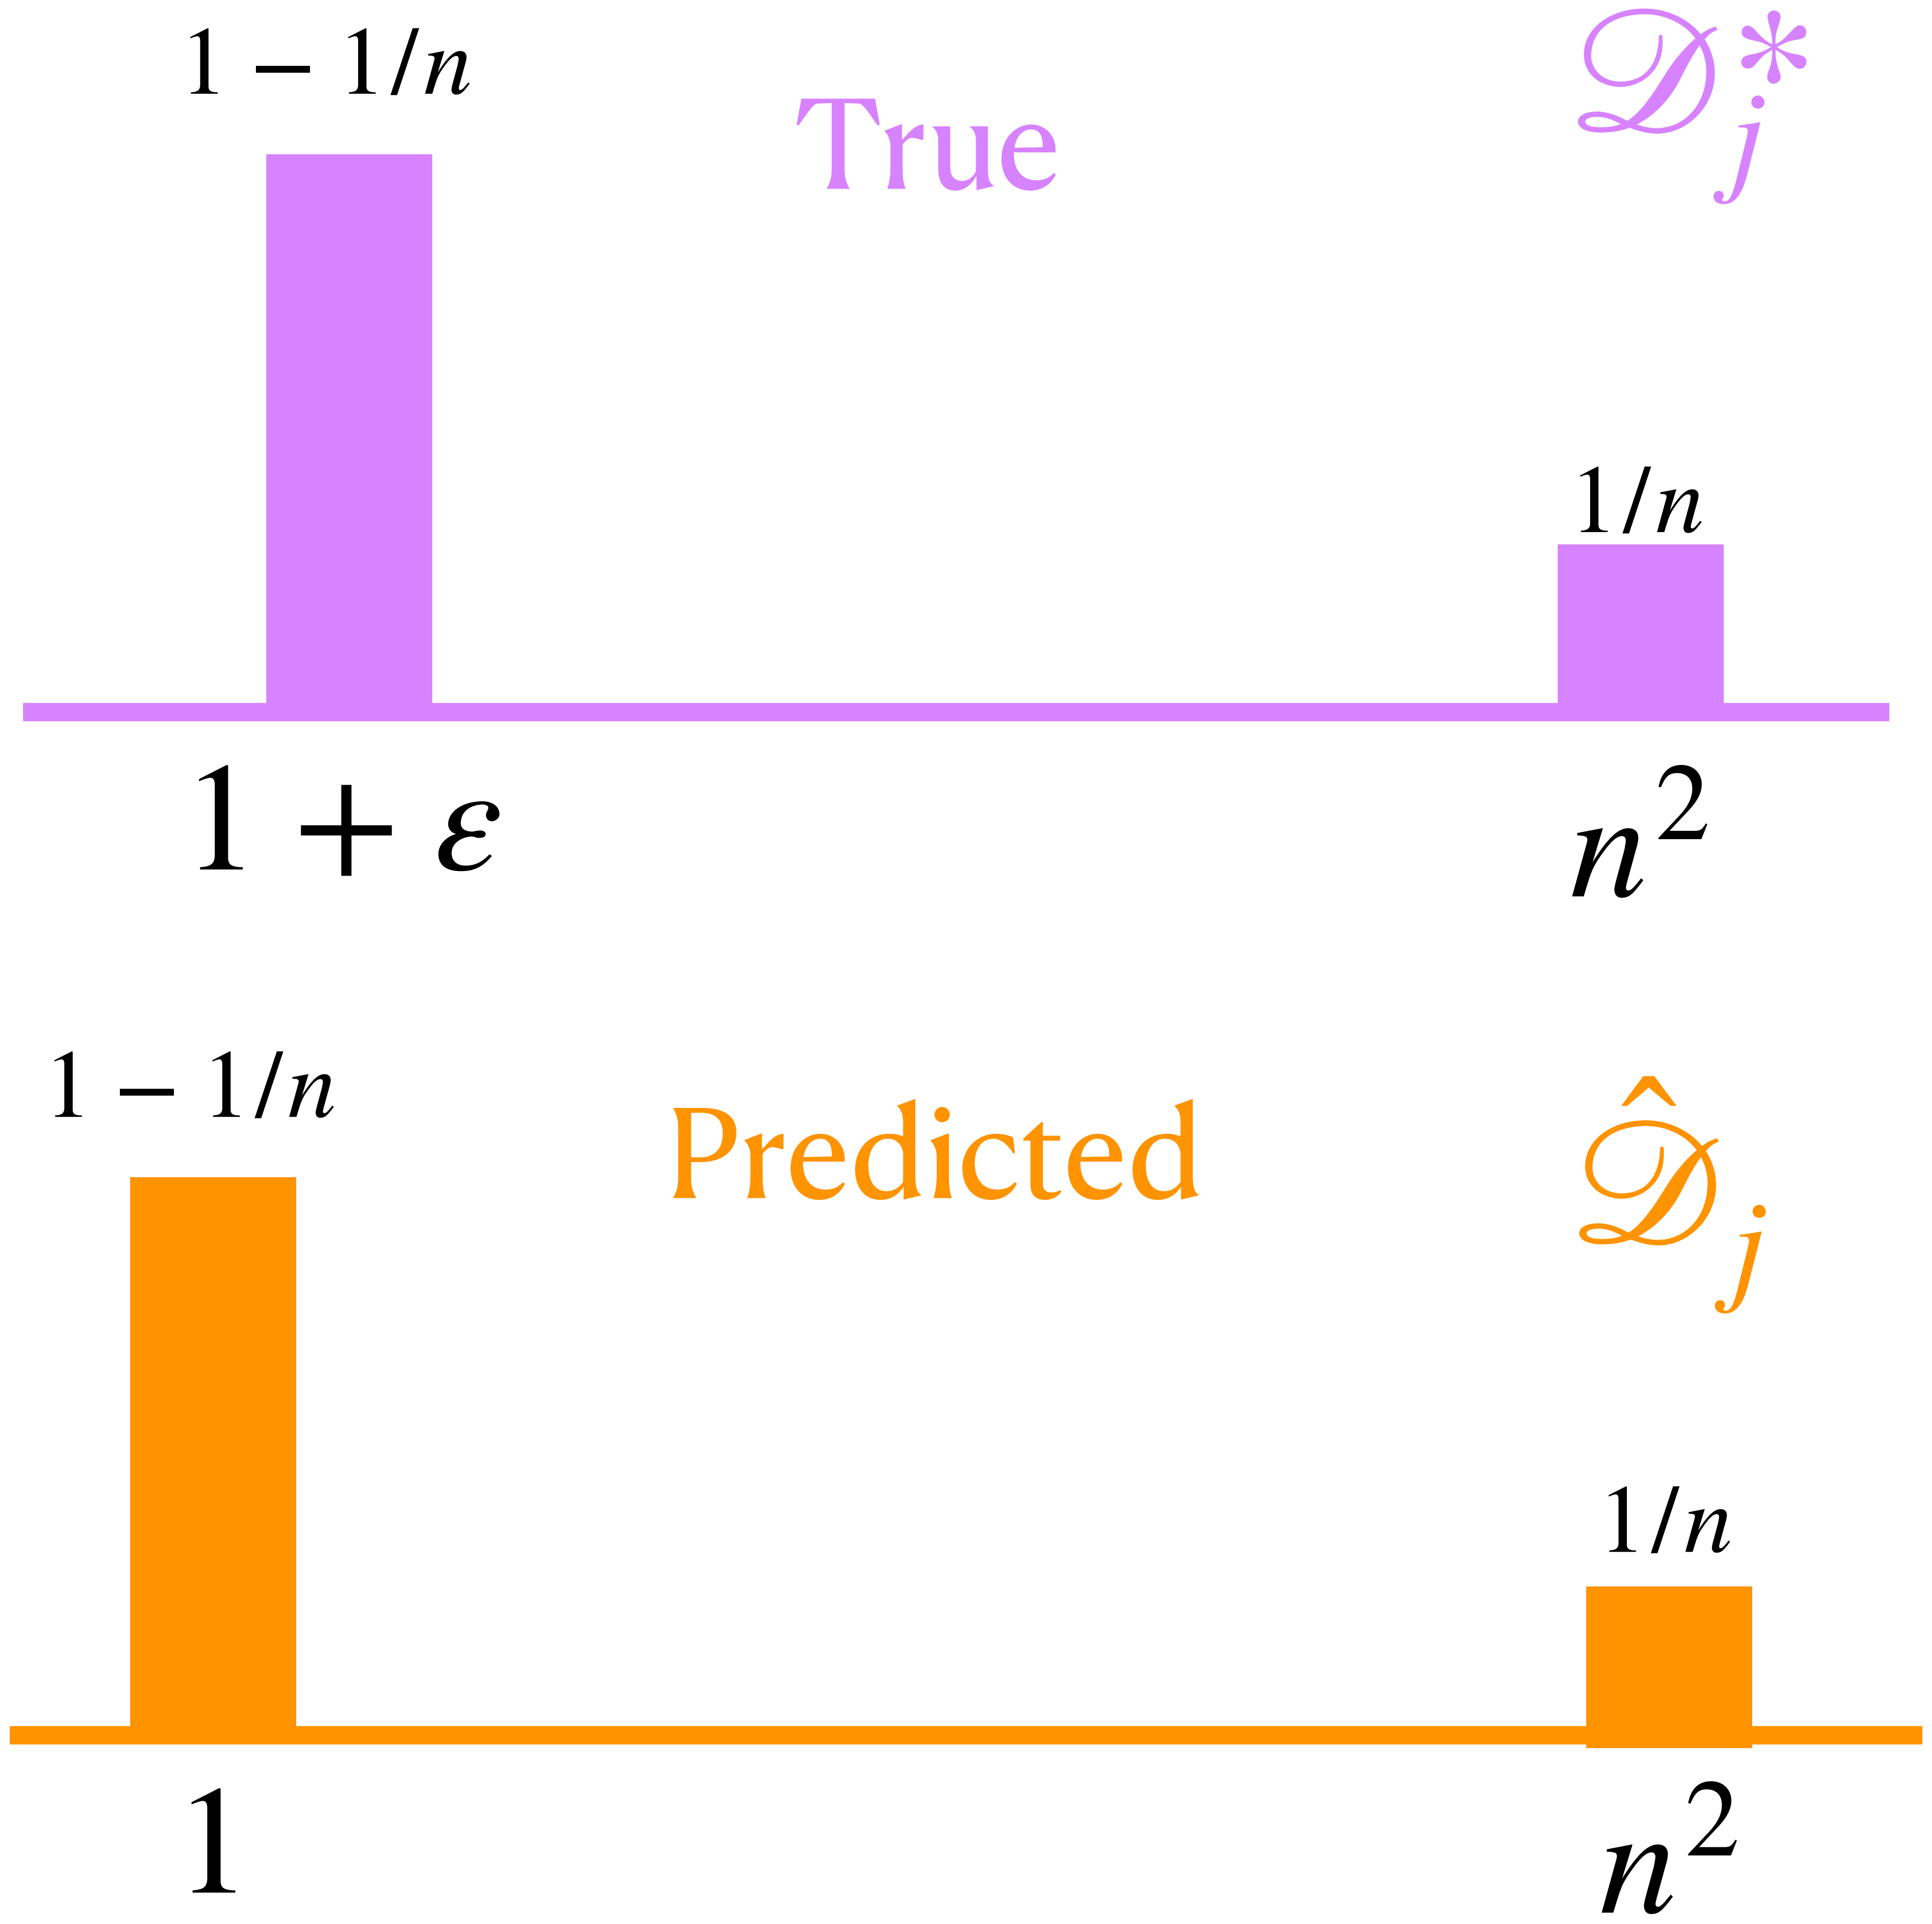

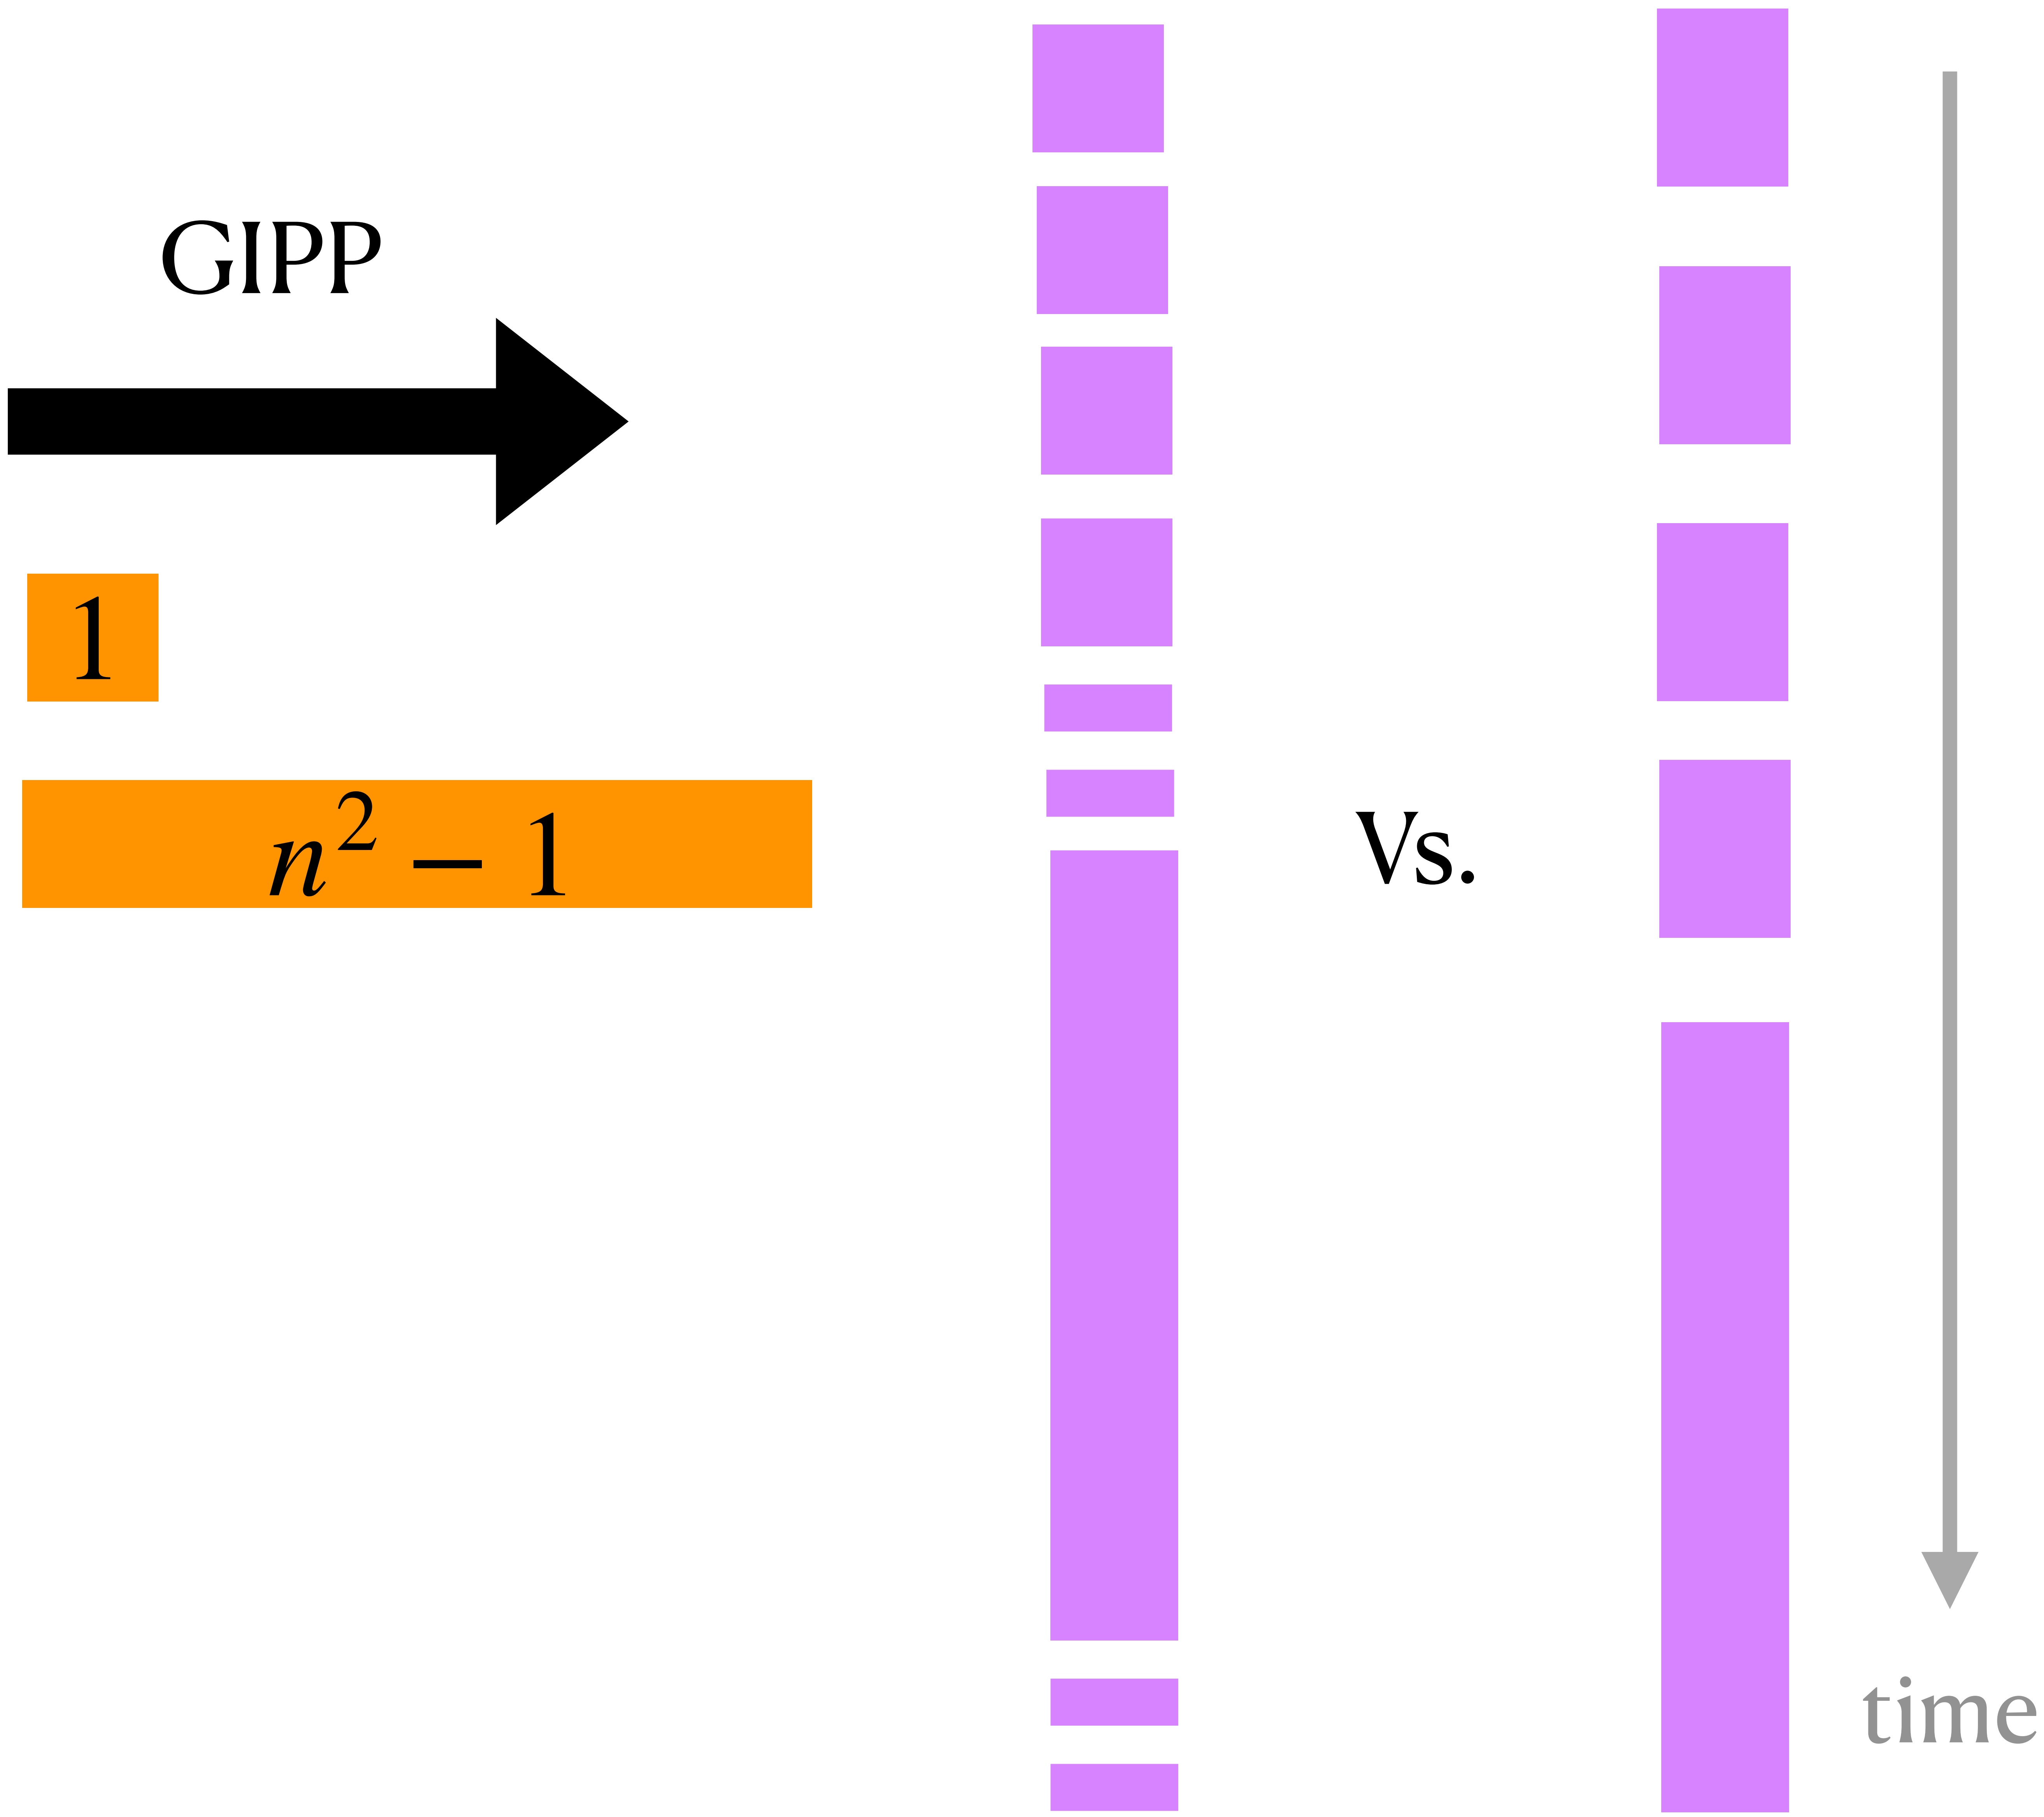

Robust Gittins (RG)

Robust Gittins (RG)

Theorem. Given an “α-close” ($\mathcal{I}* = \{\mathcal{D}_j^*\}_{j=1}^n, \hat{\mathcal{I}} = \{\hat{\mathcal{D}}_j\}_{j=1}^n$), $\alpha \geq 1$, pair of true and predicted distributions, there is a policy, Robust Gittins (RG), satisfying:

$$\mathrm{RG}(\mathcal{I}*, \hat{\mathcal{I}}) \leq \alpha^6 \cdot \mathrm{GIPP}(\mathcal{I}*, \mathcal{I}*) = \mathrm{OPT}(\mathcal{I}*)$$

Robust Gittins (RG)

Theorem. Given an "α-close" $(\mathscr{I}^* = \{\mathscr{D}_j^*\}_{j=1}^n, \hat{\mathscr{I}} = \{\hat{\mathscr{D}}_j\}_{j=1}^n)$, $\alpha \geq 1$, pair of true and predicted distributions, there is a policy, Robust Gittins (RG), satisfying:

$$\mathrm{RG}(\mathscr{I}^*, \hat{\mathscr{I}}) \leq \alpha^6 \cdot \mathrm{GIPP}(\mathscr{I}^*, \mathscr{I}^*) = \mathrm{OPT}(\mathscr{I}^*)$$

Cost of Robust Gittins on true distributions $\mathscr{D}_j^*$, given predicted distributions $\hat{D}_j$ a priori

Robust Gittins (RG)

Theorem. Given an "α-close" ($\mathscr{I}^* = \{\mathscr{D}_j^*\}_{j=1}^n, \hat{\mathscr{I}} = \{\hat{\mathscr{D}}_j\}_{j=1}^n$), $\alpha \geq 1$, pair of true and predicted distributions, there is a policy, Robust Gittins (RG), satisfying:

$$\mathrm{RG}(\mathscr{I}^*, \hat{\mathscr{I}}) \leq \alpha^6 \cdot \mathrm{GIPP}(\mathscr{I}^*, \mathscr{I}^*) = \mathrm{OPT}(\mathscr{I}^*)$$

$\alpha \downarrow 1$: approach optimality

Cost of Robust Gittins on true distributions $\mathscr{D}_j^*$, given predicted distributions $\hat{D}_j$ a priori

Robust Gittins (RG)

Theorem. Given an "α-close" ($\mathcal{I}^* = \{\mathcal{D}_j^*\}_{j=1}^n, \hat{\mathcal{I}} = \{\hat{\mathcal{D}}_j\}_{j=1}^n$), $\alpha \geq 1$, pair of true and predicted distributions, there is a policy, Robust Gittins (RG), satisfying:

Know α

$$\mathrm{RG}(\mathcal{I}^*, \hat{\mathcal{I}}) \leq \alpha^6 \cdot \mathrm{GIPP}(\mathcal{I}^*, \mathcal{I}^*) = \mathrm{OPT}(\mathcal{I}^*)$$

$\alpha \downarrow 1$: approach optimality

Cost of Robust Gittins on true distributions $\mathcal{D}_j^*$, given predicted distributions $\hat{D}_j$ a priori

Robust Gittins (RG)

Theorem. Given an "α-close" ($\mathscr{I}^* = \{\mathscr{D}_j^*\}_{j=1}^n, \hat{\mathscr{I}} = \{\hat{\mathscr{D}}_j\}_{j=1}^n$), $\alpha \geq 1$, pair of true and predicted distributions, there is a policy, Robust Gittins (RG), satisfying:

Know α

$$\mathrm{RG}(\mathscr{I}^*, \hat{\mathscr{I}}) \leq \alpha^6 \cdot \mathrm{GIPP}(\mathscr{I}^*, \mathscr{I}^*) = \mathrm{OPT}(\mathscr{I}^*)$$

$\alpha \downarrow 1$:
approach
optimality

Cost of Robust Gittins on true distributions $\mathscr{D}_j^*$, given predicted distributions $\hat{D}_j$ a priori

Robust Gittins (RG):

- Compute and order $\hat{\mathscr{I}}$ quanta
- Stretch by factor of α

Robust Gittins (RG)

Theorem. Given an “α-close” ($\mathscr{I}^* = \{\mathscr{D}_j^*\}_{j=1}^n, \hat{\mathscr{I}} = \{\hat{\mathscr{D}}_j\}_{j=1}^n$), $\alpha \geq 1$, pair of true and predicted distributions, there is a policy, Robust Gittins (RG), satisfying:

Know α

$$\mathrm{RG}(\mathscr{I}^*, \hat{\mathscr{I}}) \leq \alpha^6 \cdot \mathrm{GIPP}(\mathscr{I}^*, \mathscr{I}^*) = \mathrm{OPT}(\mathscr{I}^*)$$

$\alpha \downarrow 1$: approach optimality

Cost of Robust Gittins on true distributions $\mathscr{D}_j^*$, given predicted distributions $\hat{D}_j$ a priori

Robust Gittins (RG):

- Compute and order $\hat{\mathscr{I}}$ quanta
- Stretch by factor of α

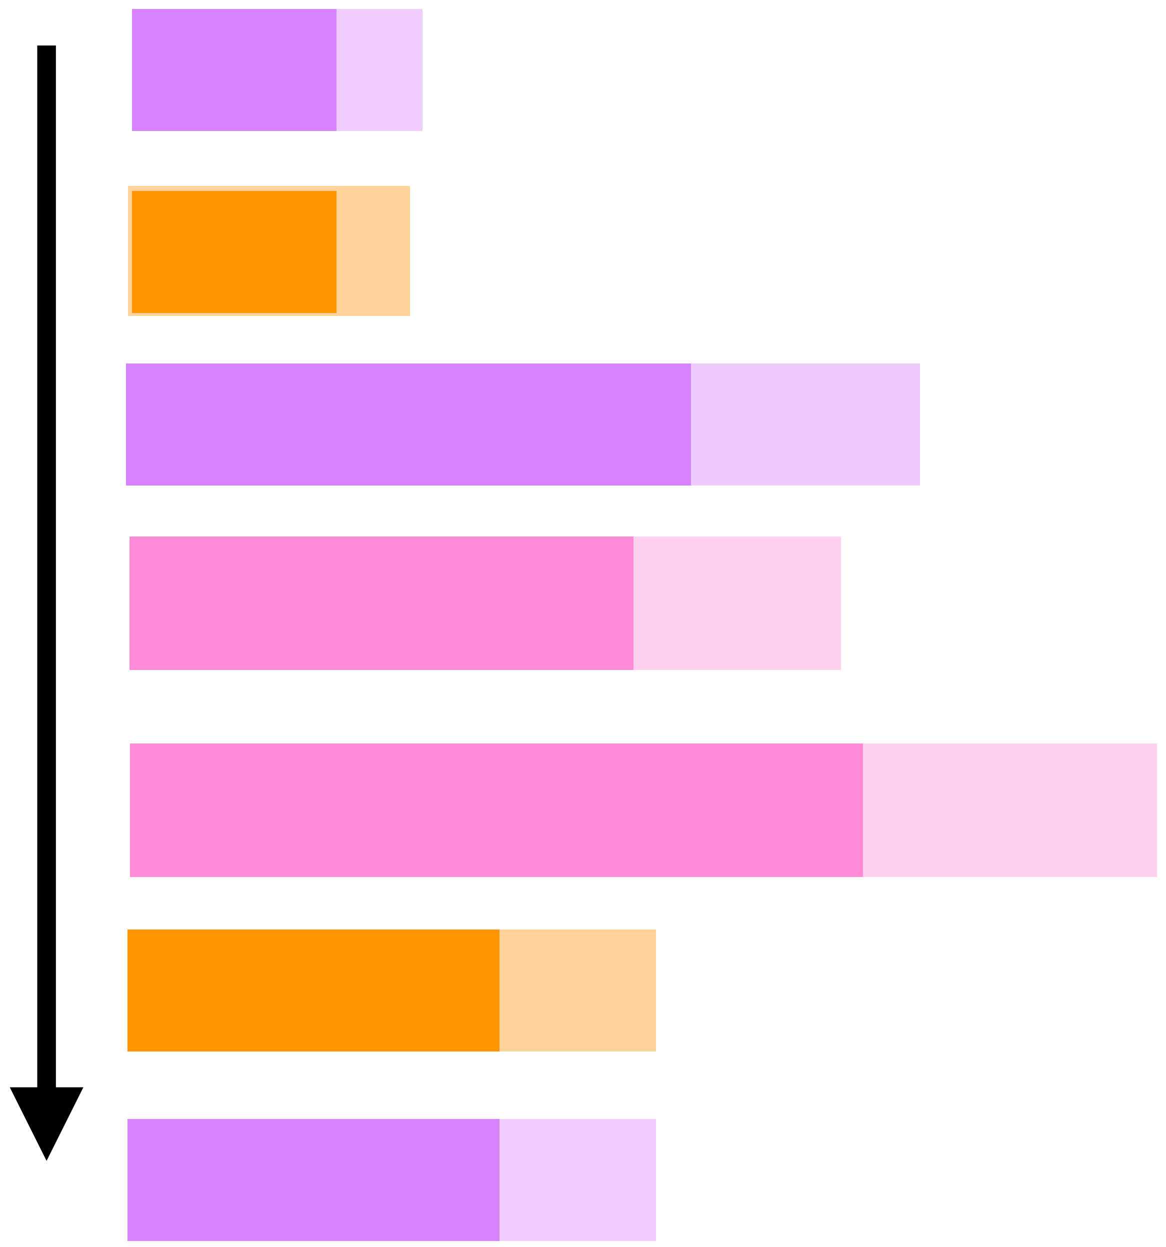

decreasing order of rank

Robust Gittins (RG)

Theorem. Given an "α-close" ($\mathscr{I}^* = \{\mathscr{D}_j^*\}_{j=1}^n, \hat{\mathscr{I}} = \{\hat{\mathscr{D}}_j\}_{j=1}^n$), $\alpha \geq 1$, pair of true and predicted distributions, there is a policy, Robust Gittins (RG), satisfying:

Know α

$$\mathrm{RG}(\mathscr{I}^*, \hat{\mathscr{I}}) \leq \alpha^6 \cdot \mathrm{GIPP}(\mathscr{I}^*, \mathscr{I}^*) = \mathrm{OPT}(\mathscr{I}^*)$$

$\alpha \downarrow 1$: approach optimality

Cost of Robust Gittins on true distributions $\mathscr{D}_j^*$, given predicted distributions $\hat{D}_j$ a priori

Robust Gittins (RG):

- Compute and order $\hat{\mathscr{I}}$ quanta
- Stretch by factor of α

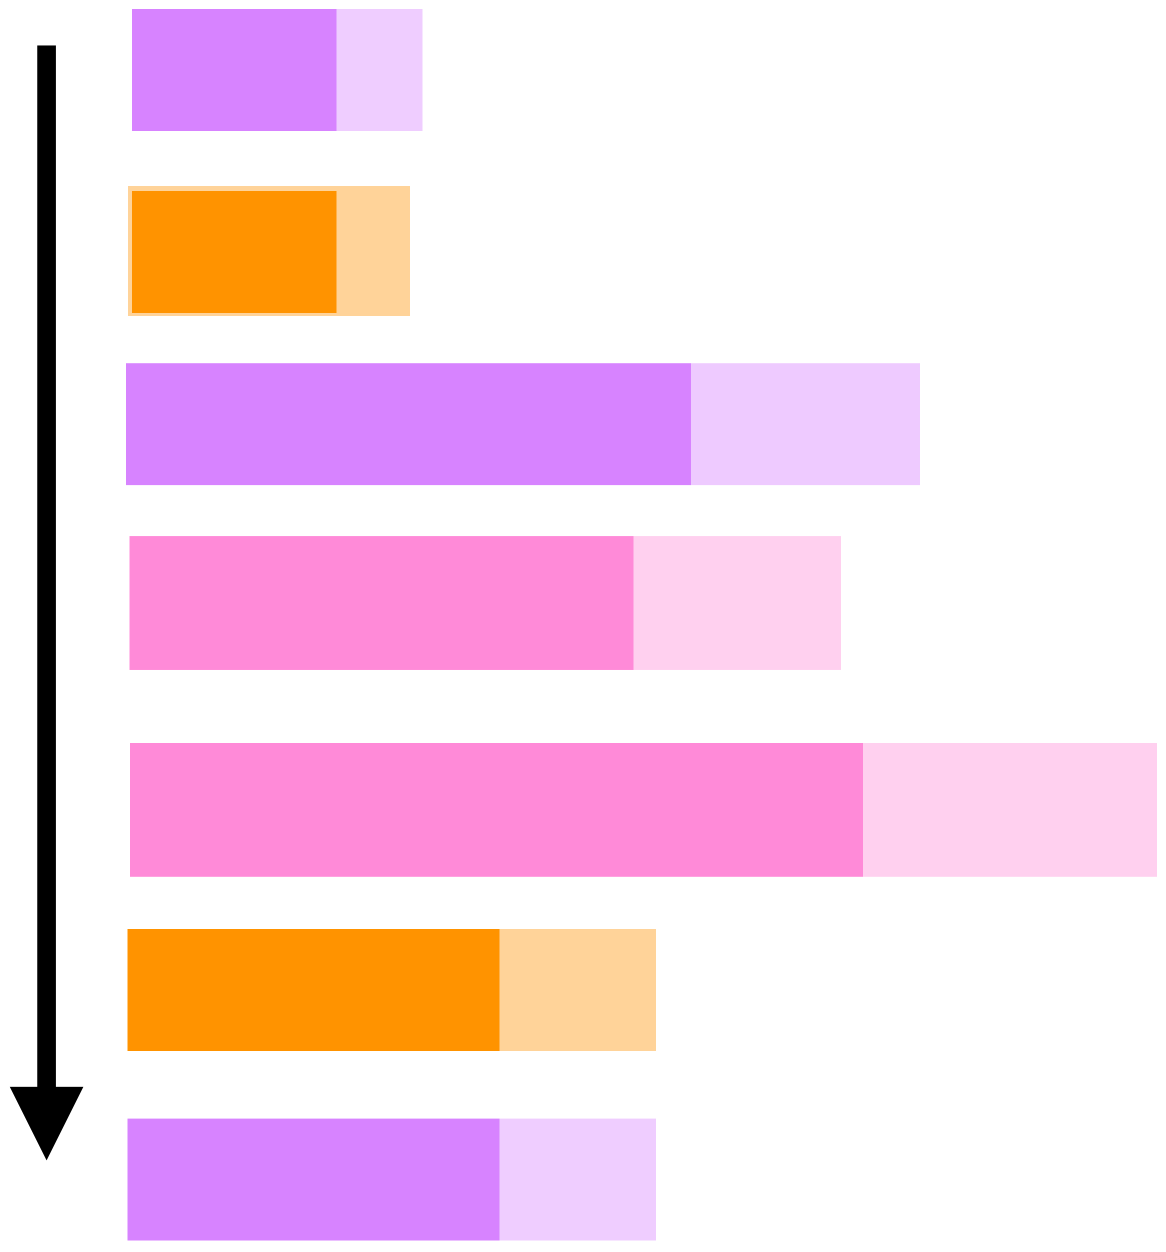

decreasing order of rank

Robust Gittins (RG)

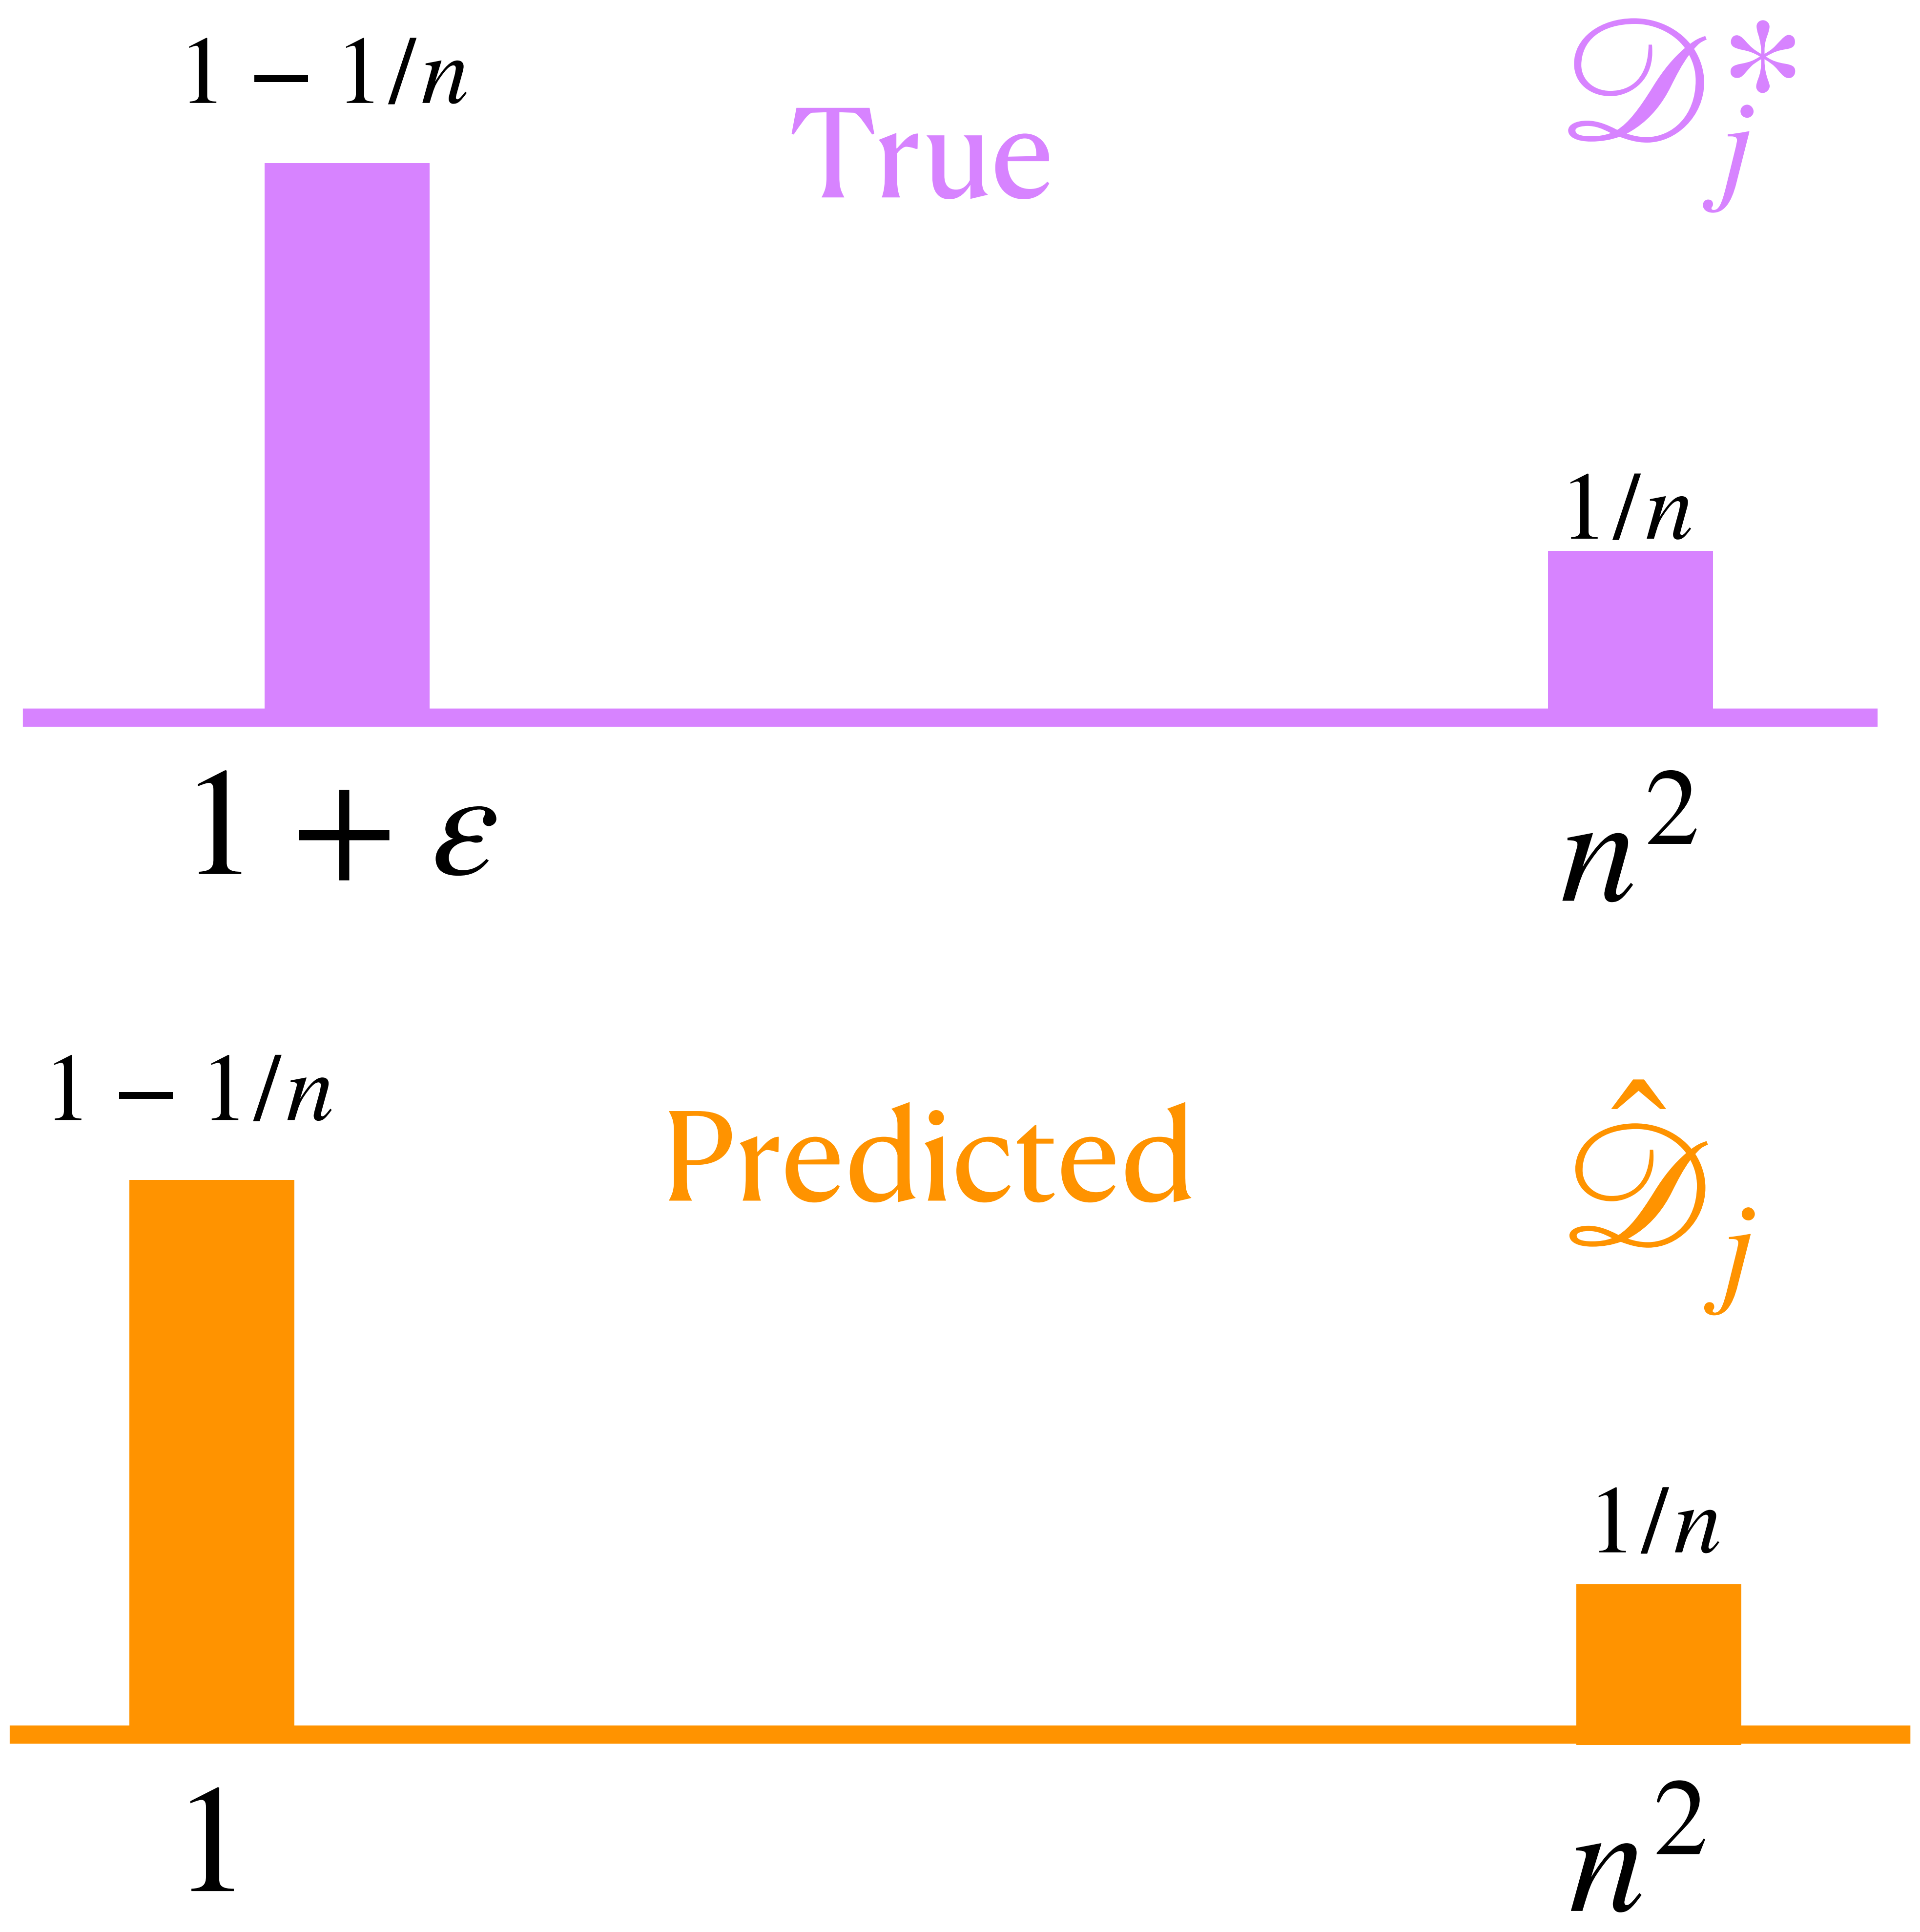

Robust Gittins (RG)

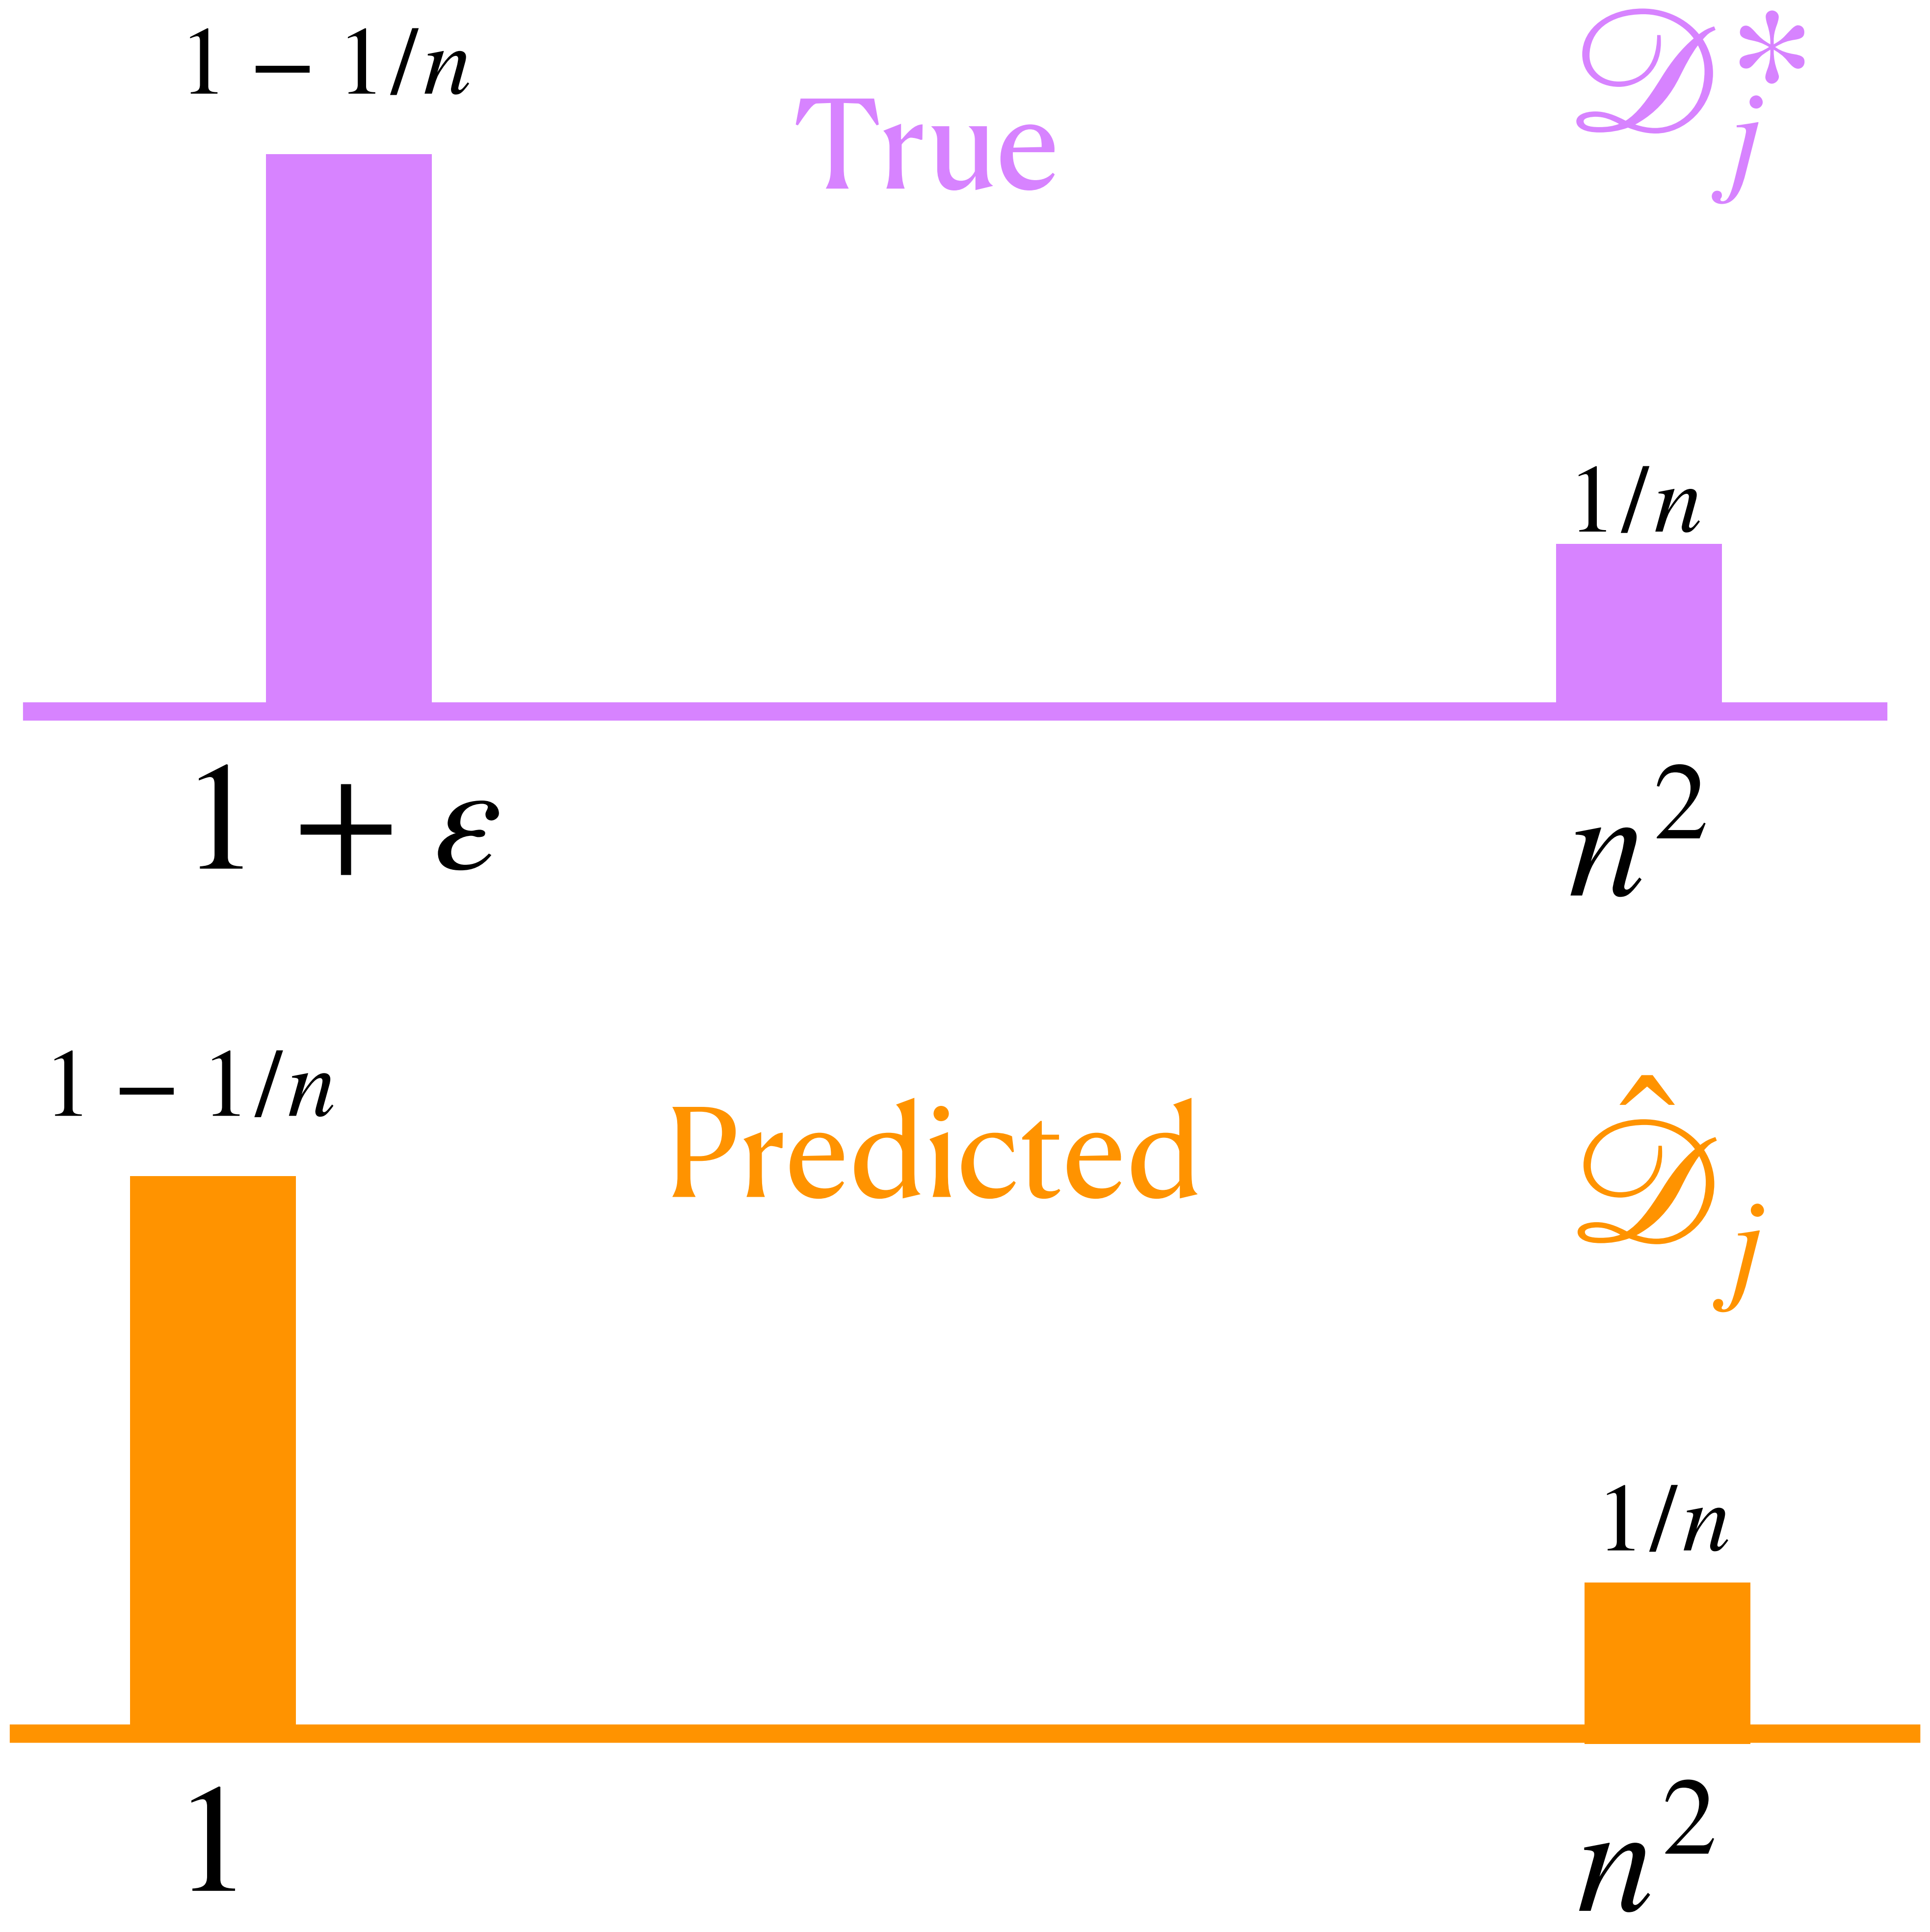

$$\alpha = 1 + \varepsilon$$

Robust Gittins (RG)

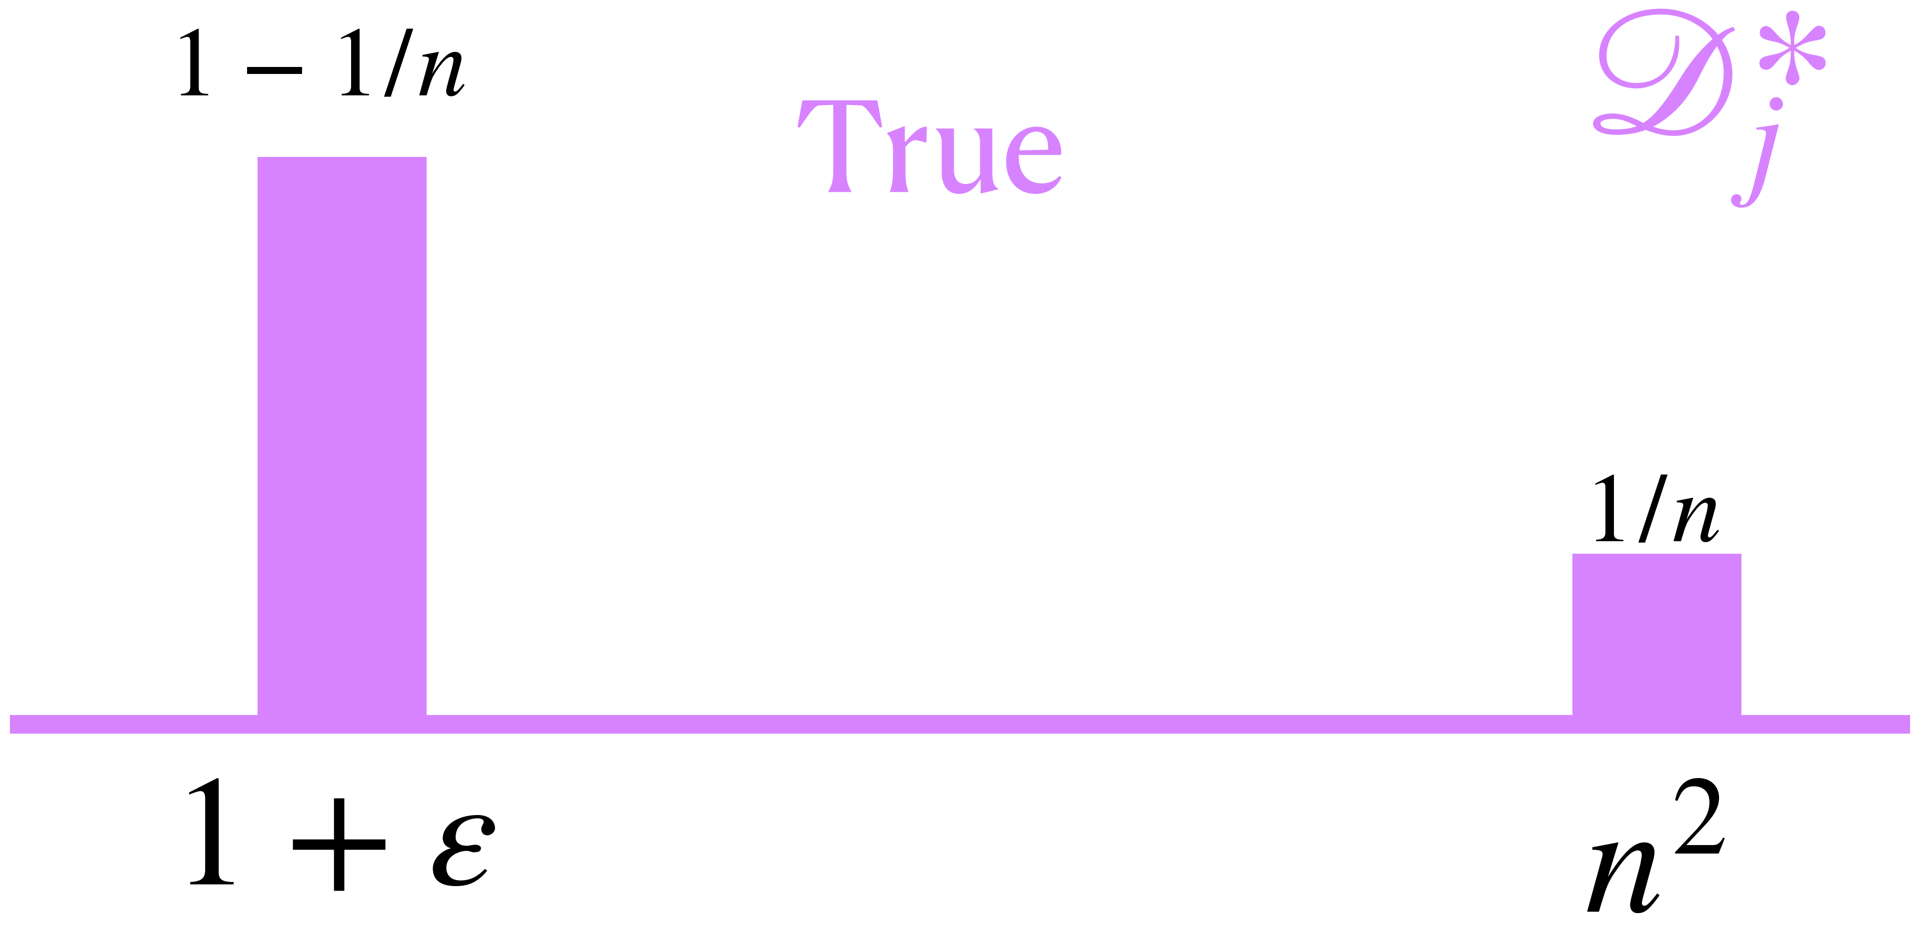

$\alpha = 1 + \varepsilon$

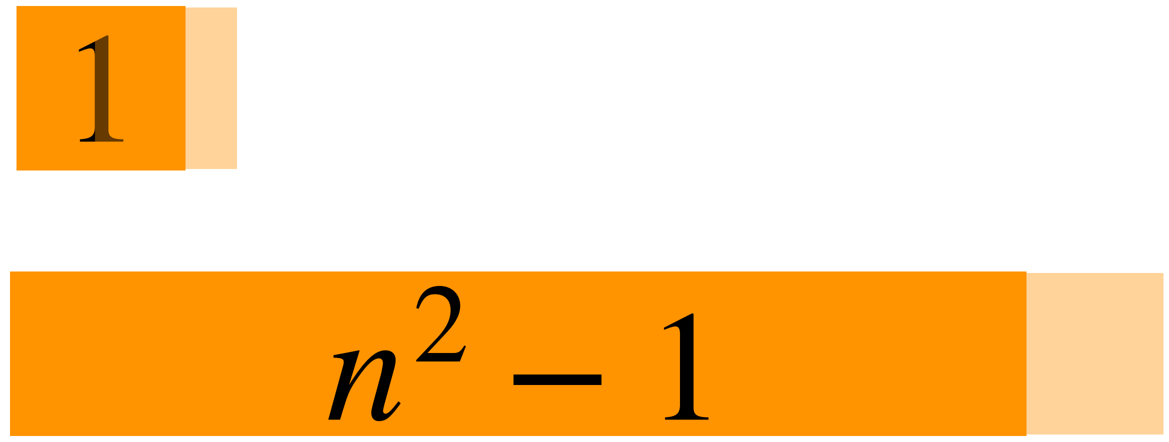

A New Error Measure

A New Error Measure

Definition. $(\mathscr{D}, \mathscr{D}')$ are ***α*-close** $(\alpha \geq 1)$ if

"multiplicative error in upper tails is small"

A New Error Measure

Definition. $(\mathscr{D}, \mathscr{D}')$ are **α-close** $(\alpha \geq 1)$ if

$$\forall x \geq 0 : \qquad \frac{1}{\alpha} \cdot \mathbb{P}_{P \sim \mathscr{D}}(P > \alpha \cdot x) \leq \mathbb{P}_{P' \sim \mathscr{D}'}(P' > x) \leq \alpha \cdot \mathbb{P}_{P \sim \mathscr{D}}(P > x/\alpha)$$

"multiplicative error in upper tails is small"

A New Error Measure

Definition. $(\mathscr{D}, \mathscr{D}')$ are **α-close** $(\alpha \geq 1)$ if

$$\forall x \geq 0 : \qquad \frac{1}{\alpha} \cdot \mathbb{P}_{P \sim \mathscr{D}}(P > \alpha \cdot x) \leq \mathbb{P}_{P' \sim \mathscr{D}'}(P' > x) \leq \alpha \cdot \mathbb{P}_{P \sim \mathscr{D}}(P > x/\alpha)$$

$(\{\mathscr{D}_j\}_{j=1}^{n}, \{\mathscr{D}'_j\}_{j=1}^{n})$ are **α-close** if every pair $(\mathscr{D}_j, \mathscr{D}'_j)$ is α-close

"multiplicative error in upper tails is small"

A New Error Measure

Definition. $(\mathscr{D}, \mathscr{D}')$ are **α-close** $(\alpha \geq 1)$ if

$$\forall x \geq 0 : \qquad \frac{1}{\alpha} \cdot \mathbb{P}_{P \sim \mathscr{D}}(P > \alpha \cdot x) \leq \mathbb{P}_{P' \sim \mathscr{D}'}(P' > x) \leq \alpha \cdot \mathbb{P}_{P \sim \mathscr{D}}(P > x/\alpha)$$

$(\{\mathscr{D}_j\}_{j=1}^n, \{\mathscr{D}_j'\}_{j=1}^n)$ are **α-close** if every pair $(\mathscr{D}_j, \mathscr{D}_j')$ is α-close

"multiplicative error in upper tails is small"

✓Symmetry: can flip roles of $\mathscr{D}, \mathscr{D}'$

A New Error Measure

Definition. $(\mathscr{D}, \mathscr{D}')$ are **α-close** $(\alpha \geq 1)$ if

$$\forall x \geq 0 : \qquad \frac{1}{\alpha} \cdot \mathbb{P}_{P \sim \mathscr{D}}(P > \alpha \cdot x) \leq \mathbb{P}_{P' \sim \mathscr{D}'}(P' > x) \leq \alpha \cdot \mathbb{P}_{P \sim \mathscr{D}}(P > x/\alpha)$$

$(\{\mathscr{D}_j\}_{j=1}^{n}, \{\mathscr{D}'_j\}_{j=1}^{n})$ are **α-close** if every pair $(\mathscr{D}_j, \mathscr{D}'_j)$ is α-close

“multiplicative error in upper tails is small”

✓Symmetry: can flip roles of $\mathscr{D}, \mathscr{D}'$

✓Monotonicity: α-close $\Longrightarrow$ α'-close, $\alpha \leq \alpha'$

A New Error Measure

Definition. $(\mathscr{D}, \mathscr{D}')$ are ***α*-close** $(\alpha \geq 1)$ if

$$\forall x \geq 0 : \qquad \frac{1}{\alpha} \cdot \mathbb{P}_{P \sim \mathscr{D}}(P > \alpha \cdot x) \leq \mathbb{P}_{P' \sim \mathscr{D}'}(P' > x) \leq \alpha \cdot \mathbb{P}_{P \sim \mathscr{D}}(P > x/\alpha)$$

$(\{\mathscr{D}_j\}_{j=1}^n, \{\mathscr{D}_j'\}_{j=1}^n)$ are ***α*-close** if every pair $(\mathscr{D}_j, \mathscr{D}_j')$ is α-close

"multiplicative error in upper tails is small"

✓Symmetry: can flip roles of $\mathscr{D}, \mathscr{D}'$

✓Monotonicity: α-close $\Longrightarrow$ α'-close, $\alpha \leq \alpha'$

✓Composition: "multiplicative triangle inequality"

A New Error Measure

Definition. $(\mathscr{D}, \mathscr{D}')$ are ***α*-close** $(\alpha \geq 1)$ if

$$\forall x \geq 0 : \qquad \frac{1}{\alpha} \cdot \mathbb{P}_{P \sim \mathscr{D}}(P > \alpha \cdot x) \leq \mathbb{P}_{P' \sim \mathscr{D}'}(P' > x) \leq \alpha \cdot \mathbb{P}_{P \sim \mathscr{D}}(P > x/\alpha)$$

$(\{\mathscr{D}_j\}_{j=1}^n, \{\mathscr{D}'_j\}_{j=1}^n)$ are ***α*-close** if every pair $(\mathscr{D}_j, \mathscr{D}'_j)$ is α-close

"multiplicative error in upper tails is small"

✓Symmetry: can flip roles of $\mathscr{D}, \mathscr{D}'$

✓Monotonicity: α-close $\Longrightarrow$ α'-close, $\alpha \leq \alpha'$

✓Composition: "multiplicative triangle inequality"

A New Error Measure

Definition. $(\mathscr{D}, \mathscr{D}')$ are ***α*-close** $(\alpha \geq 1)$ if

$$\forall x \geq 0 : \qquad \frac{1}{\alpha} \cdot \mathbb{P}_{P \sim \mathscr{D}}(P > \alpha \cdot x) \leq \mathbb{P}_{P' \sim \mathscr{D}'}(P' > x) \leq \alpha \cdot \mathbb{P}_{P \sim \mathscr{D}}(P > x/\alpha)$$

$(\{\mathscr{D}_j\}_{j=1}^n, \{\mathscr{D}'_j\}_{j=1}^n)$ are ***α*-close** if every pair $(\mathscr{D}_j, \mathscr{D}'_j)$ is α-close

"multiplicative error in upper tails is small"

- ✓Symmetry: can flip roles of $\mathscr{D}, \mathscr{D}'$
- ✓Monotonicity: α-close $\Longrightarrow \alpha'$-close, $\alpha \leq \alpha'$
- ✓Composition: "multiplicative triangle inequality"

Examples

Examples

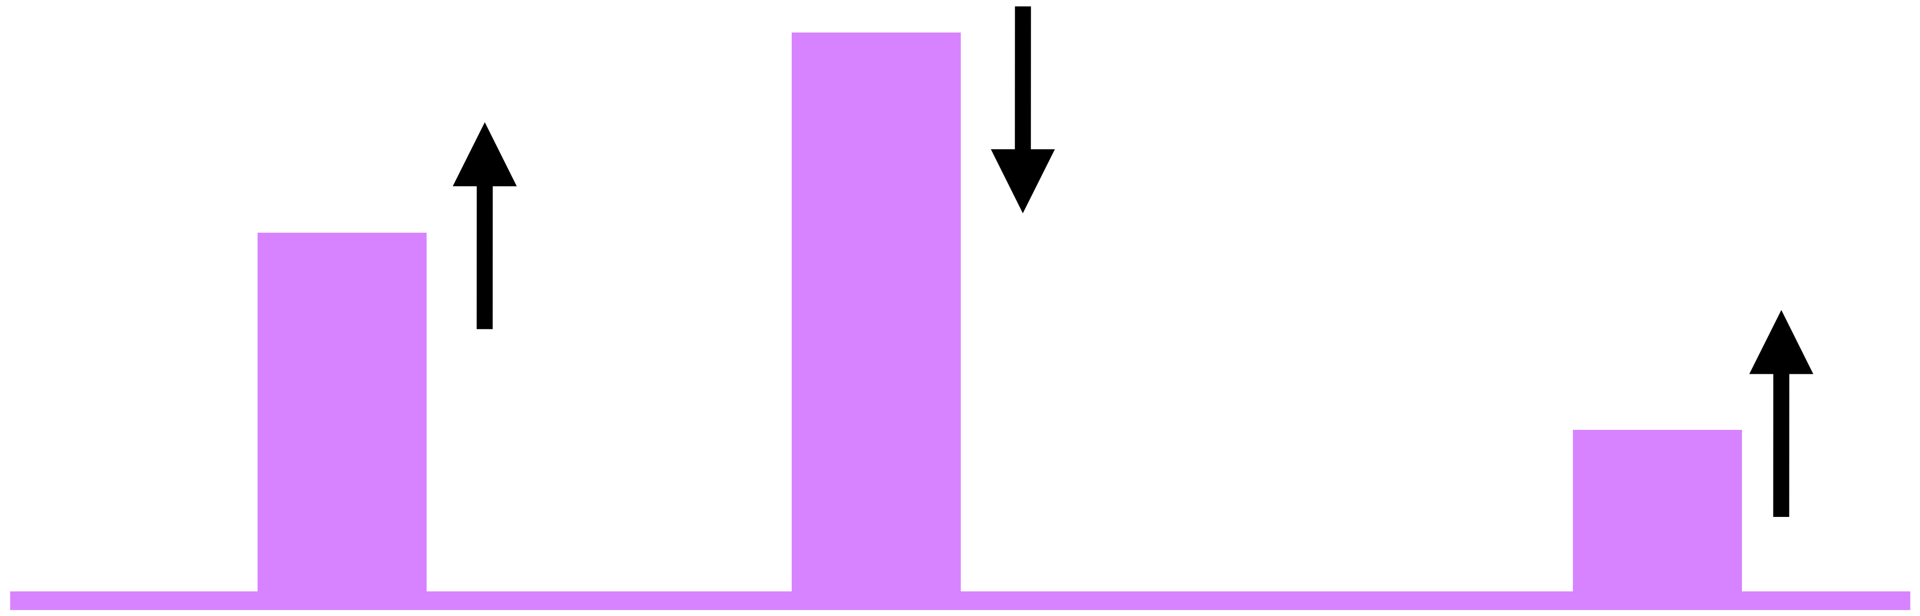

Examples

Examples

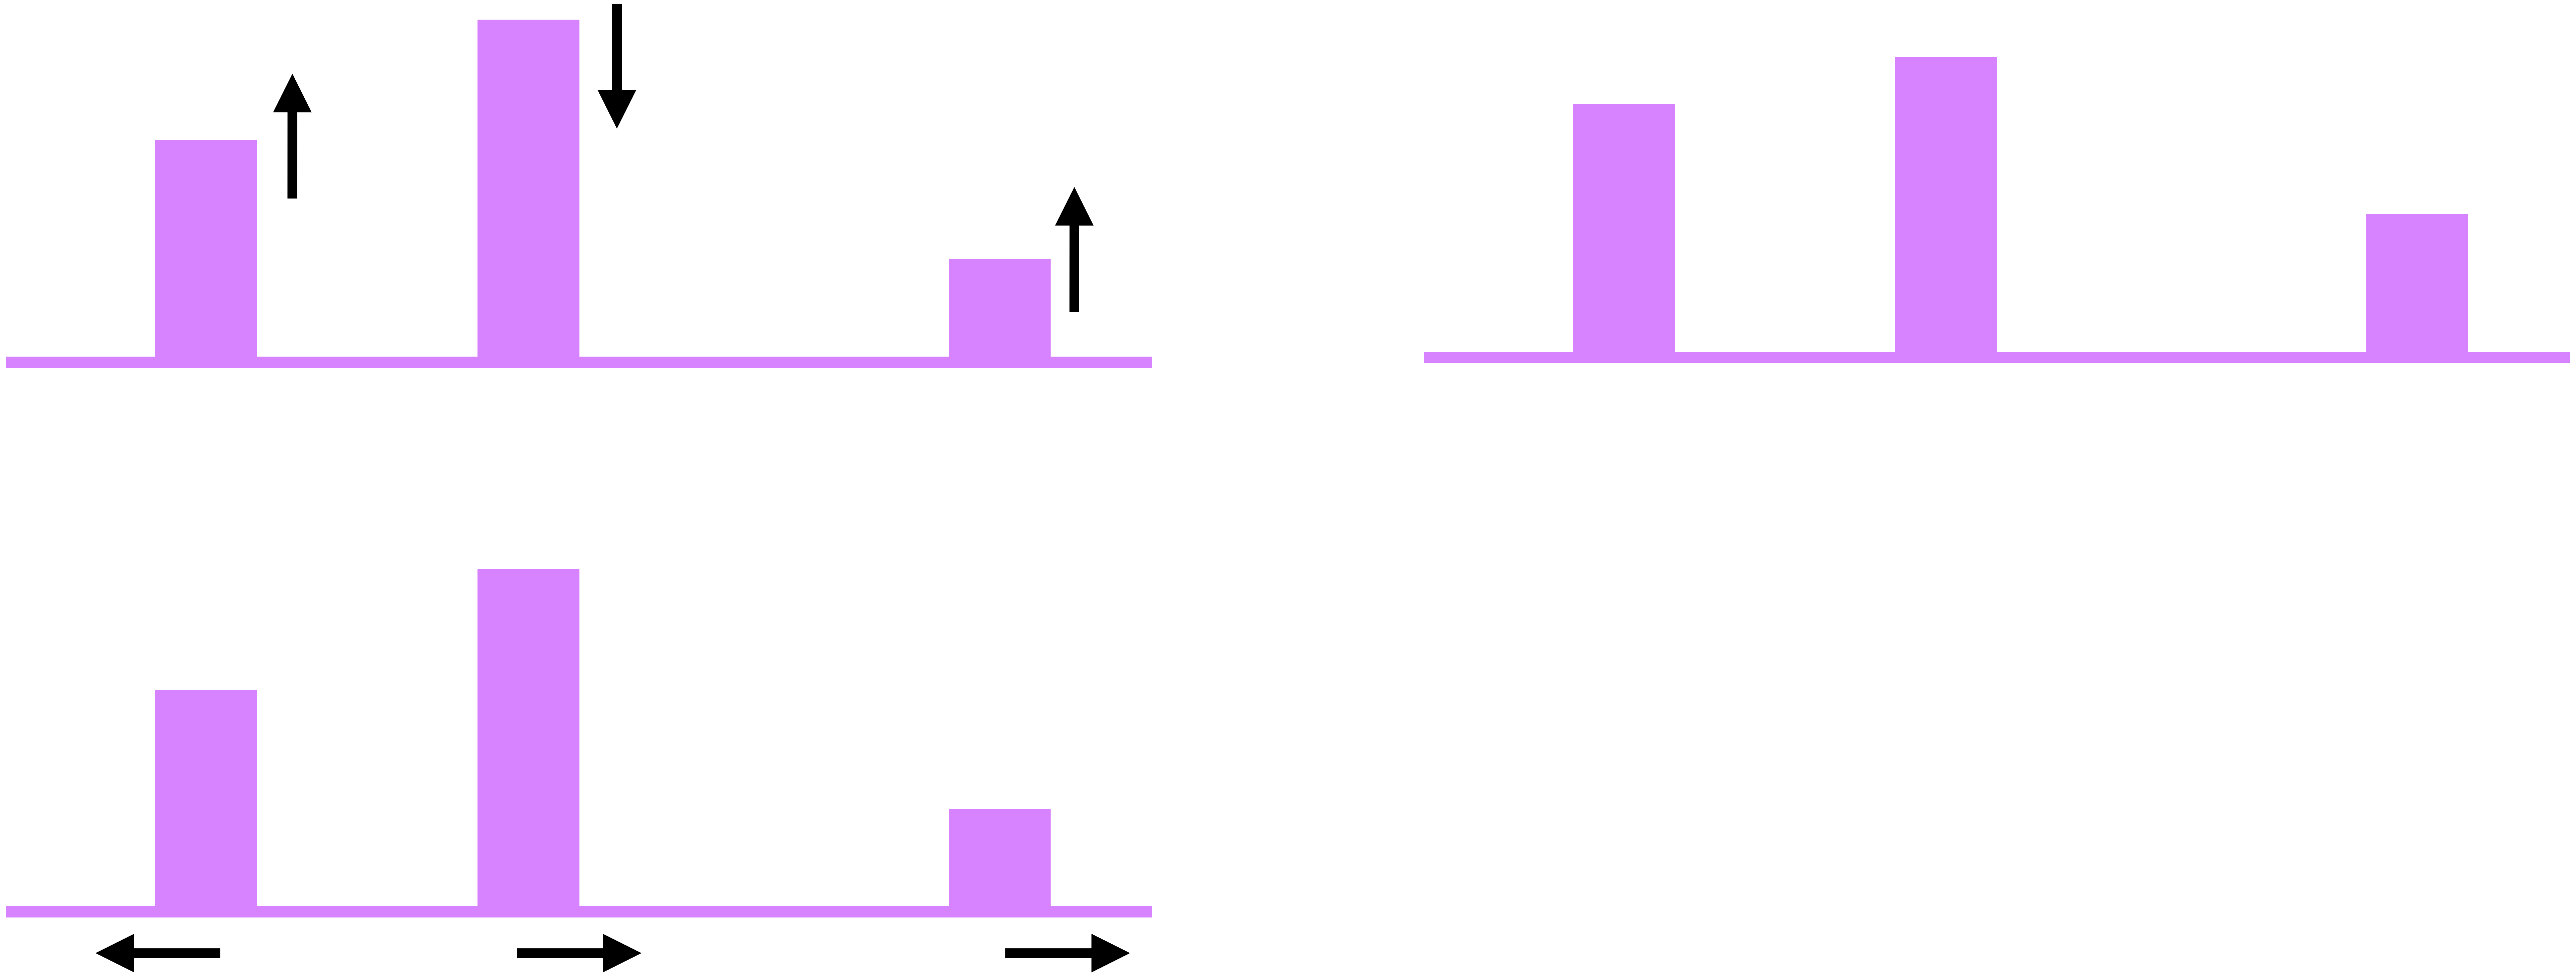

Examples

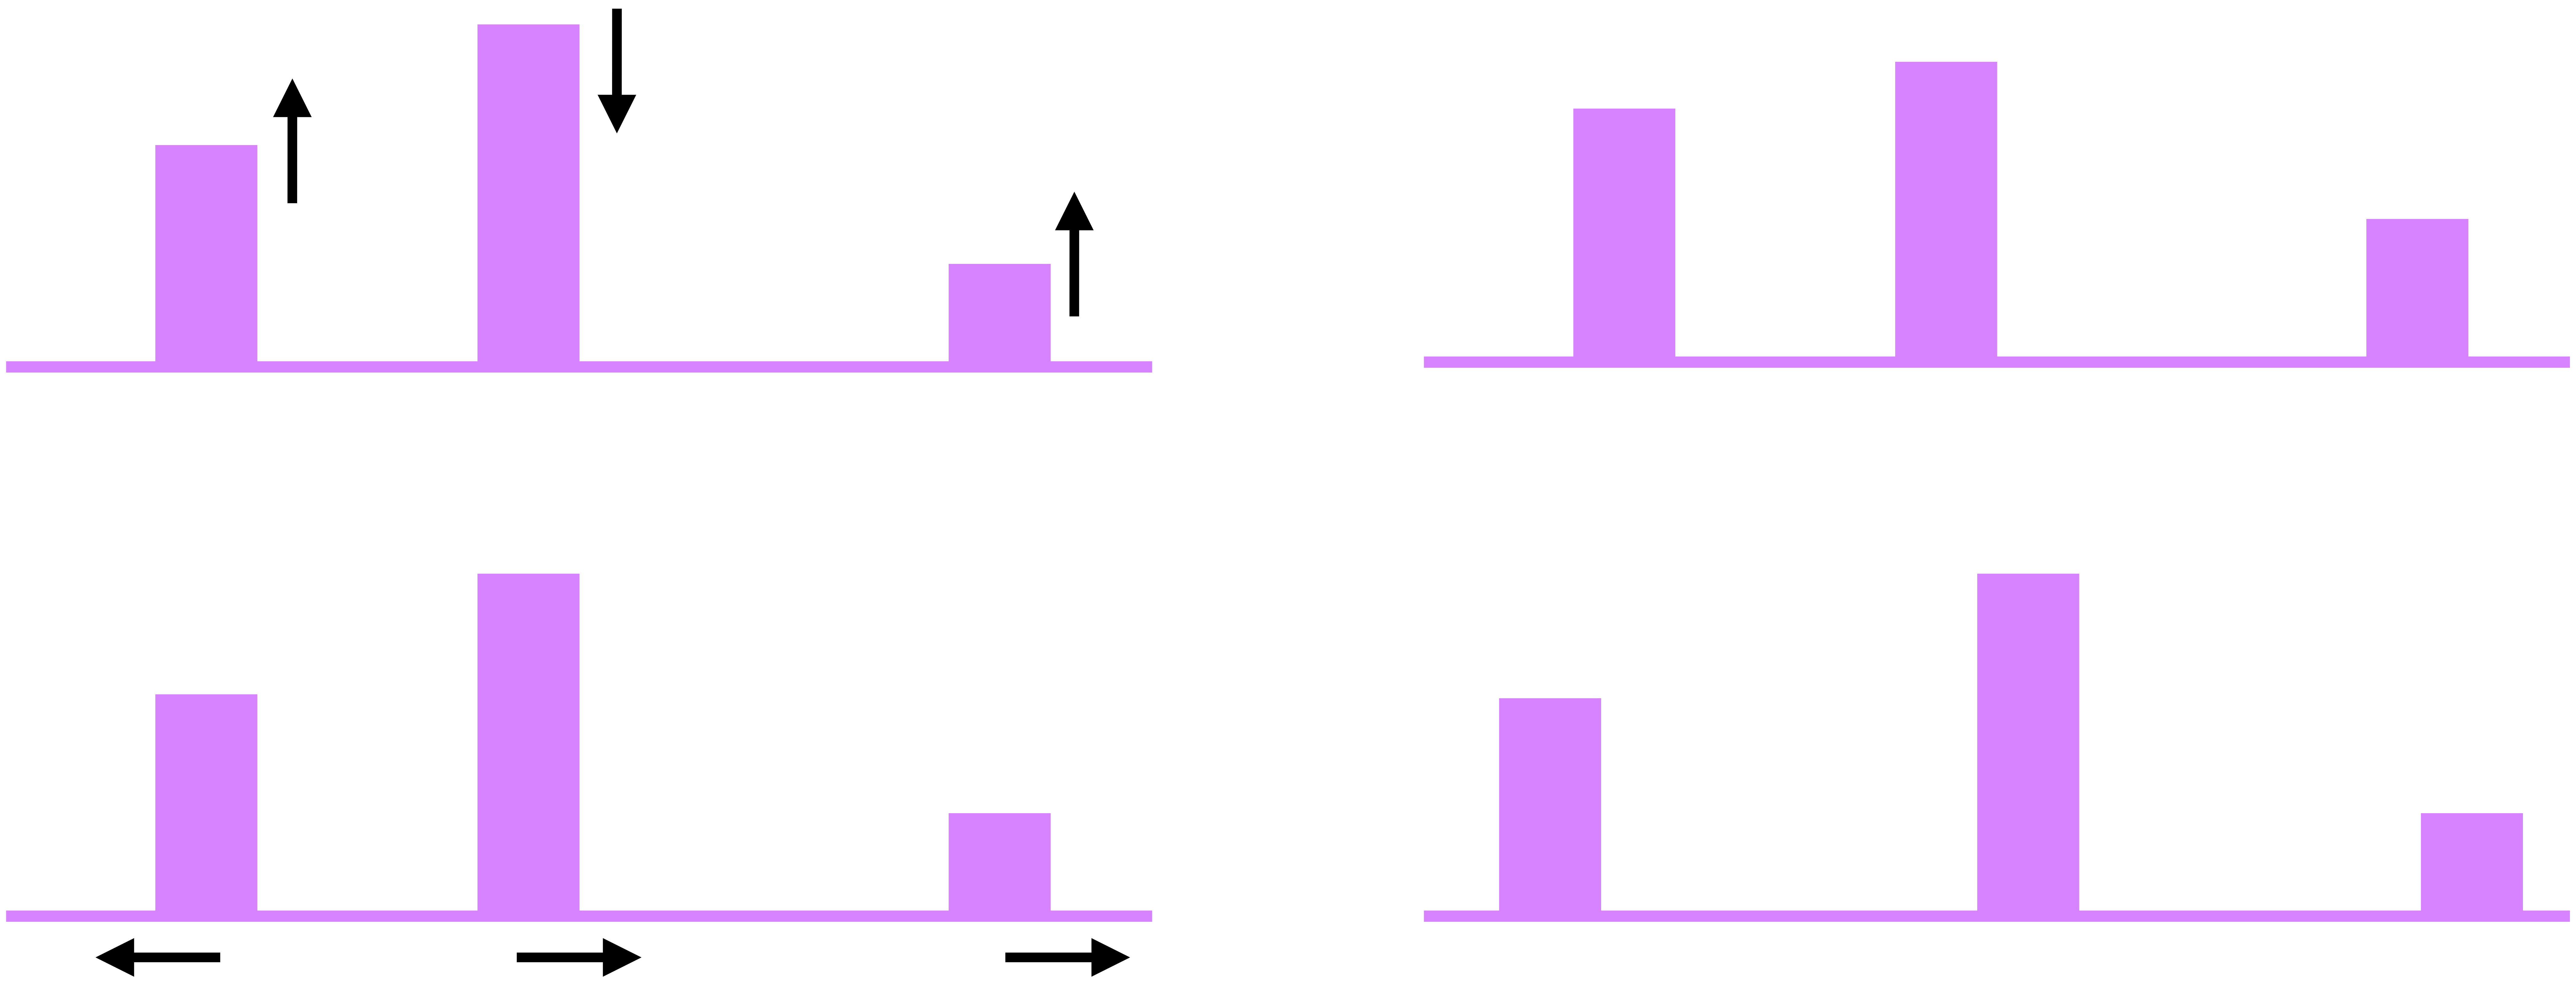

Examples

Examples

...and a less immediately obvious example

Examples

...and a less immediately obvious example

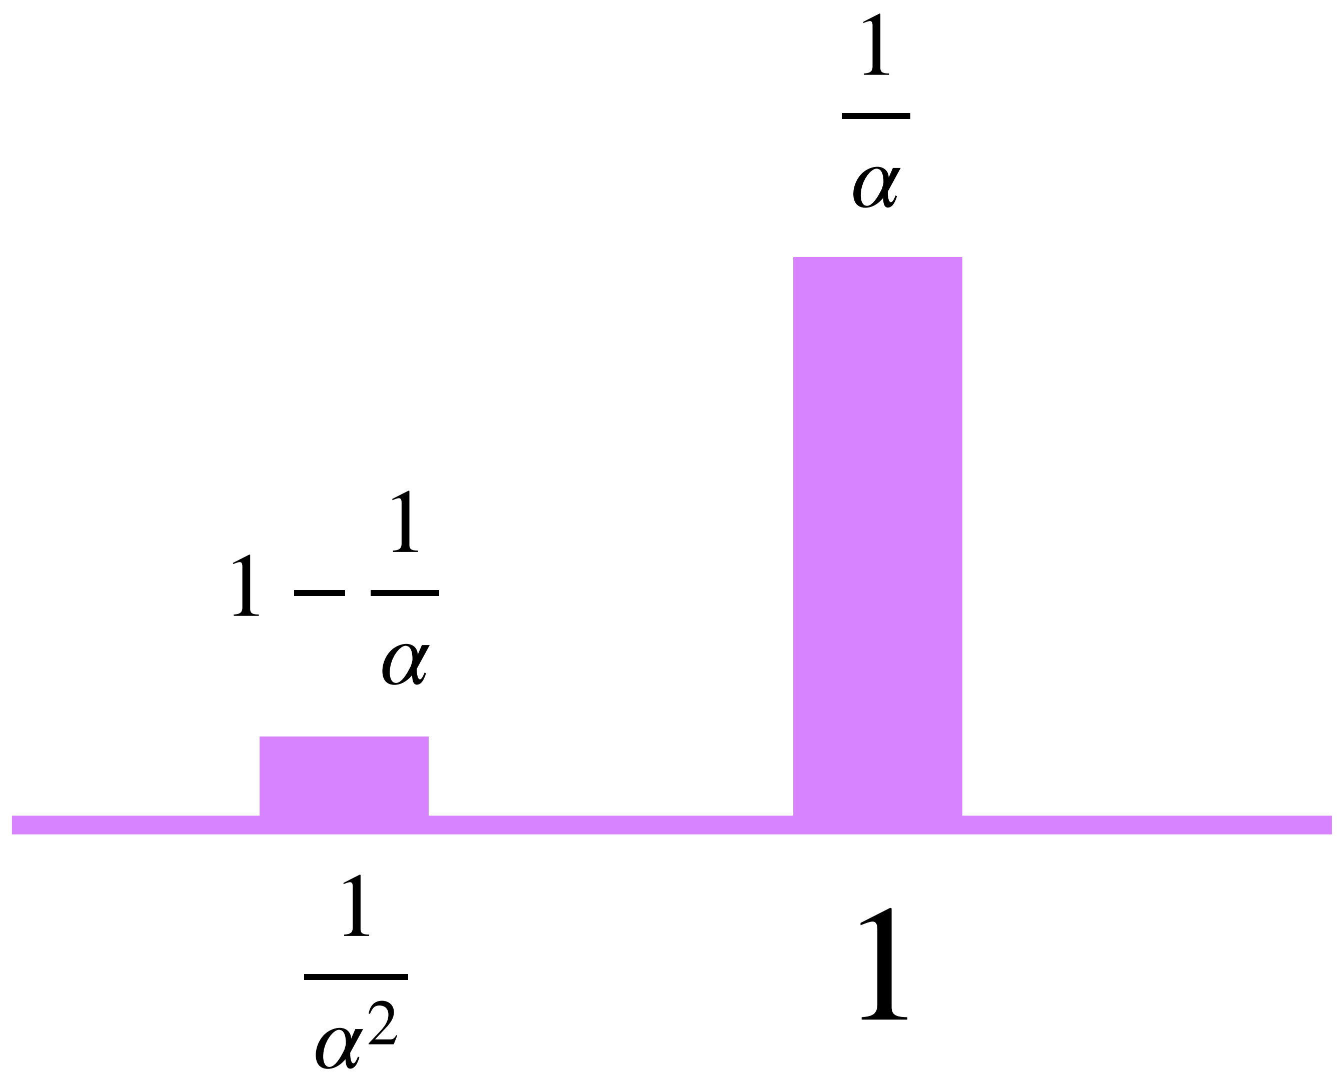

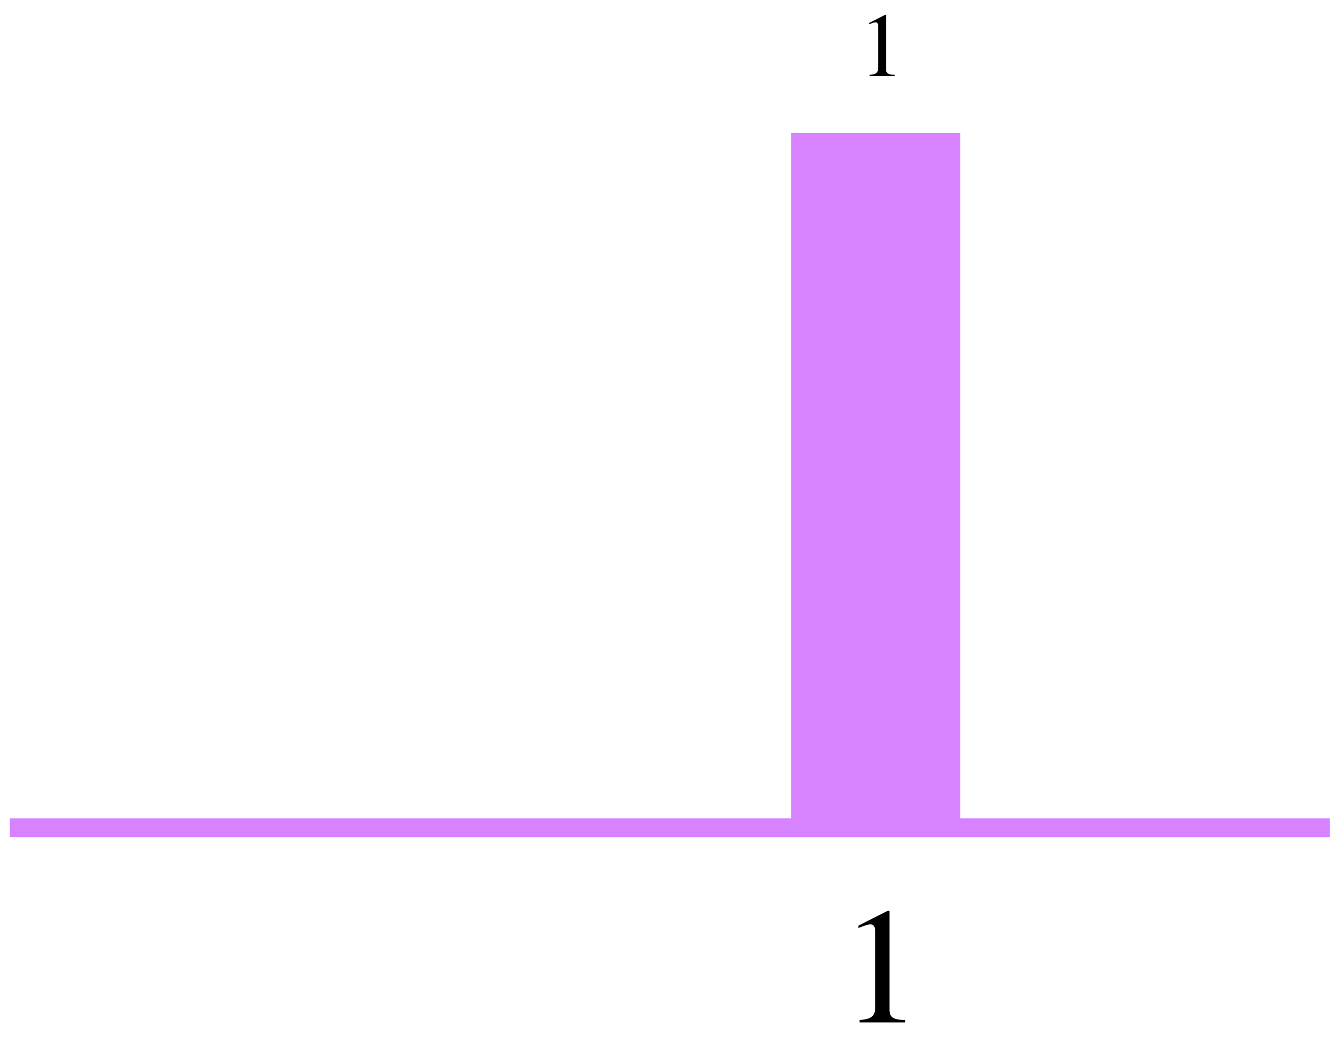

not related by a combined vertical + horizontal shift

Examples

...and a less immediately obvious example

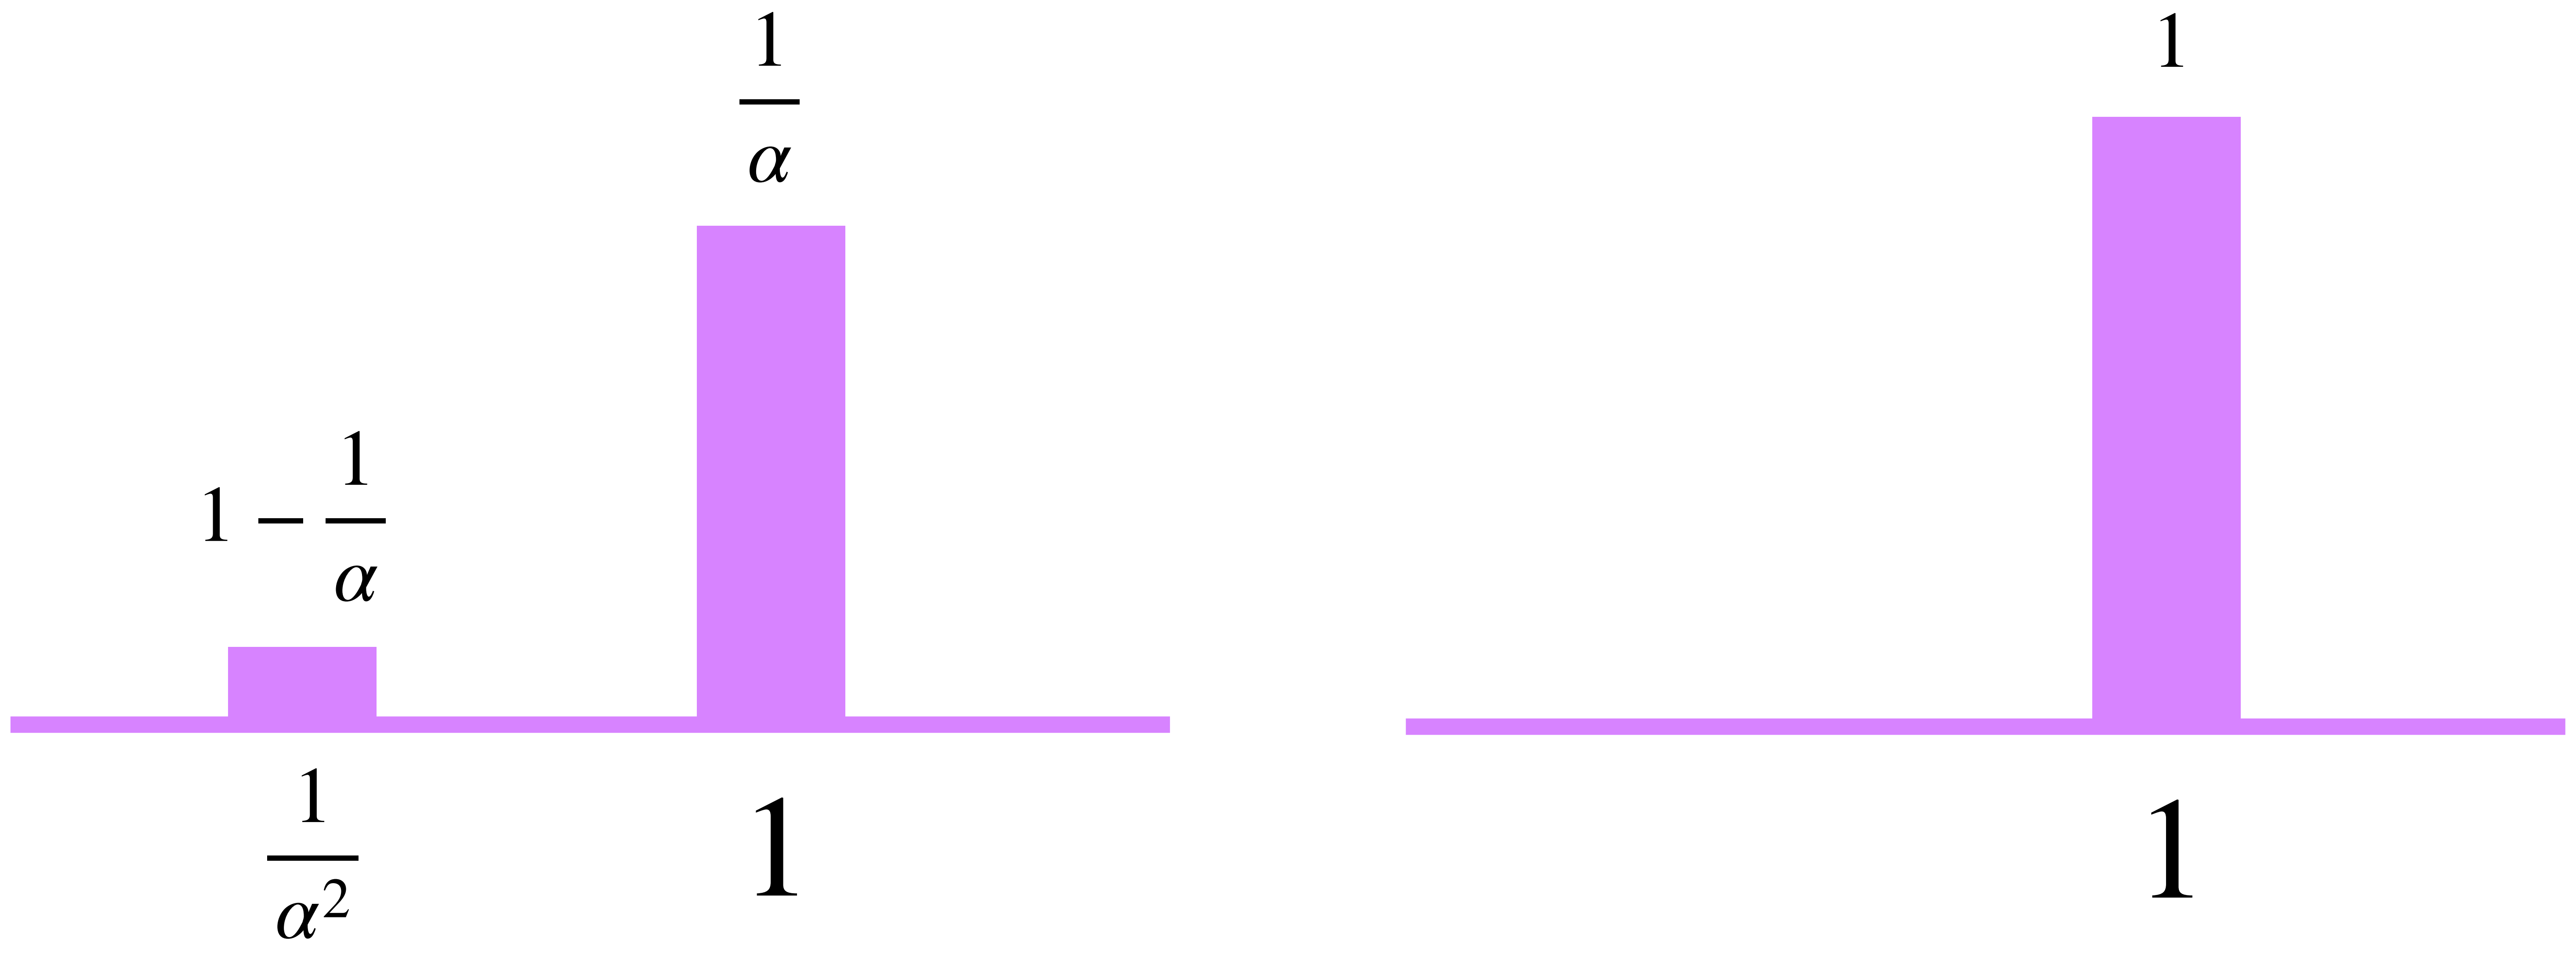

not related by a combined vertical + horizontal shift

parameter error bound → error measure bound

Examples

…and a less immediately obvious example

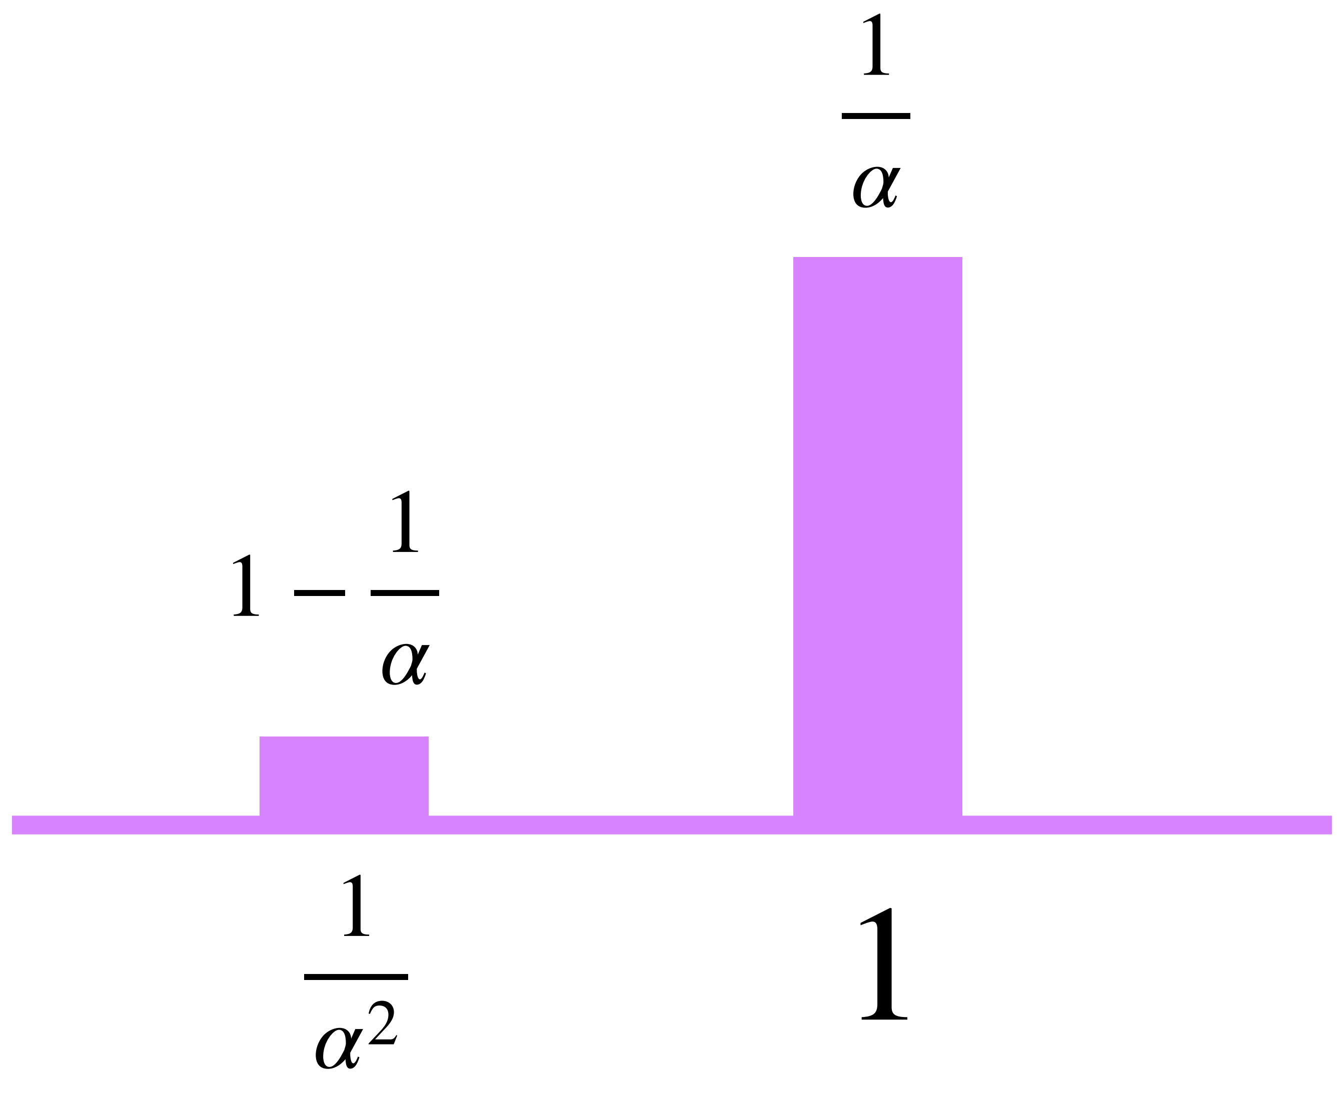

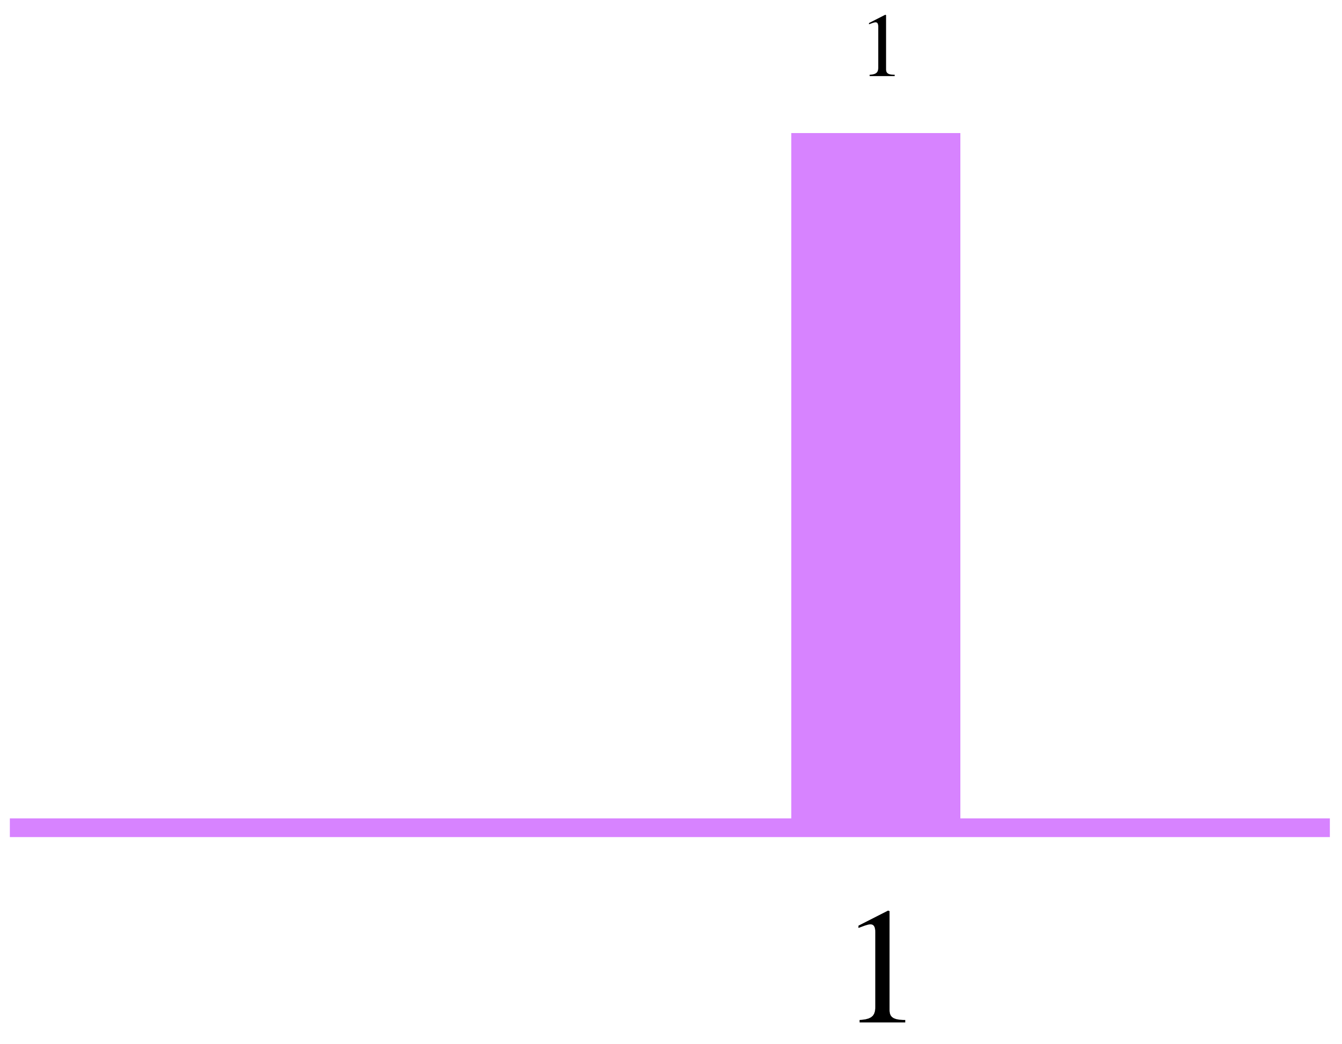

not related by a combined vertical + horizontal shift

parameter error bound → error measure bound

$Exp(\lambda)$ and $Exp(\lambda')$ are $\frac{\lambda'}{\lambda}$ - close

Relation to Other Distances

Relation to Other Distances

Lévy distance

Relation to Other Distances

Lévy distance

$$d_L(\mathscr{D}, \mathscr{D}') = \inf\{\varepsilon \geq 0 \mid F'(x-\varepsilon) - \varepsilon \leq F(x) \leq F'(x+\varepsilon) + \varepsilon \quad \forall x \in \mathbb{R}\}$$

Robust Gittins (RG) Analysis

Theorem. Given an α-close $(\mathscr{I}^* = \{\mathscr{D}_j^*\}_{j=1}^n, \hat{\mathscr{I}} = \{\hat{\mathscr{D}}_j\}_{j=1}^n)$, $\alpha \geq 1$, pair of true and predicted distributions, there is a policy, Robust Gittins (RG), satisfying:

$$\mathrm{RG}(\mathscr{I}^*, \hat{\mathscr{I}}) \leq \alpha^6 \cdot \mathrm{GIPP}(\mathscr{I}^*, \mathscr{I}^*) = \mathrm{OPT}(\mathscr{I}^*)$$

Cost of Robust Gittins on true distributions $\mathscr{D}_j^*$, given predicted distributions $\hat{\mathscr{D}}_j$ a priori

Robust Gittins (RG) Analysis

Theorem. Given an α-close $(\mathcal{I}^* = \{\mathcal{D}_j^*\}_{j=1}^n, \hat{\mathcal{I}} = \{\hat{\mathcal{D}}_j\}_{j=1}^n)$, $\alpha \geq 1$, pair of true and predicted distributions, there is a policy, Robust Gittins (RG), satisfying:

$$\mathrm{RG}(\mathcal{I}^*, \hat{\mathcal{I}}) \leq \alpha^6 \cdot \mathrm{GIPP}(\mathcal{I}^*, \mathcal{I}^*) = \mathrm{OPT}(\mathcal{I}^*)$$

Cost of Robust Gittins on true distributions $\mathcal{D}_j^*$, given predicted distributions $\hat{\mathcal{D}}_j$ a priori

$$\mathrm{RG}(\mathcal{I}^*, \hat{\mathcal{I}}) \leq \alpha^3 \cdot \mathrm{GIPP}(\hat{\mathcal{I}}, \hat{\mathcal{I}}) \leq \alpha^3 \cdot \mathrm{RG}(\hat{\mathcal{I}}, \mathcal{I}^*) \leq \alpha^6 \cdot \mathrm{GIPP}(\mathcal{I}^*, \mathcal{I}^*)$$

Robust Gittins (RG) Analysis

Theorem. Given an α-close $(\mathcal{I}^* = \{\mathcal{D}_j^*\}_{j=1}^n, \hat{\mathcal{I}} = \{\hat{\mathcal{D}}_j\}_{j=1}^n)$, $\alpha \geq 1$, pair of true and predicted distributions, there is a policy, Robust Gittins (RG), satisfying:

$$\mathrm{RG}(\mathcal{I}^*, \hat{\mathcal{I}}) \leq \alpha^6 \cdot \mathrm{GIPP}(\mathcal{I}^*, \mathcal{I}^*) = \mathrm{OPT}(\mathcal{I}^*)$$

Cost of Robust Gittins on true distributions $\mathcal{D}_j^*$, given predicted distributions $\hat{\mathcal{D}}_j$ a priori

$$\mathrm{RG}(\mathcal{I}^*, \hat{\mathcal{I}}) \overset{!}{\leq} \alpha^3 \cdot \mathrm{GIPP}(\hat{\mathcal{I}}, \hat{\mathcal{I}}) \leq \alpha^3 \cdot \mathrm{RG}(\hat{\mathcal{I}}, \mathcal{I}^*) \leq \alpha^6 \cdot \mathrm{GIPP}(\mathcal{I}^*, \mathcal{I}^*)$$

Robust Gittins (RG) Analysis

Theorem. Given an α-close $(\mathscr{I}^* = \{\mathscr{D}_j^*\}_{j=1}^n, \hat{\mathscr{I}} = \{\hat{\mathscr{D}}_j\}_{j=1}^n)$, $\alpha \geq 1$, pair of true and predicted distributions, there is a policy, Robust Gittins (RG), satisfying:

$$\mathrm{RG}(\mathscr{I}^*, \hat{\mathscr{I}}) \leq \alpha^6 \cdot \mathrm{GIPP}(\mathscr{I}^*, \mathscr{I}^*) = \mathrm{OPT}(\mathscr{I}^*)$$

Cost of Robust Gittins on true distributions $\mathscr{D}_j^*$, given predicted distributions $\hat{\mathscr{D}}_j$ a priori

$$\mathrm{RG}(\mathscr{I}^*, \hat{\mathscr{I}}) \overset{!}{\leq} \alpha^3 \cdot \underset{\substack{\parallel \\ \mathrm{OPT}(\hat{\mathscr{I}})}}{\mathrm{GIPP}(\hat{\mathscr{I}}, \hat{\mathscr{I}})} \leq \alpha^3 \cdot \mathrm{RG}(\hat{\mathscr{I}}, \mathscr{I}^*) \leq \alpha^6 \cdot \mathrm{GIPP}(\mathscr{I}^*, \mathscr{I}^*)$$

Robust Gittins (RG) Analysis

Theorem. Given an α-close $(\mathscr{I}^* = \{\mathscr{D}_j^*\}_{j=1}^n, \hat{\mathscr{I}} = \{\hat{\mathscr{D}}_j\}_{j=1}^n)$, $\alpha \geq 1$, pair of true and predicted distributions, there is a policy, Robust Gittins (RG), satisfying:

$$\mathrm{RG}(\mathscr{I}^*, \hat{\mathscr{I}}) \leq \alpha^6 \cdot \mathrm{GIPP}(\mathscr{I}^*, \mathscr{I}^*) = \mathrm{OPT}(\mathscr{I}^*)$$

Cost of Robust Gittins on true distributions $\mathscr{D}_j^*$, given predicted distributions $\hat{\mathscr{D}}_j$ a priori

$$\mathrm{RG}(\mathscr{I}^*, \hat{\mathscr{I}}) \overset{!}{\leq} \alpha^3 \cdot \underset{\overset{\|}{\mathrm{OPT}(\hat{\mathscr{I}})}}{\mathrm{GIPP}(\hat{\mathscr{I}}, \hat{\mathscr{I}})} \leq \alpha^3 \cdot \mathrm{RG}(\hat{\mathscr{I}}, \mathscr{I}^*) \underset{\uparrow}{\leq} \alpha^6 \cdot \mathrm{GIPP}(\mathscr{I}^*, \mathscr{I}^*)$$

by symmetry of error measure

Robust Gittins (RG) Analysis

Theorem. Given an α-close $(\mathcal{J}^* = \{\mathcal{D}_j^*\}_{j=1}^n, \hat{\mathcal{J}} = \{\hat{\mathcal{D}}_j\}_{j=1}^n)$, $\alpha \geq 1$, pair of true and predicted distributions, there is a policy, Robust Gittins (RG), satisfying:

$$\mathrm{RG}(\mathcal{J}^*, \hat{\mathcal{J}}) \leq \alpha^6 \cdot \mathrm{GIPP}(\mathcal{J}^*, \mathcal{J}^*) = \mathrm{OPT}(\mathcal{J}^*)$$

Cost of Robust Gittins on true distributions $\mathcal{D}_j^*$, given predicted distributions $\hat{\mathcal{D}}_j$ a priori

$$\mathrm{RG}(\mathcal{J}^*, \hat{\mathcal{J}}) \overset{!}{\leq} \alpha^3 \cdot \underset{\substack{\| \\ \mathrm{OPT}(\hat{\mathcal{J}})}}{\mathrm{GIPP}(\hat{\mathcal{J}}, \hat{\mathcal{J}})} \leq \alpha^3 \cdot \mathrm{RG}(\hat{\mathcal{J}}, \mathcal{J}^*) \underset{\substack{\uparrow \\ \text{by symmetry of} \\ \text{error measure}}}{\leq} \alpha^6 \cdot \underset{\substack{\| \\ \mathrm{OPT}(\mathcal{J}^*)}}{\mathrm{GIPP}(\mathcal{J}^*, \mathcal{J}^*)}$$

Robust Gittins (RG) Analysis

Theorem. Given an α-close $(\mathscr{J}^* = \{\mathscr{D}_j^*\}_{j=1}^n, \hat{\mathscr{J}} = \{\hat{\mathscr{D}}_j\}_{j=1}^n)$, $\alpha \geq 1$, pair of true and predicted distributions, there is a policy, Robust Gittins (RG), satisfying:

$$\mathrm{RG}(\mathscr{J}^*, \hat{\mathscr{J}}) \leq \alpha^6 \cdot \mathrm{GIPP}(\mathscr{J}^*, \mathscr{J}^*) = \mathrm{OPT}(\mathscr{J}^*)$$

Cost of Robust Gittins on true distributions $\mathscr{D}_j^*$, given predicted distributions $\hat{\mathscr{D}}_j$ a priori

$$\mathrm{RG}(\mathscr{J}^*, \hat{\mathscr{J}}) \overset{!}{\leq} \alpha^3 \cdot \underset{\overset{\|}{\mathrm{OPT}(\hat{\mathscr{J}})}}{\mathrm{GIPP}(\hat{\mathscr{J}}, \hat{\mathscr{J}})} \leq \alpha^3 \cdot \mathrm{RG}(\hat{\mathscr{J}}, \mathscr{J}^*) \underset{\overset{\uparrow}{\text{by symmetry of error measure}}}{\leq} \alpha^6 \cdot \underset{\overset{\|}{\mathrm{OPT}(\mathscr{J}^*)}}{\mathrm{GIPP}(\mathscr{J}^*, \mathscr{J}^*)}$$

Note: proves $\mathrm{OPT}(\hat{\mathscr{J}})$ and $\mathrm{OPT}(\mathscr{J}^*)$ off by factor of at most α^3

Robust Gittins (RG) Analysis

Theorem. Given an α-close ($\mathcal{I}^* = \{\mathcal{D}_j^*\}_{j=1}^n, \hat{\mathcal{I}} = \{\hat{\mathcal{D}}_j\}_{j=1}^n$), $\alpha \geq 1$, pair of true and predicted distributions, there is a policy, Robust Gittins (RG), satisfying:

$$\mathrm{RG}(\mathcal{I}^*, \hat{\mathcal{I}}) \leq \alpha^6 \cdot \mathrm{GIPP}(\mathcal{I}^*, \mathcal{I}^*) = \mathrm{OPT}(\mathcal{I}^*)$$

Cost of Robust Gittins on true distributions $\mathcal{D}_j^*$, given predicted distributions $\hat{\mathcal{D}}_j$ a priori

Robust Gittins (RG) Analysis

Theorem. Given an α-close $(\mathscr{I}^* = \{\mathscr{D}_j^*\}_{j=1}^n, \hat{\mathscr{I}} = \{\hat{\mathscr{D}}_j\}_{j=1}^n)$, $\alpha \geq 1$, pair of true and predicted distributions, there is a policy, Robust Gittins (RG), satisfying:

$$\mathrm{RG}(\mathscr{I}^*, \hat{\mathscr{I}}) \leq \alpha^6 \cdot \mathrm{GIPP}(\mathscr{I}^*, \mathscr{I}^*) = \mathrm{OPT}(\mathscr{I}^*)$$

Cost of Robust Gittins on true distributions $\mathscr{D}_j^*$, given predicted distributions $\hat{\mathscr{D}}_j$ a priori

Closed-form cost of Gittins (Megow & Vredeveld '14)

Robust Gittins (RG) Analysis

Theorem. Given an α-close $(\mathscr{I}^* = \{\mathscr{D}_j^*\}_{j=1}^n, \hat{\mathscr{I}} = \{\hat{\mathscr{D}}_j\}_{j=1}^n)$, $\alpha \geq 1$, pair of true and predicted distributions, there is a policy, Robust Gittins (RG), satisfying:

$$\mathrm{RG}(\mathscr{I}^*, \hat{\mathscr{I}}) \leq \alpha^6 \cdot \mathrm{GIPP}(\mathscr{I}^*, \mathscr{I}^*) = \mathrm{OPT}(\mathscr{I}^*)$$

Cost of Robust Gittins on true distributions $\mathscr{D}_j^*$, given predicted distributions $\hat{\mathscr{D}}_j$ a priori

Closed-form cost of Gittins (Megow & Vredeveld '14)

$$\mathrm{GIPP}(\mathscr{I}, \mathscr{I}) = \sum_{j=1}^{n} \sum_{i=1}^{n_j} \sum_{(k, q_{k,l}) \in H'(j,i)} \mathbb{E}\left[\mathbf{1}_{\{P_j > y_{j,i}\}} \cdot \mathbf{1}_{\{P_k > y_{k,l}\}} \cdot \min\{P_k - y_{k,l}, q_{k,l}\}\right]$$

Robust Gittins (RG) Analysis

Robust Gittins (RG) Analysis

Upshot: new stochastic scheduling model robust to errors in predicted distributions

Robust Gittins (RG) Analysis

Upshot: new stochastic scheduling model robust to errors in predicted distributions

✓Unknown α ?

Robust Gittins (RG) Analysis

Upshot: new stochastic scheduling model robust to errors in predicted distributions

✓Unknown α ?

✓Other problems where (something like) Gittins index is optimal?