Recursive Rational Inattention Is Entropic

Henrique De Oliveira

São Paulo School

Tetsuya Hoshino

ITAM

AMES 2022

Information allows DM to make a better decision

Information may consume DM's attention

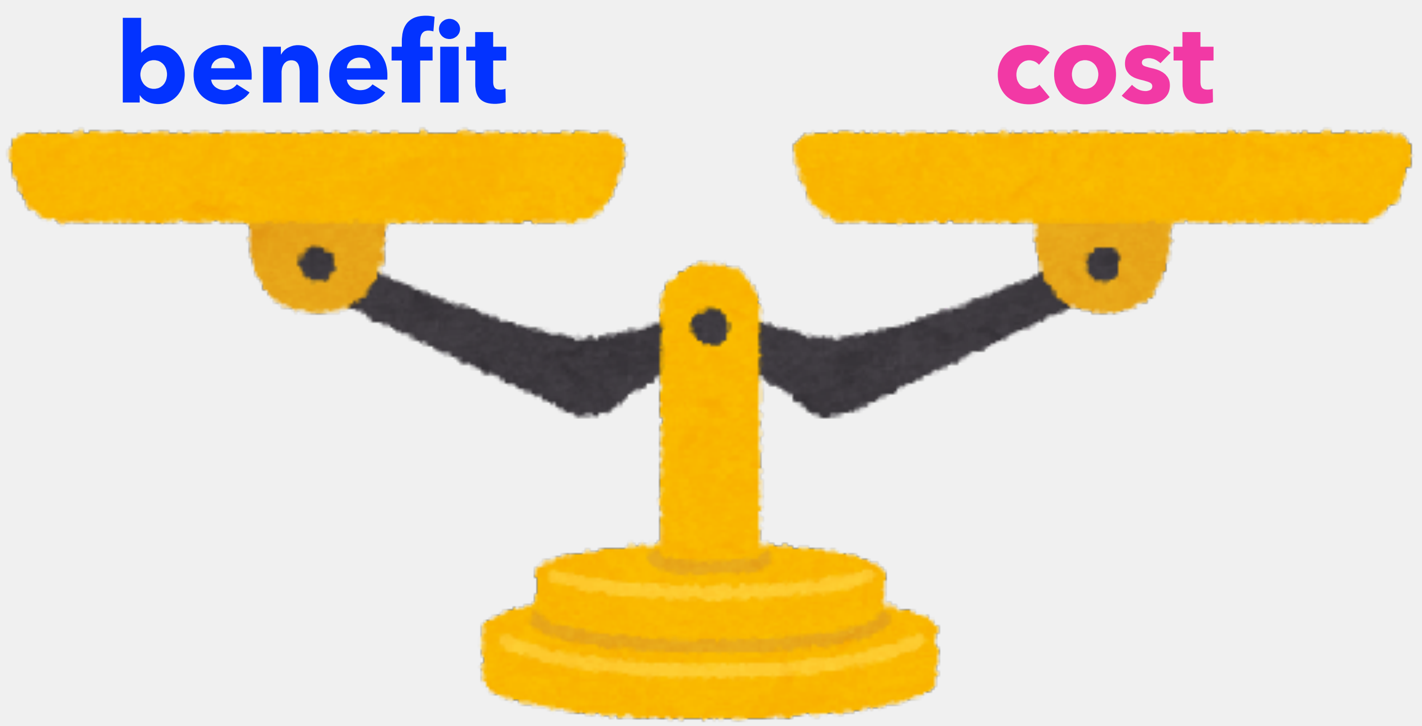

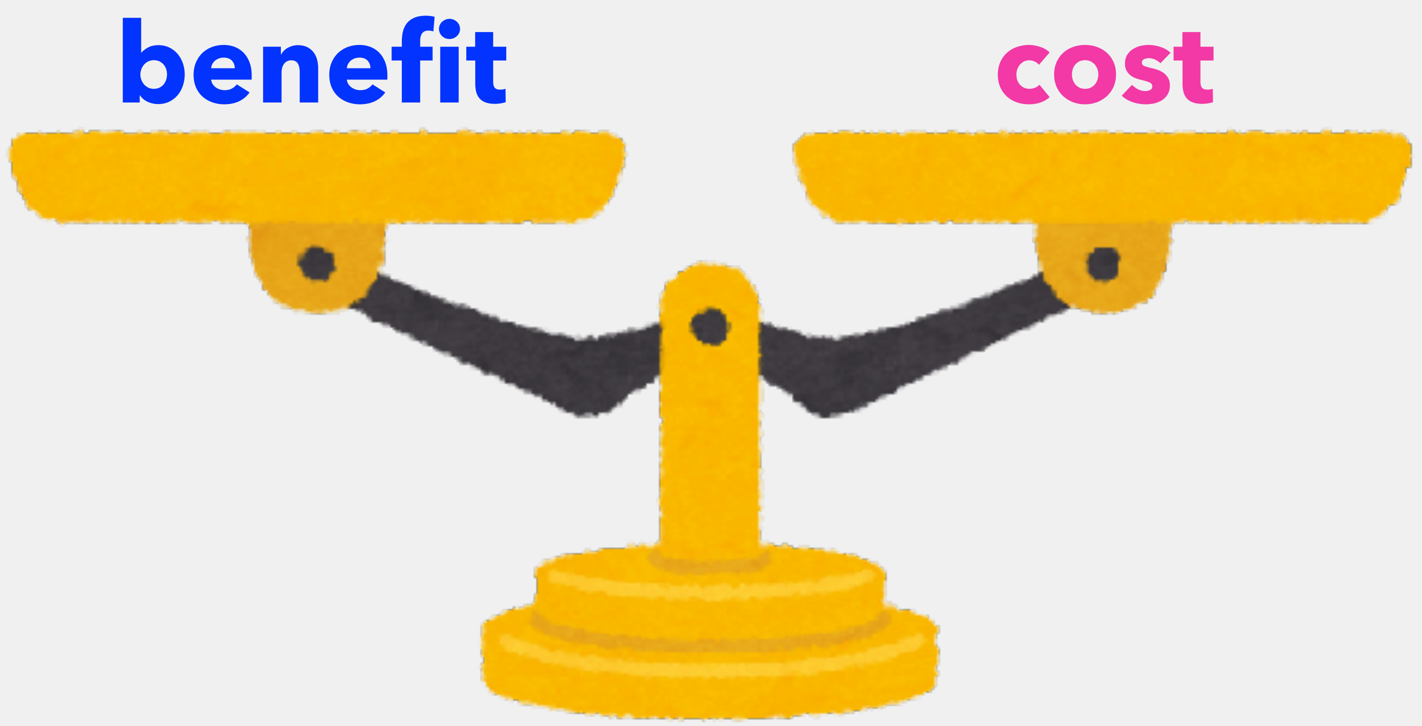

DM balances the **benefit** and **cost**

DM acquires costly information about a payoff-relevant state before choosing an action

$$\mathbb{E}\left[\sum_{\text{time}} \text{discounted}\left(\text{utilities} - \text{information costs}\right)\right]$$

DM acquires costly information about a payoff-relevant state before choosing an action

$$\mathbb{E}\left[\sum_{\text{time}} \text{discounted}\left(\text{utilities} - \text{information costs}\right)\right]$$

DM acquires costly information about a payoff-relevant state before choosing an action

$$\mathbb{E}\left[\sum_{\text{time}} \text{discounted}\left(\text{utilities} - \text{information costs}\right)\right]$$

How should we model the costs?

cost function

preference/choice

cost function

preference/choice

What cost function satisfies desirable properties?

- Sims (2003)
- Bloedel–Zhong (2021)
- Hebert–Woodford (2020)
- Pomatto–Strack–Tamuz (2022)

cost function

preference/choice

What cost function rationalizes given preference/choice?

- Caplin–Dean (2015)
- Caplin–Dean–Leahy (2022)
- De Oliveora–Denti–Mihm–Ozbek (2017)
- De Oliveora (2017)
- Denti (2022)

Existing Approaches

cost function

preference/choice

static

cost function

preference/choice

static

dynamic

What cost functions make the value functions **recursive**?

What costs allow us to use the dynamic programming?

Example: Entropic Costs

Main Result

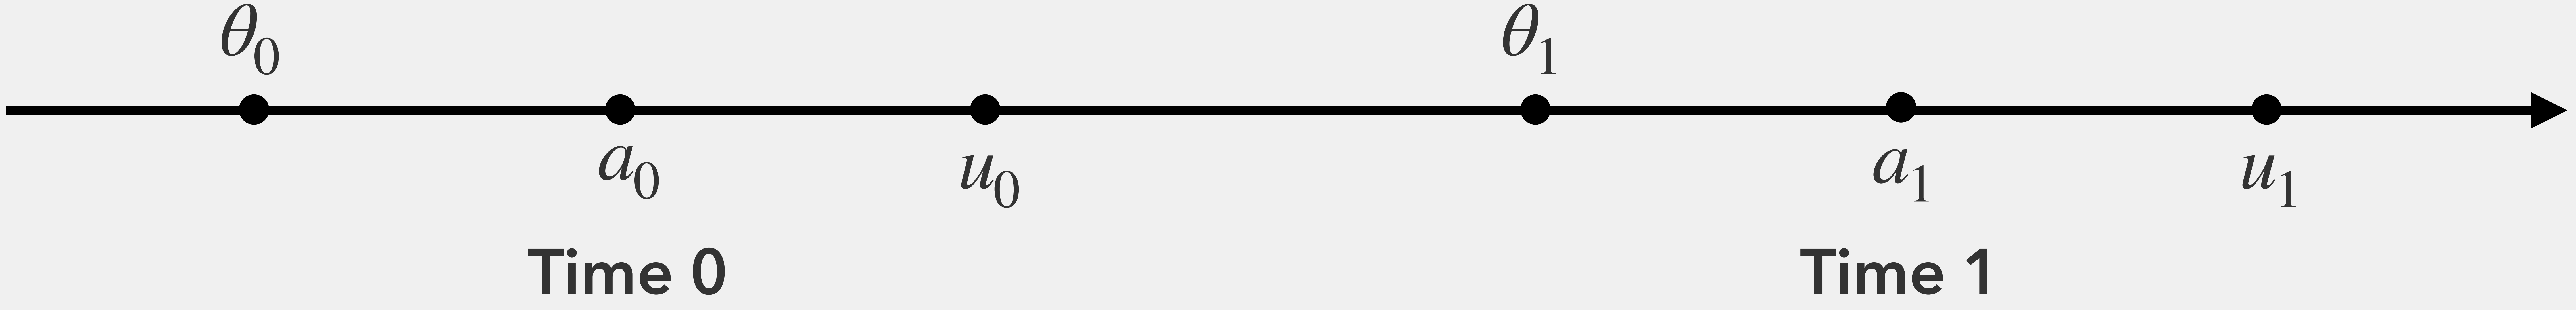

Time is discrete: $t = 0,1,2,\ldots$

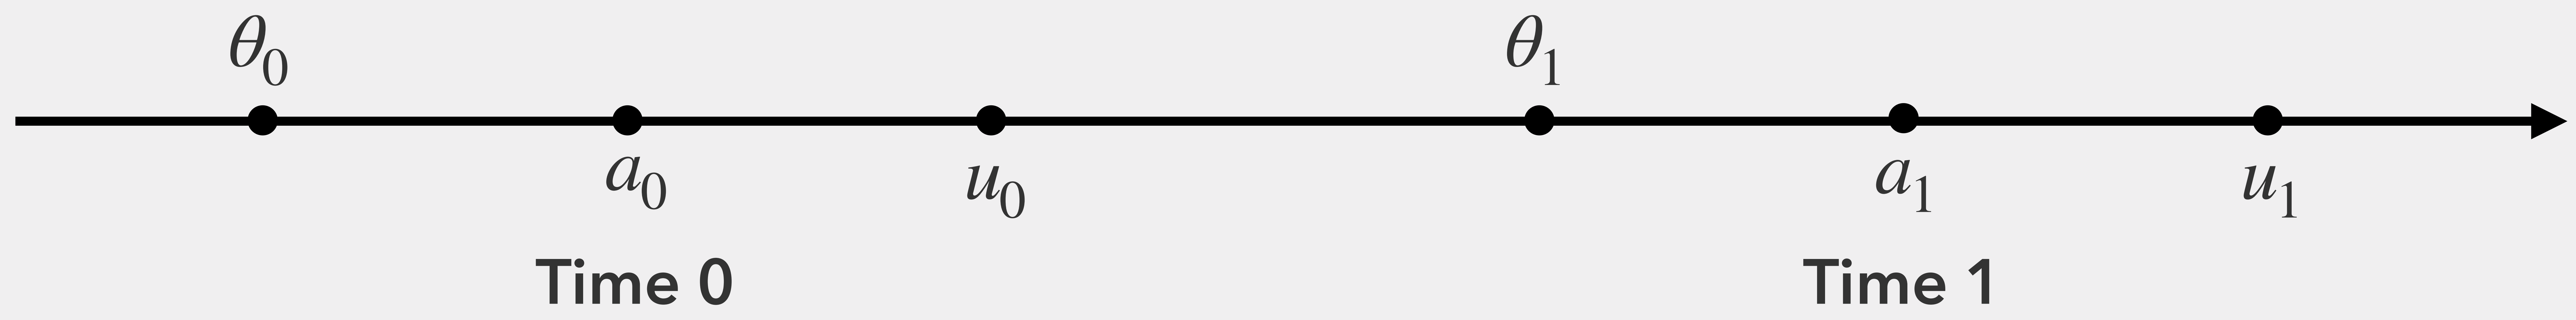

Nature draws a state $\boldsymbol{\theta}_0$ according to a prior $\boldsymbol{\pi}$

- DM does **not** observe a state realization θ_0

DM learns a **costly signal** $\boldsymbol{x}_0$ about θ_0

In today's talk, I focus on a "direct" signal $\boldsymbol{a}_0$ rather than an arbitrary signal $\boldsymbol{x}_0$

- DM chooses distributions $p_0(a_0 \mid \theta_0) \in \Delta(A)$ of **action recommendations** $\boldsymbol{a}_0$
- DM takes the recommended action a_0

Payoff $u_0(\theta_0, a_0)$ is realized

- DM observes the payoff $u_0 \equiv u_0(\theta_0, a_0)$

At the end of time 0 (= the beginning of time 1), DM knows a decision node $z^1 = (a_0, u_0)$

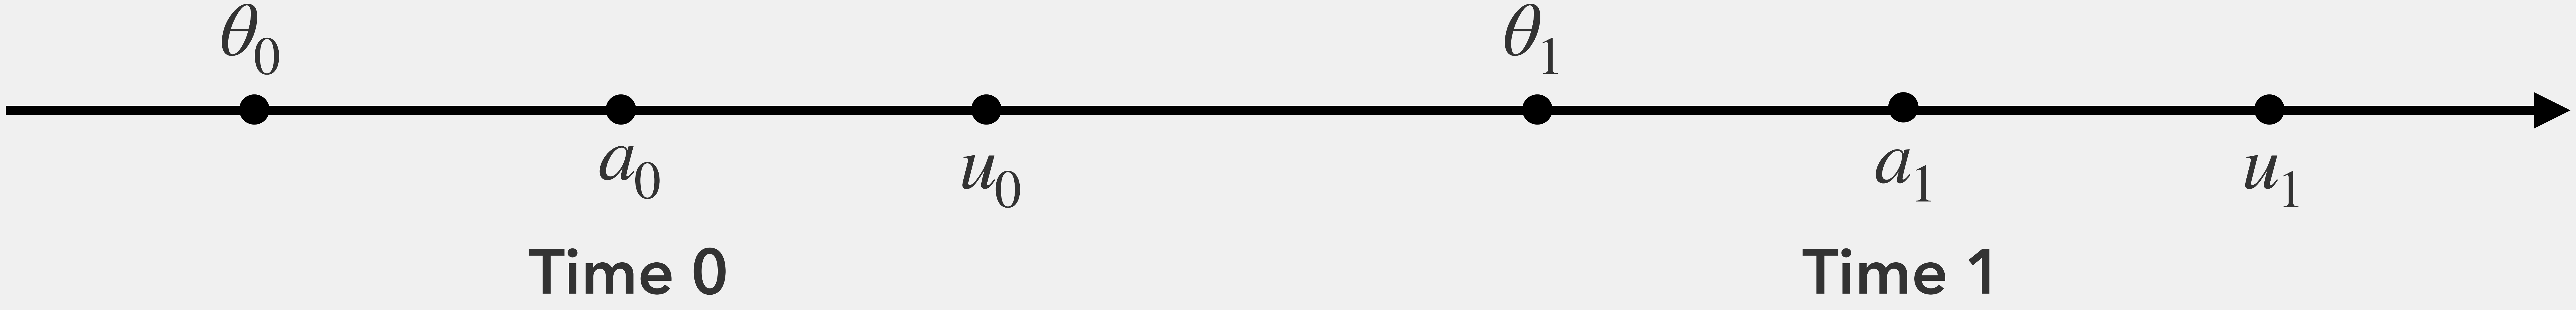

Nature draws a state $\boldsymbol{\theta}_1$ according to a prior $\pi(\cdot \mid \theta_0)$

- DM does **not** observe a state realization $\theta^1 = (\theta_0, \theta_1)$

DM learns an **action recommendation $\boldsymbol{a}_1$**

- DM chooses distributions $p_1(a_1 \mid \theta^1, z^1) \in \Delta(A)$
- DM takes the recommended action a_1

Payoff $u_1(\theta^1, a^1)$ is realized

- DM observes the payoff $u_1 \equiv u_1(\theta^1, a^1)$

At the end of time 1 (= the beginning of time 2), DM knows a decision node $z^2 = (a^1, u^1)$

The **entropic cost** is linear in the mutual information I:

$$\left(\text{unit cost } \lambda\right) \times I\left(\boldsymbol{\theta}^t, \boldsymbol{a}_t \mid z^{t-1}\right)$$

The **mutual information** I quantifies the "amount of information" about states $\boldsymbol{\theta}^t$ that DM obtains from observing signal $\boldsymbol{a}_t$ conditional on z^{t-1}

Formally

$$I\left(\boldsymbol{\theta}^t, \boldsymbol{a}_t \mid z^{t-1}\right) = \mathbb{E}_{\boldsymbol{a}_t}\left[H\left(\boldsymbol{\theta}^t \mid z^{t-1}\right) - H\left(\boldsymbol{\theta}^t \mid z^{t-1}, \boldsymbol{a}_t\right)\right]$$

where entropy H is a measure of uncertainty:

$$H(\boldsymbol{w}) = -\sum_w \mathbb{P}(w)\log \mathbb{P}(w)$$

No discounting (for simplicity)
- $\sum_t u_t$ must be well-defined

$$\max \quad \mathbb{E}\left[\sum_{t=0}^{\infty}\Big(u_t(\boldsymbol{\theta}^t, \boldsymbol{a}^t) - \lambda\, I\big(\boldsymbol{\theta}^t, \boldsymbol{a}_t \mid z^{t-1}\big)\Big)\right]$$

No discounting (for simplicity)

- $\sum_t u_t$ must be well-defined

$$\max \quad \mathbb{E}\left[\sum_{t=0}^{\infty}\Big(u_t(\boldsymbol{\theta}^t, \boldsymbol{a}^t) - \lambda\, I(\boldsymbol{\theta}^t, \boldsymbol{a}_t \mid z^{t-1})\Big)\right]$$

By the properties of the entropic costs the dynamic problem is reduced to a collection of static problems

$$\max \quad \mathbb{E}\Big[\Big(u_t(\boldsymbol{\theta}^t, \boldsymbol{a}^t) + \mathbb{E}_{\boldsymbol{\theta}_{t+1}}\big[V(\boldsymbol{\theta}^{t+1}, \boldsymbol{z}^t) \mid \boldsymbol{\theta}^t, \boldsymbol{z}^t\big]\Big) - \lambda\, I(\boldsymbol{\theta}^t, \boldsymbol{x}_t \mid \boldsymbol{z}^{t-1}) \mid z^{t-1}\Big]$$

$$\max \quad \mathbb{E}\left[u_t\left(\boldsymbol{\theta}^t, \boldsymbol{a}^t\right) + \mathbb{E}_{\boldsymbol{\theta}_{t+1}}\left[V\left(\boldsymbol{\theta}^{t+1}, \boldsymbol{z}^t\right) \mid \boldsymbol{\theta}^t, \boldsymbol{z}^t\right] - \lambda\, I\left(\boldsymbol{\theta}^t, \boldsymbol{x}_t \mid \boldsymbol{z}^{t-1}\right) \mid z^{t-1}\right]$$

The logit-like solution:

$$p_t(a_t \mid \theta^t, z^{t-1}) \propto q_t(a_t \mid z^{t-1}) \exp\left[\frac{u_t(\theta^t, a^t) + \mathbb{E}_{\boldsymbol{\theta}_{t+1}}\left[V(\theta^{t+1}, z^t) \mid \theta^t, z^t\right]}{\lambda}\right]$$

$$\max \quad \mathbb{E}\Big[u_t(\boldsymbol{\theta}^t, \boldsymbol{a}^t) + \mathbb{E}_{\boldsymbol{\theta}_{t+1}}\big[V(\boldsymbol{\theta}^{t+1}, \boldsymbol{z}^t) \mid \boldsymbol{\theta}^t, \boldsymbol{z}^t\big] - \lambda\, I\big(\boldsymbol{\theta}^t, \boldsymbol{x}_t \mid \boldsymbol{z}^{t-1}\big) \mid z^{t-1}\Big]$$

The **default rule** $q_t(a_t \mid z^{t-1}) = \mathbb{E}_{\boldsymbol{\theta}^t}\big[p_t(a_t \mid \boldsymbol{\theta}^t, z^{t-1}) \mid z^{t-1}\big]$ is the unconditional choice probability consistent with p_t

The logit-like solution:

$$p_t(a_t \mid \theta^t, z^{t-1}) \propto q_t(a_t \mid z^{t-1}) \exp\left[\frac{u_t(\theta^t, a^t) + \mathbb{E}_{\boldsymbol{\theta}_{t+1}}\big[V(\theta^{t+1}, z^t) \mid \theta^t, z^t\big]}{\lambda}\right]$$

$$\max \quad \mathbb{E}\left[u_t(\boldsymbol{\theta}^t, \boldsymbol{a}^t) + \mathbb{E}_{\boldsymbol{\theta}_{t+1}}\left[V(\boldsymbol{\theta}^{t+1}, \boldsymbol{z}^t) \mid \boldsymbol{\theta}^t, \boldsymbol{z}^t\right] - \lambda\, I\left(\boldsymbol{\theta}^t, \boldsymbol{x}_t \mid \boldsymbol{z}^{t-1}\right) \mid z^{t-1}\right]$$

The **default rule** $q_t(a_t \mid z^{t-1}) = \mathbb{E}_{\boldsymbol{\theta}^t}\left[p_t(a_t \mid \boldsymbol{\theta}^t, z^{t-1}) \mid z^{t-1}\right]$ is the unconditional choice probability consistent with p_t

The logit-like solution:

$$p_t(a_t \mid \theta^t, z^{t-1}) \propto q_t(a_t \mid z^{t-1}) \exp\left[\frac{u_t(\theta^t, a^t) + \mathbb{E}_{\boldsymbol{\theta}_{t+1}}\left[V(\theta^{t+1}, z^t) \mid \theta^t, z^t\right]}{\lambda}\right]$$

The log-sum-exp value function:

$$V\left(\theta^t, z^{t-1}\right) = \lambda \log \sum_{a_t} q_t\left(a_t \mid z^{t-1}\right) \exp\left[\frac{u_t(\theta^t, a^t) + \mathbb{E}_{\boldsymbol{\theta}_{t+1}}\left[V(\theta^{t+1}, z^t) \mid \theta^t, z^t\right]}{\lambda}\right]$$

Entropic Cost

↓

Recursive Value

Entropic Cost

some assumption

Recursive Value

Choice-based recursivity:

$$\text{stage payoff} \;=\; \text{utility}\left(\theta^t, a^t\right) + \mathbb{E}\left[\text{value}\left(\theta^t, z^t\right)\right]$$

Steiner–Stewart–Matějka (2017)

Belief-based recursivity:

$$\text{stage payoff} \;=\; \text{utility}\left(\theta^t, a^t\right) + \mathbb{E}\left[\text{value}\left(\text{belief}\right)\right]$$

Miao–Xing (2020) study this DP, focusing on “Markov” environments

Example: Entropic Costs

Main Result

An **environment** E consists of:

- $\boldsymbol{\theta} = (\boldsymbol{\theta}_0, \boldsymbol{\theta}_1, \ldots)$ is a stochastic process: $\boldsymbol{\theta} \sim \pi$
  - π is a prior probability measure
- $\boldsymbol{a} = (\boldsymbol{a}_0, \boldsymbol{a}_1, \ldots)$ is a stochastic process: $\boldsymbol{a} \sim q$
  - q is a default rule
- $u = (u_0, u_1, \ldots)$ is payoff functions
- No discounting (for simplicity)
  - $\sum_t u_t$ must be well-defined

This default rule q "should" be endogenous but in today's talk, we will take it as given.

How do we define the value of an environment E?

1. Transform E to a default random utility:

A "plain" definition is…

$$\text{value}(E) = \max \; \mathbb{E}\left[\sum_{t=0}^{\infty}\left(u_t(\boldsymbol{\theta}^t, \boldsymbol{a}^t) - \text{cost}(\boldsymbol{\theta}^t, \boldsymbol{x}_t \mid z^{t-1})\right)\right]$$

But this definition is hard to work with…

- T… …ed by (π, q)

2. De… …$\mathbb{R}$

- $\mathbb{V}[U]$ is an abstract formulation of the value of E

How do we define the value of an environment E?

1. Transform E to a **default random utility**:

$$\mathbf{U} = \sum_{t=0}^{\infty} u_t(\boldsymbol{\theta}^t, \boldsymbol{a}^t)$$

   - The random vector $(\boldsymbol{\theta}, \boldsymbol{a})$ follow the joint distributions induced by (π, q)

2. Define a value functional $\mathbb{V} : \{\text{default random utilities } \mathbf{U}\} \to \mathbb{R}$

   - $\mathbb{V}[\mathbf{U}]$ is an abstract formulation of the value of E

Control-problem interpretation: how to control a process $(u_t(\boldsymbol{\theta}^t, \boldsymbol{a}^t))_t$

- DM controls actions against random states
- control costs depend on the deviation from the default rule q

Static rational inattention with entropic costs

$$\max \quad \mathbb{E}\left[u(\boldsymbol{\theta},\boldsymbol{a}) - \lambda I(\boldsymbol{\theta},\boldsymbol{x})\right]$$

Value at state θ

$$\lambda \log \sum_a q(a) \exp\left\{\frac{u(\theta,a)}{\lambda}\right\}$$

Static rational inattention with entropic costs

$$\max \quad \mathbb{E}\left[u(\boldsymbol{\theta},\boldsymbol{a}) - \lambda I(\boldsymbol{\theta},\boldsymbol{x})\right]$$

Value at state θ

$$\lambda \log \sum_a q(a)\exp\left\{\frac{u(\theta,a)}{\lambda}\right\}$$

Risk-loving CARA utility

$$u^\gamma(z) = \exp\{\gamma z\}$$

Certainty-equivalent of exogenous $\boldsymbol{w}$

$$\frac{1}{\gamma}\log \mathbb{E}\left[\exp\{\gamma \boldsymbol{w}\}\right]$$

Static rational inattention with entropic costs

$$\max \quad \mathbb{E}\left[u(\boldsymbol{\theta}, \boldsymbol{a}) - \lambda I(\boldsymbol{\theta}, \boldsymbol{x})\right]$$

Risk-loving CARA utility

$$u^{\gamma}(z) = \exp\{\gamma z\}$$

Value at state θ

Certainty-equivalent of exogenous $\boldsymbol{w}$

$$\lambda \log \sum_a q(a) \exp\left\{\frac{u(\theta, a)}{\lambda}\right\} \overset{\gamma = \frac{1}{\lambda}}{\underset{\boldsymbol{w} \sim u(\theta, \cdot)}{=\!=\!=}} \frac{1}{\gamma} \log \mathbb{E}\left[\exp\{\gamma \boldsymbol{w}\}\right]$$

Static rational inattention with entropic costs

$$\max \quad \mathbb{E}\left[u(\boldsymbol{\theta},\boldsymbol{a}) - \lambda I(\boldsymbol{\theta},\boldsymbol{x})\right]$$

Risk-loving CARA utility

$$u^{\gamma}(z) = \exp\{\gamma z\}$$

Value at state θ

Certainty-equivalent of exogenous $\boldsymbol{w}$

$$\lambda \log \sum_{a} q(a)\exp\left\{\frac{u(\theta,a)}{\lambda}\right\} \;\underset{\boldsymbol{w}\,\sim\,u(\theta,\,\cdot\,)}{\overset{\gamma = \frac{1}{\lambda}}{=\!=\!=}}\; \frac{1}{\gamma}\log \mathbb{E}\left[\exp\{\gamma \boldsymbol{w}\}\right]$$

We use finance techniques to study RI models (with non-entropic costs)

Rational inattention

- exogenous $\boldsymbol{\theta}$
- endogenous $\boldsymbol{a}$

Finance

- exogenous $\boldsymbol{w}$
- **risk-aversion** in the literature

A value functional $\mathbb{V}$ is **recursive** if:

- for an environment E at time 0

$$\mathbb{V}\big[\mathbf{U}\big] \quad = \quad \mathbb{V}\Big[\, u_0(\boldsymbol{\theta}^0, \boldsymbol{a}^0) \; + \; \mathbb{V}\Big[\mathbf{U} \mid \boldsymbol{\theta}^0, \boldsymbol{z}^0\Big]\, \Big]$$

A value functional $\mathbb{V}$ is **recursive** if:

- for an environment E at time 0

dynamic **static**

$$\mathbb{V}[\mathbf{U}] = \mathbb{V}\Big[u_0(\boldsymbol{\theta}^0, \boldsymbol{a}^0) + \mathbb{V}\big[\mathbf{U} \mid \boldsymbol{\theta}^0, \boldsymbol{z}^0\big]\Big]$$

present utility future

A value functional $\mathbb{V}$ is **recursive** if:

- for an environment E at time 0

dynamic / **static**

$$\mathbb{V}[\mathbf{U}] = \mathbb{V}\left[u_0(\boldsymbol{\theta}^0, \boldsymbol{a}^0) + \mathbb{V}\left[\mathbf{U} \mid \boldsymbol{\theta}^0, \boldsymbol{z}^0\right]\right]$$

present utility / future

Recursivity at time t:

$$\mathbb{V}\left[\mathbf{U} \mid z^{t-1}\right] = \mathbb{V}\left[u_t(\boldsymbol{\theta}^t, \boldsymbol{a}^t) + \mathbb{V}\left[\mathbf{U} \mid \boldsymbol{\theta}^t, \boldsymbol{z}^t\right] \mid z^{t-1}\right]$$

How can we compare two environments E and $\tilde{E}$?

- Compare the default random utilities $\mathbf{U}$ and $\tilde{\mathbf{U}}$

The value functional $\mathbb{V}$ is **monotone** if

- If $\mathbf{U} \gtrsim_{\text{FOSD}} \tilde{\mathbf{U}}$ regardless of a state, then "value of $\mathbf{U} \ \geq$ value of $\tilde{\mathbf{U}}$"
- If $\tilde{\mathbf{U}} \gtrsim_{\text{FOSD}} \mathbf{U}$ regardless of a state, then "value of $\tilde{\mathbf{U}} \ \geq$ value of $\mathbf{U}$"

How can we compare two environments E and $\tilde{E}$?

- Compare the default random utilities $\mathbf{U}$ and $\tilde{\mathbf{U}}$

The value functional $\mathbb{V}$ is **monotone** if

- If $\mathbf{U} \gtrsim_{\text{FOSD}} \tilde{\mathbf{U}}$ regardless of a state, then "value of $\mathbf{U} \geq$ value of $\tilde{\mathbf{U}}$"
- If $\tilde{\mathbf{U}} \gtrsim_{\text{FOSD}} \mathbf{U}$ regardless of a state, then "value of $\tilde{\mathbf{U}} \geq$ value of $\mathbf{U}$"

Required only for a **simple** environment $\tilde{E}$ such that

$$\tilde{\mathbf{U}} = \tilde{u}_0(\tilde{\boldsymbol{\theta}}_0) + \sum_{t=1}^{\tau} \left(\epsilon \tilde{\boldsymbol{\theta}}_t \tilde{\boldsymbol{a}}_t - \eta\right)$$

random walk

Time-0:
- payoff $\tilde{u}_0(\tilde{\theta}_0)$ is independent of action $\tilde{a}_0$

Time-$t \geq 1$:
- binary states: $\tilde{\theta}_t = \pm 1$
- binary actions: $\tilde{a}_t = \pm 1$
- stage payoffs: $\epsilon \tilde{\theta}_t \tilde{a}_t - \eta$ for some η
- it ends at a bounded stopping time τ

$\mathbb{V}$ is monotone and recursive

$$\Downarrow$$

$\mathbb{V}$ has the log-sum-exp formula

$$\Downarrow$$

$\mathbb{V}$ has the entropic costs

$\mathbb{V}$ is monotone and recursive

$\Downarrow$

$\mathbb{V}$ has the log-sum-exp formula

Today's focus

$\Downarrow$

$\mathbb{V}$ has the entropic costs

If the value functional $\mathbb{V}$ is monotone and recursive then there is some $\lambda > 0$ such that

$$\mathbb{V}[\mathbf{U}] = \mathbb{E}_{\boldsymbol{\theta}^0}\left[\lambda \log \sum_{a_0} q_0(a_0) \exp\left[\frac{u_0(\boldsymbol{\theta}^0, a^0) + \mathbb{V}[\mathbf{U} \mid \boldsymbol{\theta}^0, z^0]}{\lambda}\right]\right]$$

If the value functional $\mathbb{V}$ is monotone and recursive then there is some $\lambda > 0$ such that

$$\mathbb{V}[\mathbf{U}] = \mathbb{E}_{\boldsymbol{\theta}^0}\left[\lambda \log \sum_{a_0} q_0(a_0) \exp\left[\frac{u_0(\boldsymbol{\theta}^0, a^0) + \mathbb{V}[\mathbf{U} \mid \boldsymbol{\theta}^0, z^0]}{\lambda}\right]\right]$$

Time-t result is analogous

How to evaluate the value $\mathbb{V}[\mathbf{U}]$ of the default random utility $\mathbf{U}$ induced from an env. E?

- Approximate $\mathbf{U}$ with a random walk $\tilde{\mathbf{U}}$ (for each possible state)

How to evaluate the value $\mathbb{V}[\mathbf{U}]$ of the default random utility $\mathbf{U}$ induced from an env. E?

- Approximate $\mathbf{U}$ with a random walk $\tilde{\mathbf{U}}$ (for each possible state)

$$\mathbf{U}\,|_{\theta_0} \precsim_{\text{FOSD}} \tilde{\mathbf{U}}\,|_{\tilde{\theta}_0}$$

Proof Idea

How to evaluate the value $\mathbb{V}[\mathbf{U}]$ of the default random utility $\mathbf{U}$ induced from an env. E?

- Approximate $\mathbf{U}$ with a random walk $\tilde{\mathbf{U}}$ (for each possible state)

$$\mathbf{U}\,|_{\theta_0} \precsim_{\text{FOSD}} \tilde{\mathbf{U}}\,|_{\tilde{\theta}_0}$$

monotonicity ↓

$$\mathbb{V}[\mathbf{U}] \quad \leq \quad \mathbb{V}[\tilde{\mathbf{U}}]$$

How to evaluate the value $\mathbb{V}[\mathbf{U}]$ of the default random utility $\mathbf{U}$ induced from an env. E?

- Approximate $\mathbf{U}$ with a random walk $\tilde{\mathbf{U}}$ (for each possible state)

$$\mathbf{U}|_{\theta_0} \precsim_{\text{FOSD}} \tilde{\mathbf{U}}|_{\tilde{\theta}_0}$$

monotonicity ↓

$= 0$ by **recursivity** for some η

$$\mathbb{V}[\mathbf{U}] \leq \mathbb{V}[\tilde{\mathbf{U}}] = \mathbb{E}\left[\tilde{u}_0(\tilde{\boldsymbol{\theta}}_0)\right] + \mathbb{V}\left[\sum_{t=1}^{\tau}\left(\epsilon\tilde{\boldsymbol{\theta}}_t\tilde{\boldsymbol{a}}_t - \eta\right)\right]$$

$$\tilde{\mathbf{U}} = \tilde{u}_0(\tilde{\boldsymbol{\theta}}_0) + \sum_{t=1}^{\tau}\left(\epsilon\tilde{\boldsymbol{\theta}}_t\tilde{\boldsymbol{a}}_t - \eta\right)$$

How to evaluate the value $\mathbb{V}[\mathbf{U}]$ of the default random utility $\mathbf{U}$ induced from an env. E?

- Approximate $\mathbf{U}$ with a random walk $\tilde{\mathbf{U}}$ (for each possible state)

$$\mathbf{U}\,|_{\theta_0} \precsim_{\text{FOSD}} \tilde{\mathbf{U}}\,|_{\tilde{\theta}_0}$$

monotonicity ↓

$= 0$ by **recursivity** for some η

$$\mathbb{V}[\mathbf{U}] \quad \leq \quad \mathbb{V}[\tilde{\mathbf{U}}] = \mathbb{E}\left[\tilde{u}_0(\tilde{\boldsymbol{\theta}}_0)\right] + \mathbb{V}\left[\sum_{t=1}^{\tau}\left(\epsilon\tilde{\boldsymbol{\theta}}_t\tilde{\boldsymbol{a}}_t - \eta\right)\right]$$

$$\tilde{\mathbf{U}} = \tilde{u}_0(\tilde{\boldsymbol{\theta}}_0) + \sum_{t=1}^{\tau}\left(\epsilon\tilde{\boldsymbol{\theta}}_t\tilde{\boldsymbol{a}}_t - \eta\right)$$

$$\approx \mathbb{E}\left[\frac{1}{\gamma}\log\mathbb{E}\left[e^{\gamma\mathbf{U}} \mid \boldsymbol{\theta}_0\right]\right]$$

by the **connection to CARA**

Proof Idea

How to evaluate the value $\mathbb{V}[\mathbf{U}]$ of the default random utility $\mathbf{U}$ induced from an env. E?

- Approximate $\mathbf{U}$ with a random walk $\tilde{\mathbf{U}}$ (for each possible state)

$$\tilde{\mathbf{U}}'|_{\tilde{\theta}'_0} \precsim_{\text{FOSD}} \mathbf{U}|_{\theta_0} \precsim_{\text{FOSD}} \tilde{\mathbf{U}}|_{\tilde{\theta}_0}$$

monotonicity ↓

$$\mathbb{V}[\tilde{\mathbf{U}}'] \leq \mathbb{V}[\mathbf{U}] \leq \mathbb{V}[\tilde{\mathbf{U}}] = \mathbb{E}\left[\tilde{u}_0(\tilde{\boldsymbol{\theta}}_0)\right] + \mathbb{V}\left[\sum_{t=1}^{\tau}\left(\epsilon\tilde{\boldsymbol{\theta}}_t\tilde{\boldsymbol{a}}_t - \eta\right)\right]$$

$= 0$ by **recursivity** for some η

$$\tilde{\mathbf{U}} = \tilde{u}_0(\tilde{\boldsymbol{\theta}}_0) + \sum_{t=1}^{\tau}\left(\epsilon\tilde{\boldsymbol{\theta}}_t\tilde{\boldsymbol{a}}_t - \eta\right)$$

$$\approx \mathbb{E}\left[\frac{1}{\gamma}\log\mathbb{E}\left[e^{\gamma\mathbf{U}} \mid \boldsymbol{\theta}_0\right]\right]$$

by the **connection to CARA**

How to evaluate the value $\mathbb{V}[\mathbf{U}]$ of the default random utility $\mathbf{U}$ induced from an env. E?

- Approximate $\mathbf{U}$ with a random walk $\tilde{\mathbf{U}}$ (for each possible state)

$$\tilde{\mathbf{U}}'|_{\tilde{\theta}_0'} \precsim_{\text{FOSD}} \mathbf{U}|_{\theta_0} \precsim_{\text{FOSD}} \tilde{\mathbf{U}}|_{\tilde{\theta}_0}$$

← Kupper-Schachermayer-Skorokhod Embedding Theorem (modified version)

↓ monotonicity

$$\mathbb{V}[\tilde{\mathbf{U}}'] \le \mathbb{V}[\mathbf{U}] \le \mathbb{V}[\tilde{\mathbf{U}}] = \mathbb{E}[\tilde{u}_0(\tilde{\boldsymbol{\theta}}_0)] + \mathbb{V}\left[\sum_{t=1}^{\tau}(\epsilon\tilde{\boldsymbol{\theta}}_t\tilde{\boldsymbol{a}}_t - \eta)\right]$$

$$\tilde{\mathbf{U}} = \tilde{u}_0(\tilde{\boldsymbol{\theta}}_0) + \sum_{t=1}^{\tau}(\epsilon\tilde{\boldsymbol{\theta}}_t\tilde{\boldsymbol{a}}_t - \eta)$$

$= 0$ by **recursivity** for some η

$$\approx \mathbb{E}\left[\frac{1}{\gamma}\log \mathbb{E}\left[e^{\gamma \mathbf{U}} \mid \boldsymbol{\theta}_0\right]\right]$$

by the **connection to CARA**

Entropic Cost

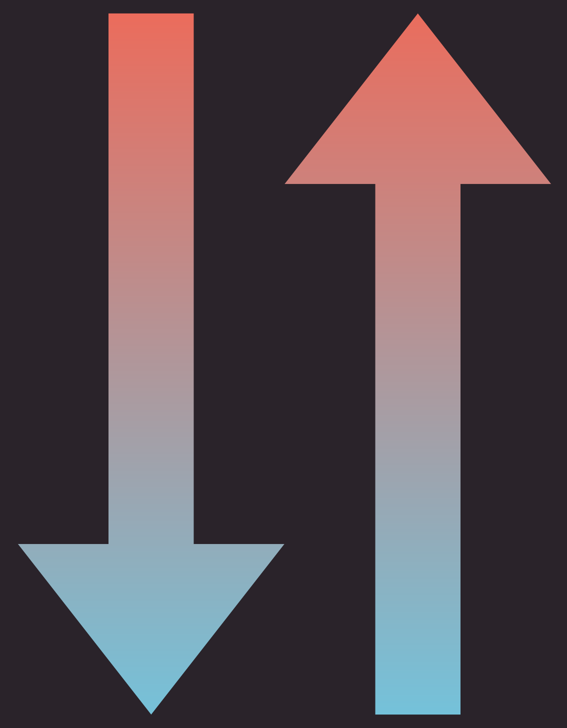

Monotonicity

Recursive Value

Approximate the default random utility $\mathbf{U}$ of a static environment E by a random walk:

$$\tilde{\mathbf{U}} = \tilde{u}_0(\tilde{\boldsymbol{\theta}}_0) + \sum_{t=1}^{\tau} \left(\epsilon \tilde{\boldsymbol{\theta}}_t \tilde{\boldsymbol{a}}_t - \eta\right)$$

How can we find such a "nice" $\tilde{\mathbf{U}}$?

- For any small ϵ, we have three parameters of choice
  - the drift parameter η
  - the stopping time τ
  - the initial value $\tilde{u}_0(\tilde{\theta}_0)$ for each $\tilde{\theta}_0$

Drift Parameter and Martingale-ish Property

$$\tilde{\mathrm{U}} = \tilde{u}_0(\tilde{\boldsymbol{\theta}}_0) + \sum_{t=1}^{\tau} \left(\epsilon \tilde{\boldsymbol{\theta}}_t \tilde{\boldsymbol{a}}_t - \eta\right)$$

How to find the drift parameter η

- "value at time $t \geq 1$" = 0
- fee η = "value of the game to match action a_t with state θ_t"

$$\eta = \mathbb{V}\left[\epsilon \, \tilde{\boldsymbol{\theta}}_1 \tilde{\boldsymbol{a}}_1\right]$$

The **martingale-ish** property:

$$\mathbb{V}[\tilde{\mathrm{U}}] \quad = \quad \mathbb{E}\left[\tilde{u}_0(\tilde{\boldsymbol{\theta}}_0)\right] \quad + \quad \mathbb{V}\left[\sum_{t=1}^{\tau} \left(\epsilon \tilde{\boldsymbol{\theta}}_t \tilde{\boldsymbol{a}}_t - \eta\right)\right] \quad = \quad \mathbb{E}\left[\tilde{u}_0(\tilde{\boldsymbol{\theta}}_0)\right]$$

$$\tilde{\mathrm{U}} = \tilde{u}_0\left(\tilde{\boldsymbol{\theta}}_0\right) + \sum_{t=1}^{\tau} \left(\epsilon \tilde{\boldsymbol{\theta}}_t \tilde{\boldsymbol{a}}_t - \eta\right)$$

How to find the drift parameter η

- "value at time $t \geq 1$" $= 0$
- fee η $=$ "value of the game to match action a_t with state θ_t"

$$\eta = \mathbb{V}\left[\epsilon\, \tilde{\boldsymbol{\theta}}_1 \tilde{\boldsymbol{a}}_1\right]$$

The **martingale-ish** property:

$$\mathbb{V}\left[\tilde{\mathrm{U}}\right] = \mathbb{E}\left[\tilde{u}_0\left(\tilde{\boldsymbol{\theta}}_0\right)\right] + \mathbb{V}\left[\sum_{t=1}^{\tau} \left(\epsilon \tilde{\boldsymbol{\theta}}_t \tilde{\boldsymbol{a}}_t - \eta\right)\right] = \mathbb{E}\left[\tilde{u}_0\left(\tilde{\boldsymbol{\theta}}_0\right)\right]$$

exogenous

$= 0$ by **recursivity**

Connection to CARA Utilities

Drift parameter

$$\mathbb{V}\left[\epsilon\,\tilde{\boldsymbol{\theta}}_1\tilde{\boldsymbol{a}}_1\right] = \eta$$

let γ equate them

Risk-loving CARA utility

$$\frac{1}{\gamma}\log\mathbb{E}\left[e^{\gamma\epsilon\mathbf{B}}\right]$$

$\mathbf{B} = \tilde{\boldsymbol{\theta}}_1\tilde{\boldsymbol{a}}_1$ is a Bernoulli with values ± 1

Drift parameter

$$\mathbb{V}\left[\epsilon\,\tilde{\boldsymbol{\theta}}_1\tilde{\boldsymbol{a}}_1\right] = \eta$$

let γ equate them

Risk-loving CARA utility

$$\frac{1}{\gamma}\log\mathbb{E}\left[e^{\gamma\epsilon\mathbf{B}}\right]$$

$\mathbf{B} = \tilde{\boldsymbol{\theta}}_1\tilde{\boldsymbol{a}}_1$ is a Bernoulli with values ± 1

CARA utility: $u^\gamma(z) = e^{\gamma z}$

martingale: $\mathbb{E}\left[u^\gamma\left(\tilde{\mathbf{U}}\right) \mid \tilde{\theta}_0\right] = u^\gamma\left(\tilde{u}\left(\tilde{\theta}_0\right)\right)$

immediate from the definitions of parameters η and γ

Modified Kupper–Schachermayer–Skorokhod Embedding Theorem

For a default random utility $\mathbf{U}$, we have a simple environment $\tilde{E}$ with default random utility

$$\tilde{\mathrm{U}} = \tilde{u}_0\big(\tilde{\boldsymbol{\theta}}_0\big) + \sum_{t=1}^{\tau} \big(\epsilon \tilde{\boldsymbol{\theta}}_t \tilde{\boldsymbol{a}}_t - \eta\big)$$

such that for each θ_0

- $\mathbf{U}|_{\theta_0} \approx \tilde{\mathrm{U}}|_{\tilde{\theta}_0}$
- $\tilde{u}\big(\tilde{\theta}_0\big) \approx \big(u^\gamma\big)^{-1}\big(\mathbb{E}\big[u^\gamma(\mathbf{U}) \mid \theta_0\big]\big)$

Modified Kupper–Schachermayer–Skorokhod Embedding Theorem

$\mathbf{U}$ takes either value $\underline{u} < \overline{u}$

- the initial value

$$u^\gamma\big(\tilde{u}(\tilde{\theta}_0)\big) \approx \mathbb{E}\big[u^\gamma(\mathbf{U}) \mid \theta_0\big]$$

- the martingale property

$$\mathbb{E}\big[u^\gamma\big(\tilde{\mathbf{U}}\big) \mid \tilde{\theta}_0\big] = u^\gamma\big(\tilde{u}_0(\tilde{\theta}_0)\big)$$

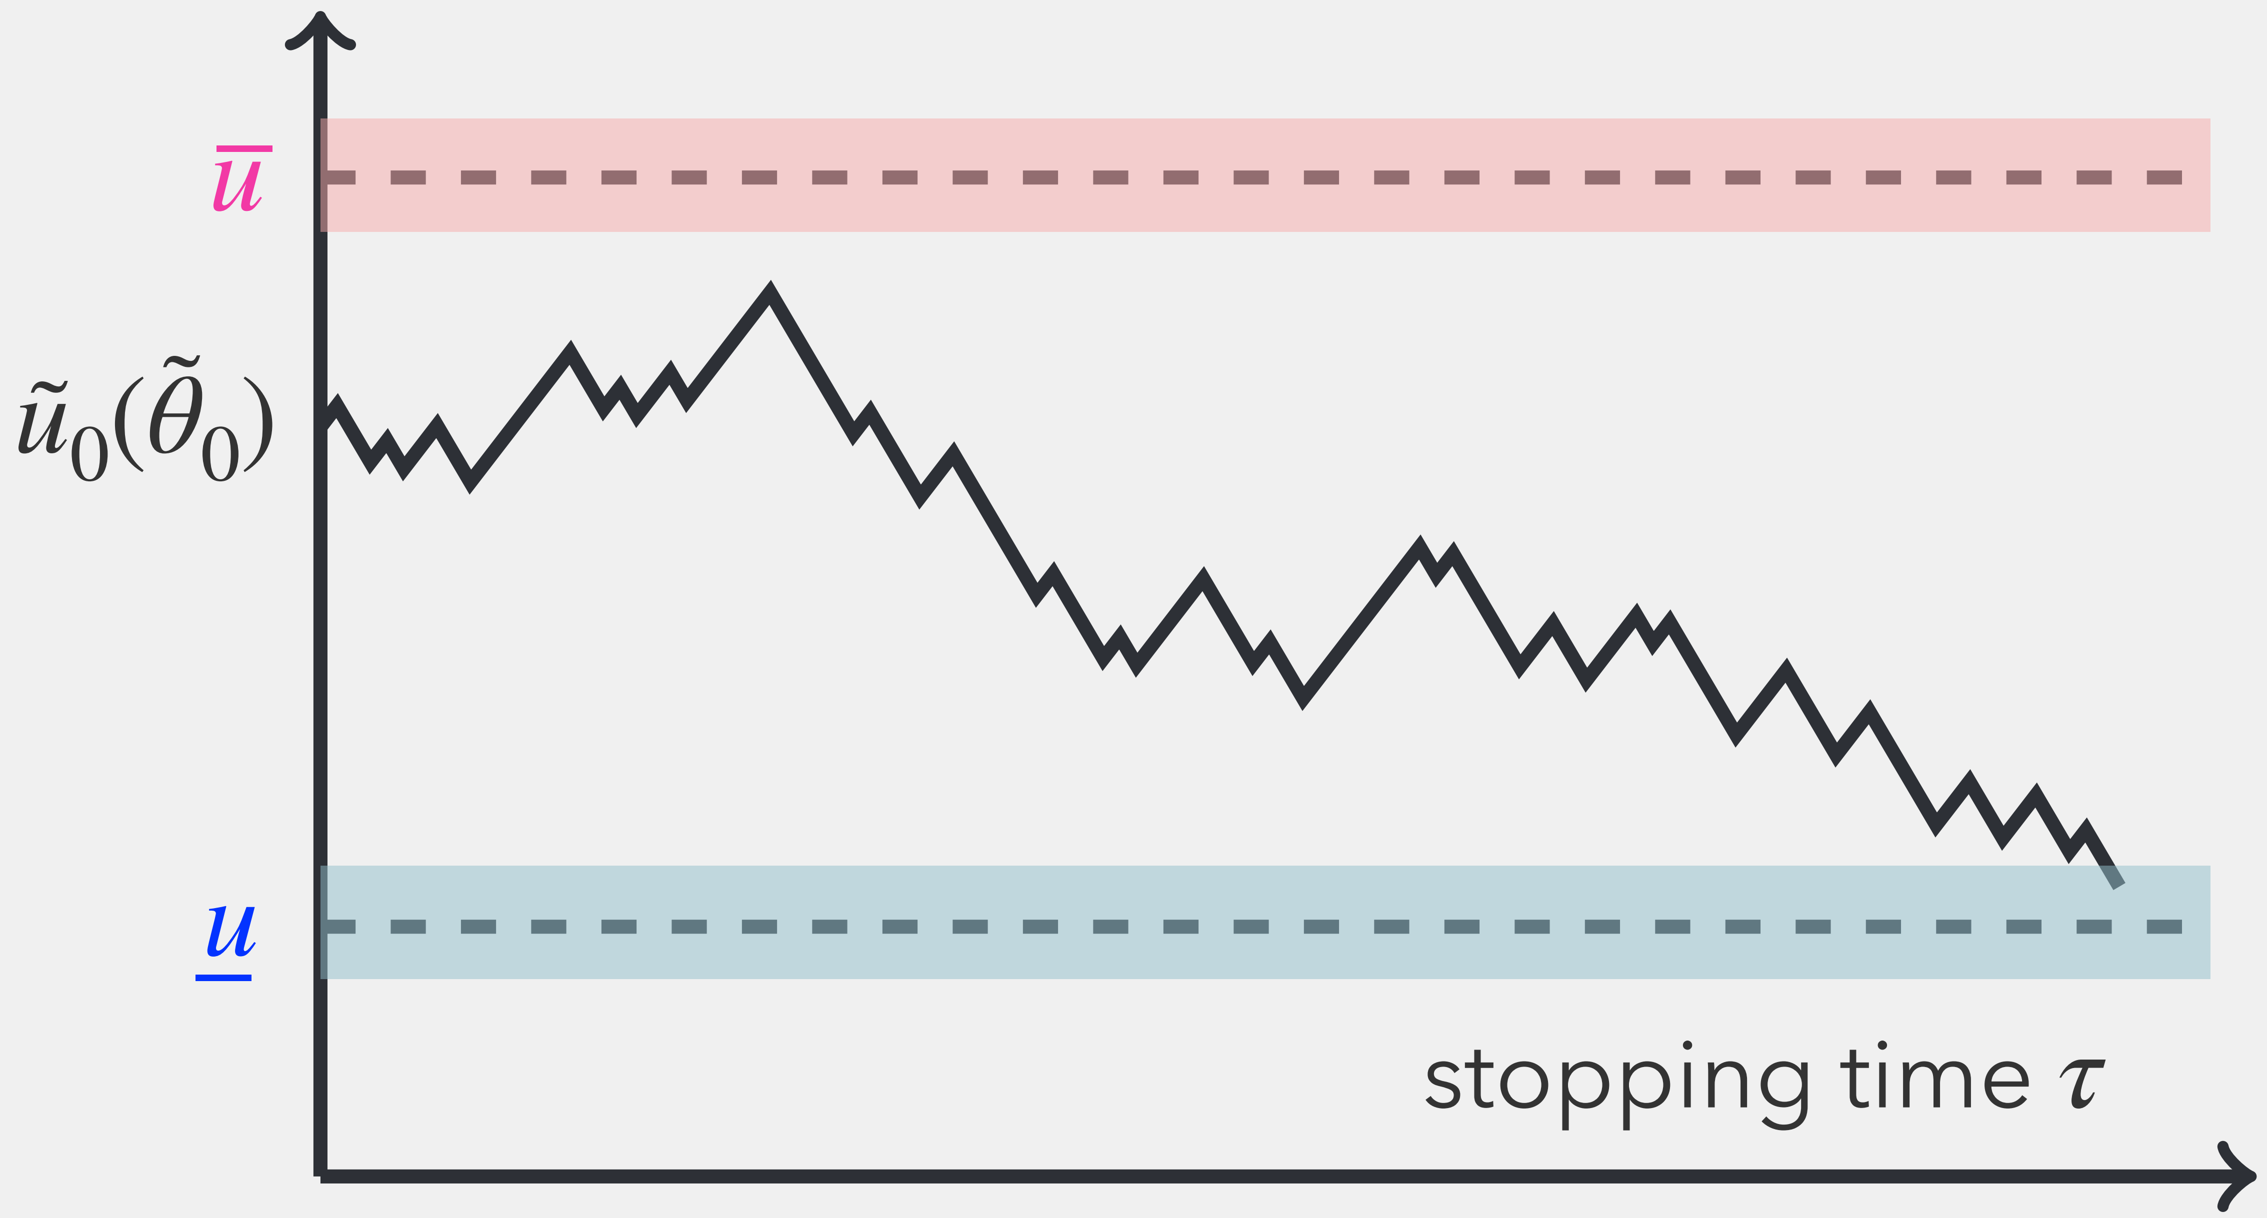

Modified Kupper–Schachermayer–Skorokhod Embedding Theorem

$\mathbf{U}$ takes either value $\underline{u} < \overline{u}$

- the initial value

$$u^\gamma\big(\tilde{u}(\tilde{\theta}_0)\big) \approx \mathbb{E}\big[u^\gamma(\mathbf{U}) \mid \theta_0\big]$$

- the martingale property

$$\mathbb{E}\big[u^\gamma(\tilde{\mathbf{U}}) \mid \tilde{\theta}_0\big] = u^\gamma\big(\tilde{u}_0(\tilde{\theta}_0)\big)$$

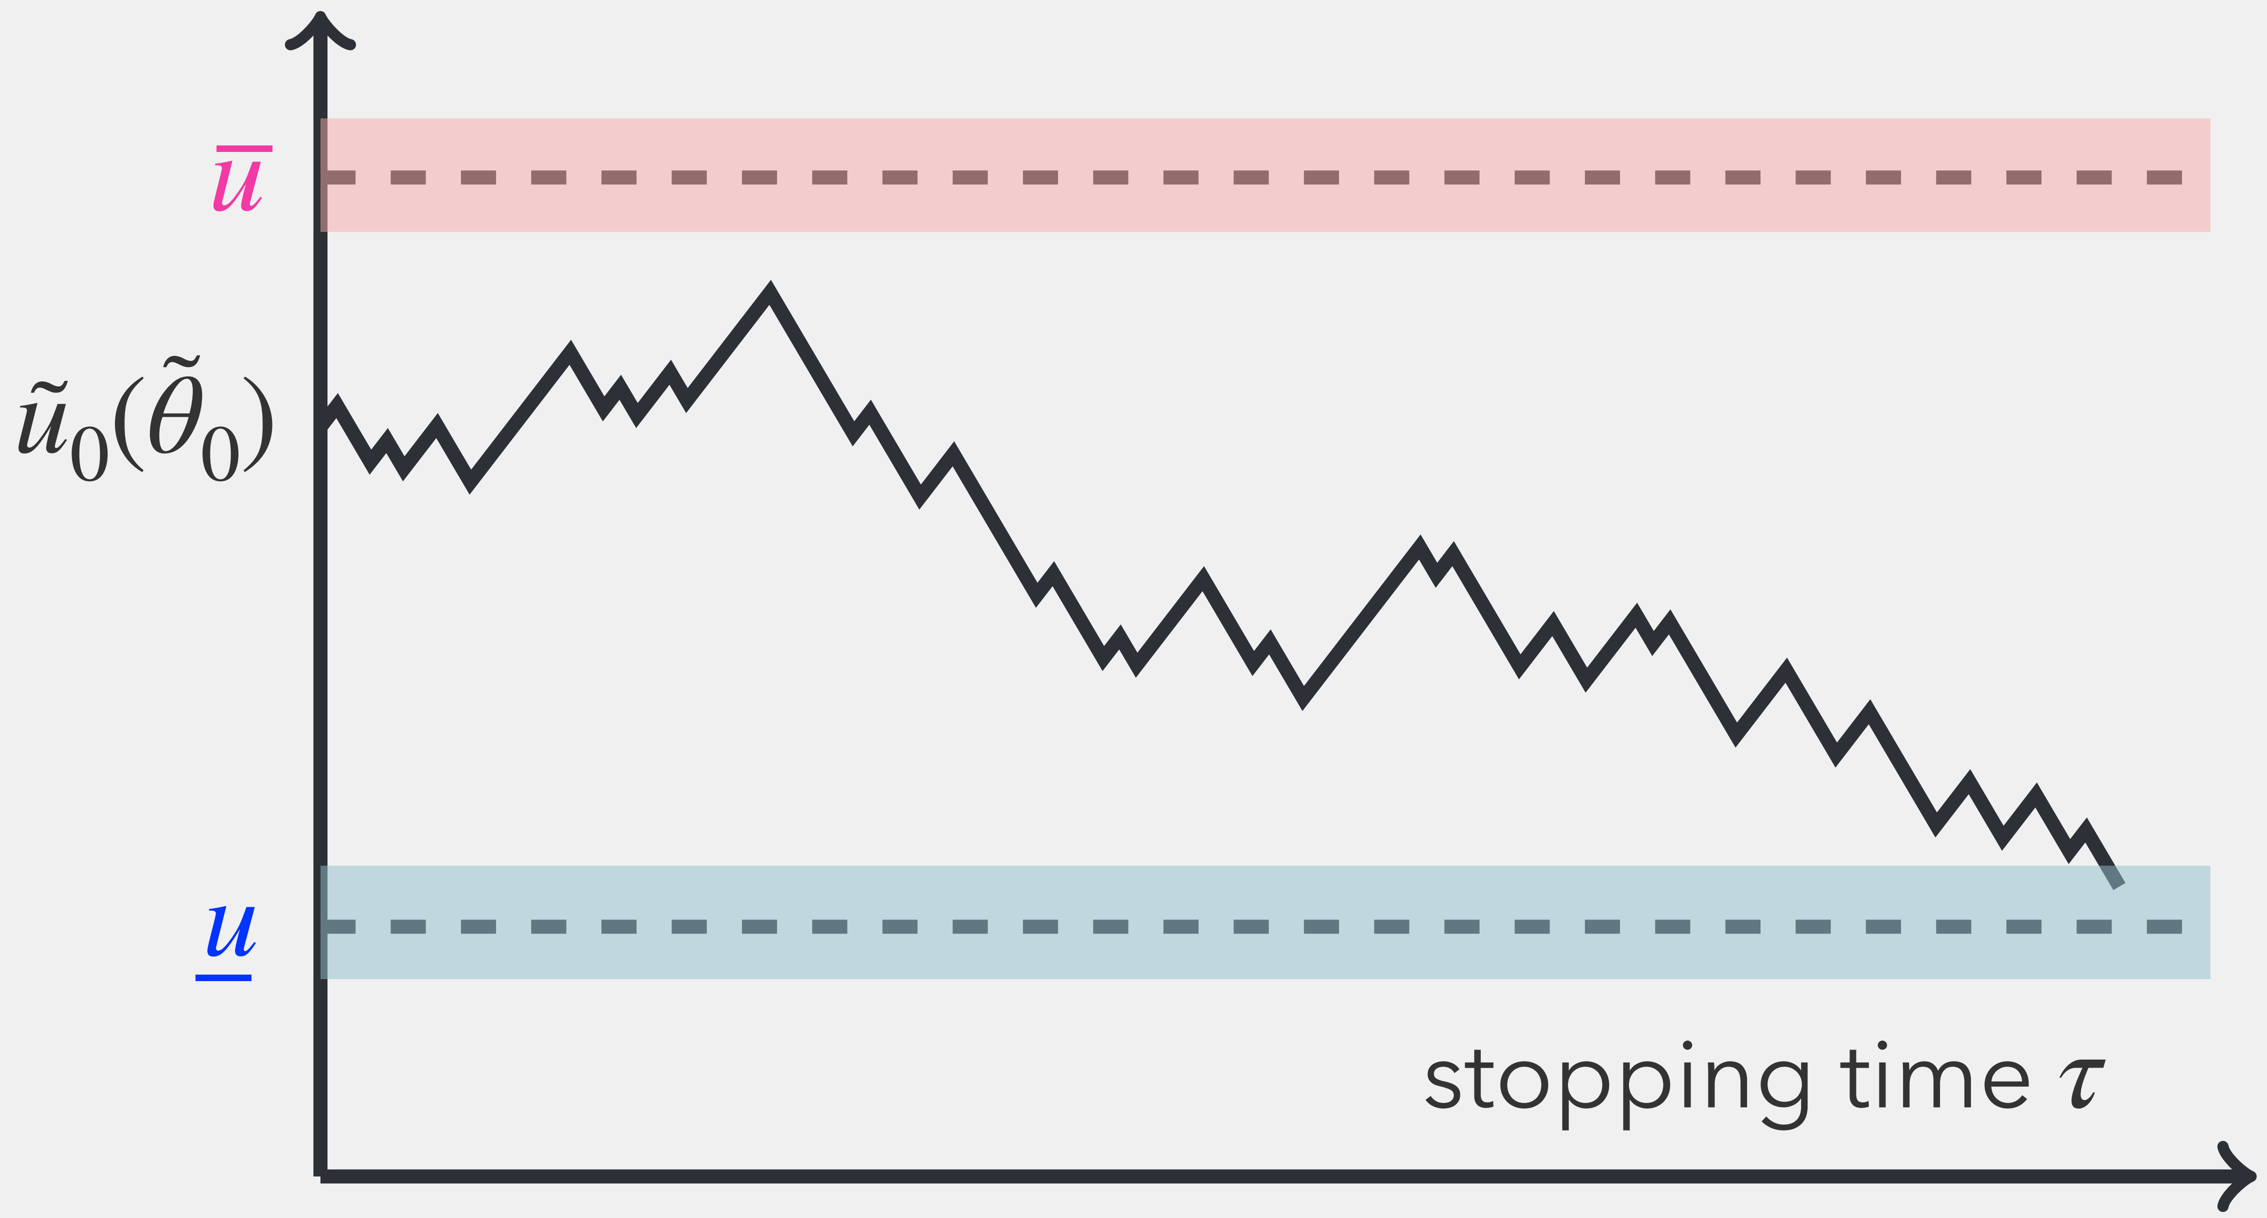

$$\mathbb{E}\big[u^\gamma(\mathbf{U}) \mid \theta_0\big]$$

$$\approx \mathbb{E}\big[u^\gamma(\tilde{\mathbf{U}}) \mid \tilde{\theta}_0\big]$$

Modified Kupper–Schachermayer–Skorokhod Embedding Theorem

$\mathbf{U}$ takes either value $\underline{u} < \overline{u}$

- the initial value

$$u^\gamma\big(\tilde{u}(\tilde{\theta}_0)\big) \approx \mathbb{E}\big[u^\gamma(\mathbf{U}) \mid \theta_0\big]$$

- the martingale property

$$\mathbb{E}\big[u^\gamma\big(\tilde{\mathbf{U}}\big) \mid \tilde{\theta}_0\big] = u^\gamma\big(\tilde{u}_0(\tilde{\theta}_0)\big)$$

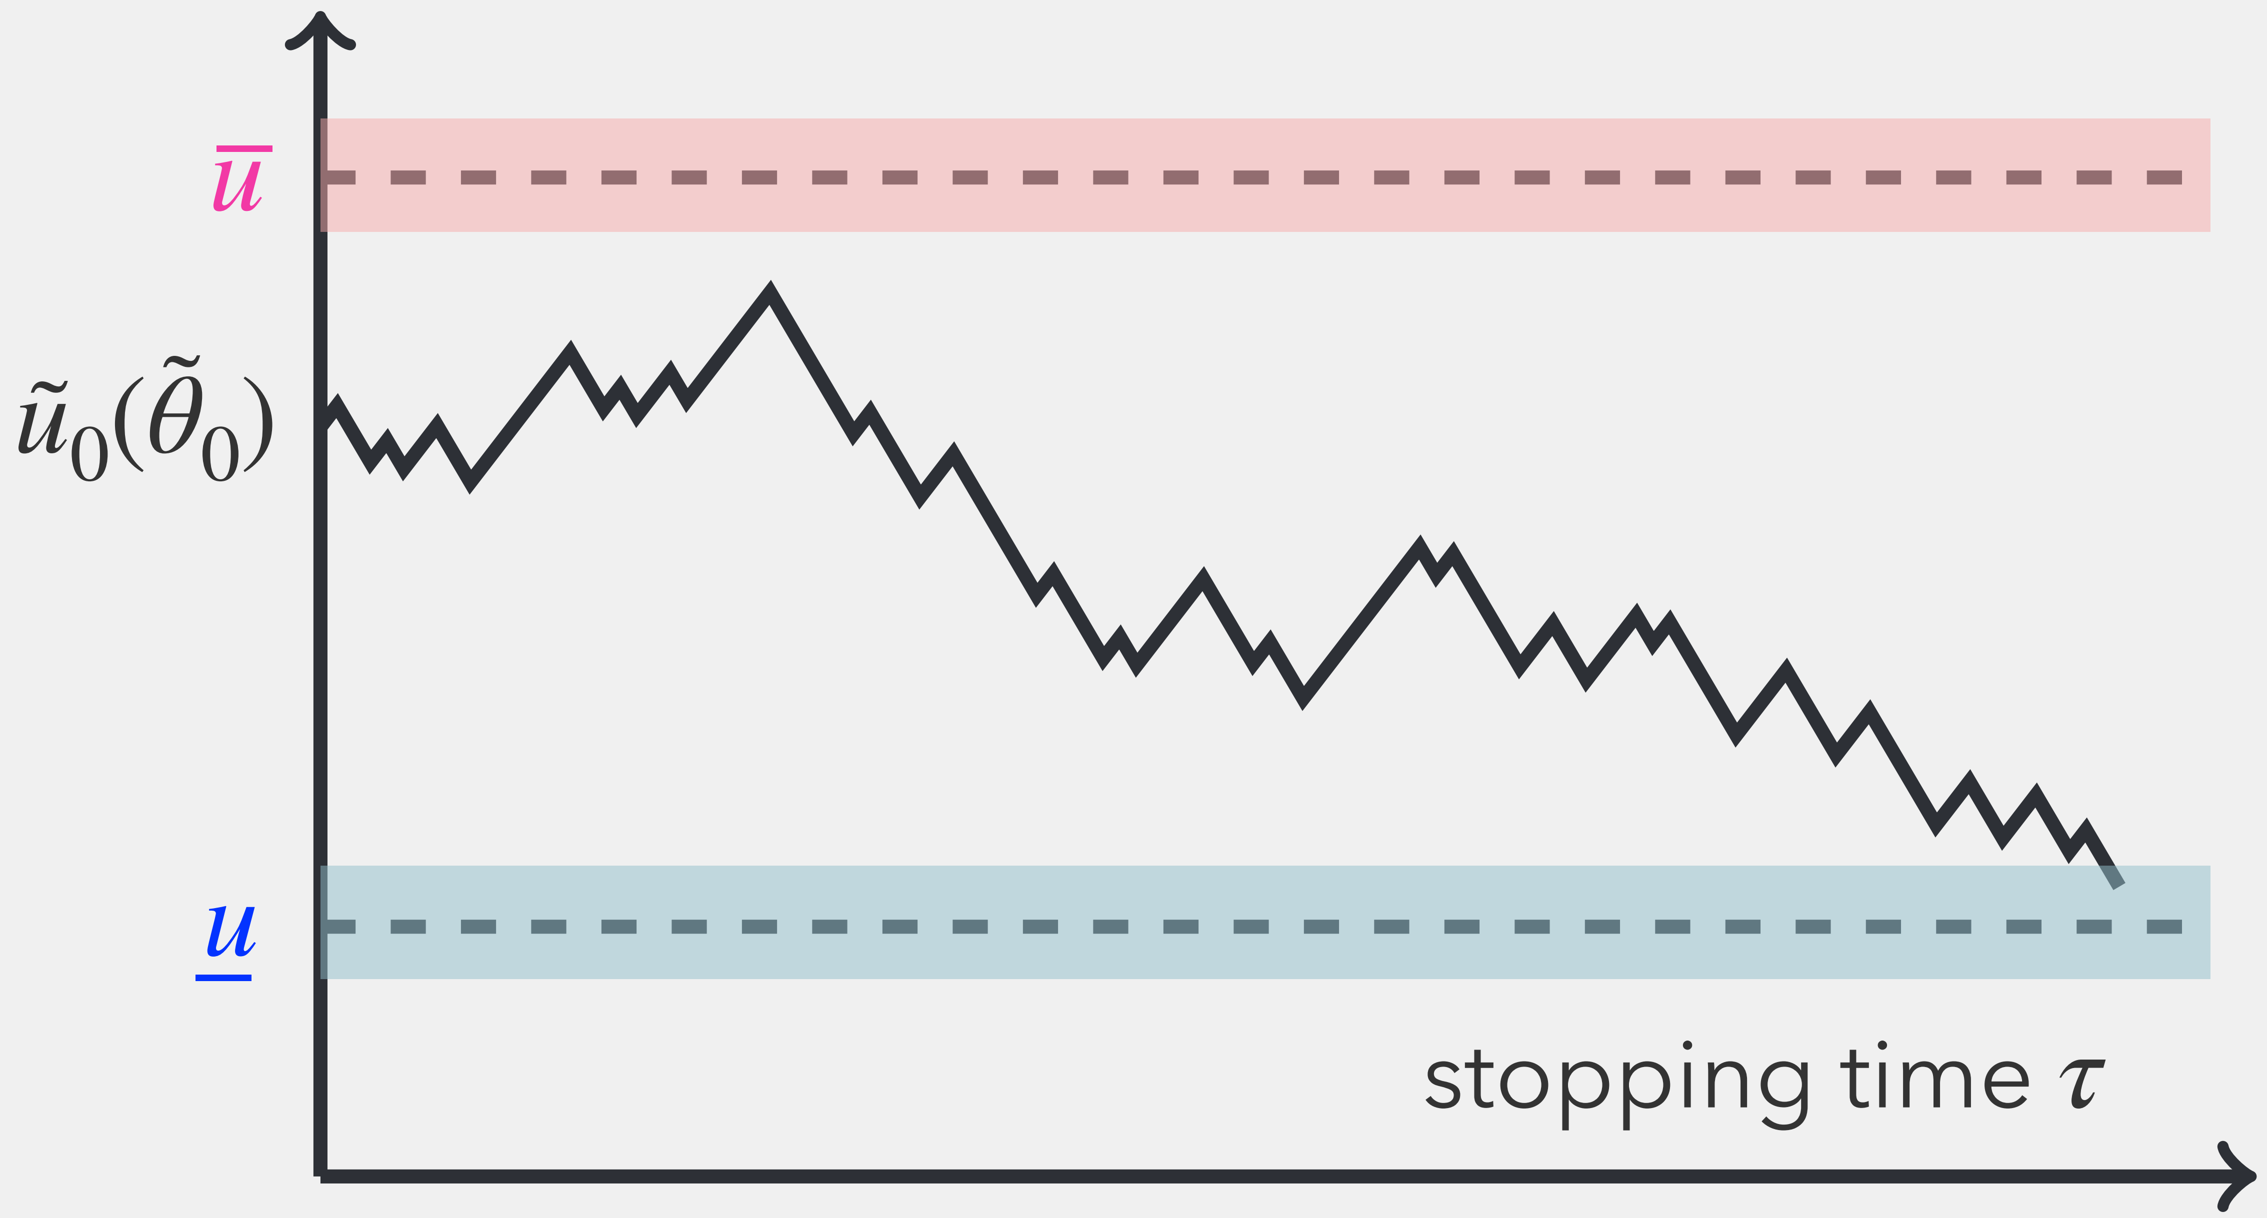

$$\mathbb{E}\big[u^\gamma(\mathbf{U}) \mid \theta_0\big] = \mathbb{P}\big[\mathbf{U} = \overline{u} \mid \theta_0\big]\, u^\gamma(\overline{u}) + \mathbb{P}\big[\mathbf{U} = \underline{u} \mid \theta_0\big]\, u^\gamma(\underline{u})$$

$$\approx \mathbb{E}\big[u^\gamma(\tilde{\mathbf{U}}) \mid \tilde{\theta}_0\big]$$

Modified Kupper–Schachermayer–Skorokhod Embedding Theorem

$\mathbf{U}$ takes either value $\underline{u} < \overline{u}$

- the initial value

$$u^\gamma\big(\tilde{u}(\tilde{\theta}_0)\big) \approx \mathbb{E}\big[u^\gamma(\mathbf{U}) \mid \theta_0\big]$$

- the martingale property

$$\mathbb{E}\big[u^\gamma\big(\tilde{\mathbf{U}}\big) \mid \tilde{\theta}_0\big] = u^\gamma\big(\tilde{u}_0(\tilde{\theta}_0)\big)$$

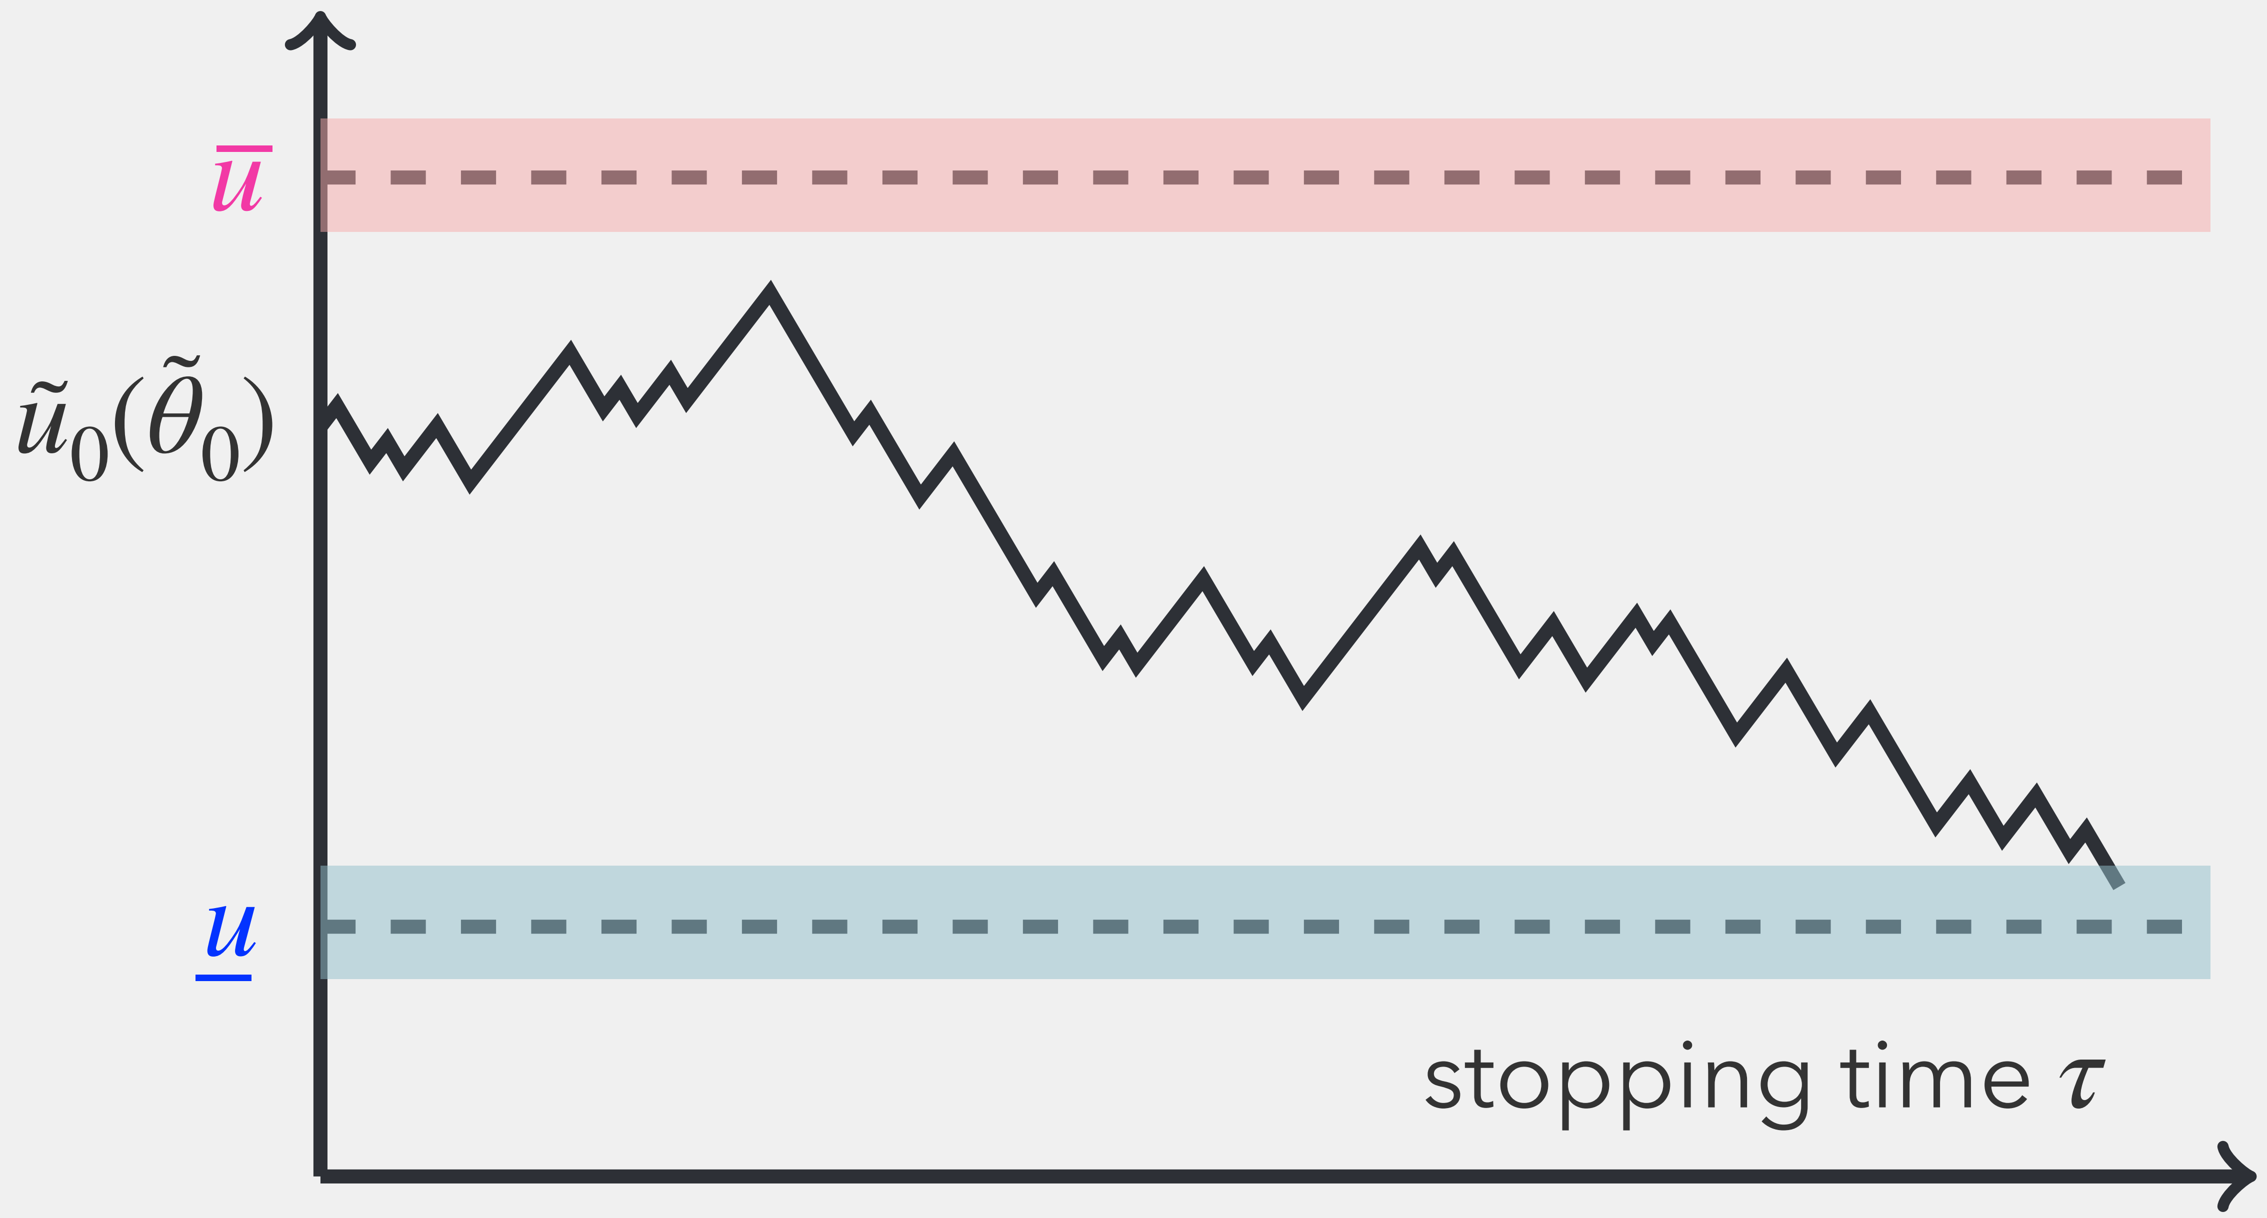

$$\begin{aligned}
&\mathbb{E}\big[u^\gamma(\mathbf{U}) \mid \theta_0\big] = \mathbb{P}\big[\mathbf{U} = \overline{u} \mid \theta_0\big]\, u^\gamma(\overline{u}) + \mathbb{P}\big[\mathbf{U} = \underline{u} \mid \theta_0\big]\, u^\gamma(\underline{u}) \\
\approx\ &\mathbb{E}\big[u^\gamma(\tilde{\mathbf{U}}) \mid \tilde{\theta}_0\big] = \mathbb{P}\big[\tilde{\mathbf{U}} \approx \overline{u} \mid \tilde{\theta}_0\big]\, u^\gamma(\overline{u}) + \mathbb{P}\big[\tilde{\mathbf{U}} \approx \underline{u} \mid \tilde{\theta}_0\big]\, u^\gamma(\underline{u})
\end{aligned}$$

Modified Kupper–Schachermayer–Skorokhod Embedding Theorem

$\mathbf{U}$ takes either value $\underline{u} < \overline{u}$

- the initial value

$$u^\gamma\big(\tilde{u}(\tilde{\theta}_0)\big) \approx \mathbb{E}\big[u^\gamma(\mathbf{U}) \mid \theta_0\big]$$

- the martingale property

$$\mathbb{E}\big[u^\gamma(\tilde{\mathbf{U}}) \mid \tilde{\theta}_0\big] = u^\gamma\big(\tilde{u}_0(\tilde{\theta}_0)\big)$$

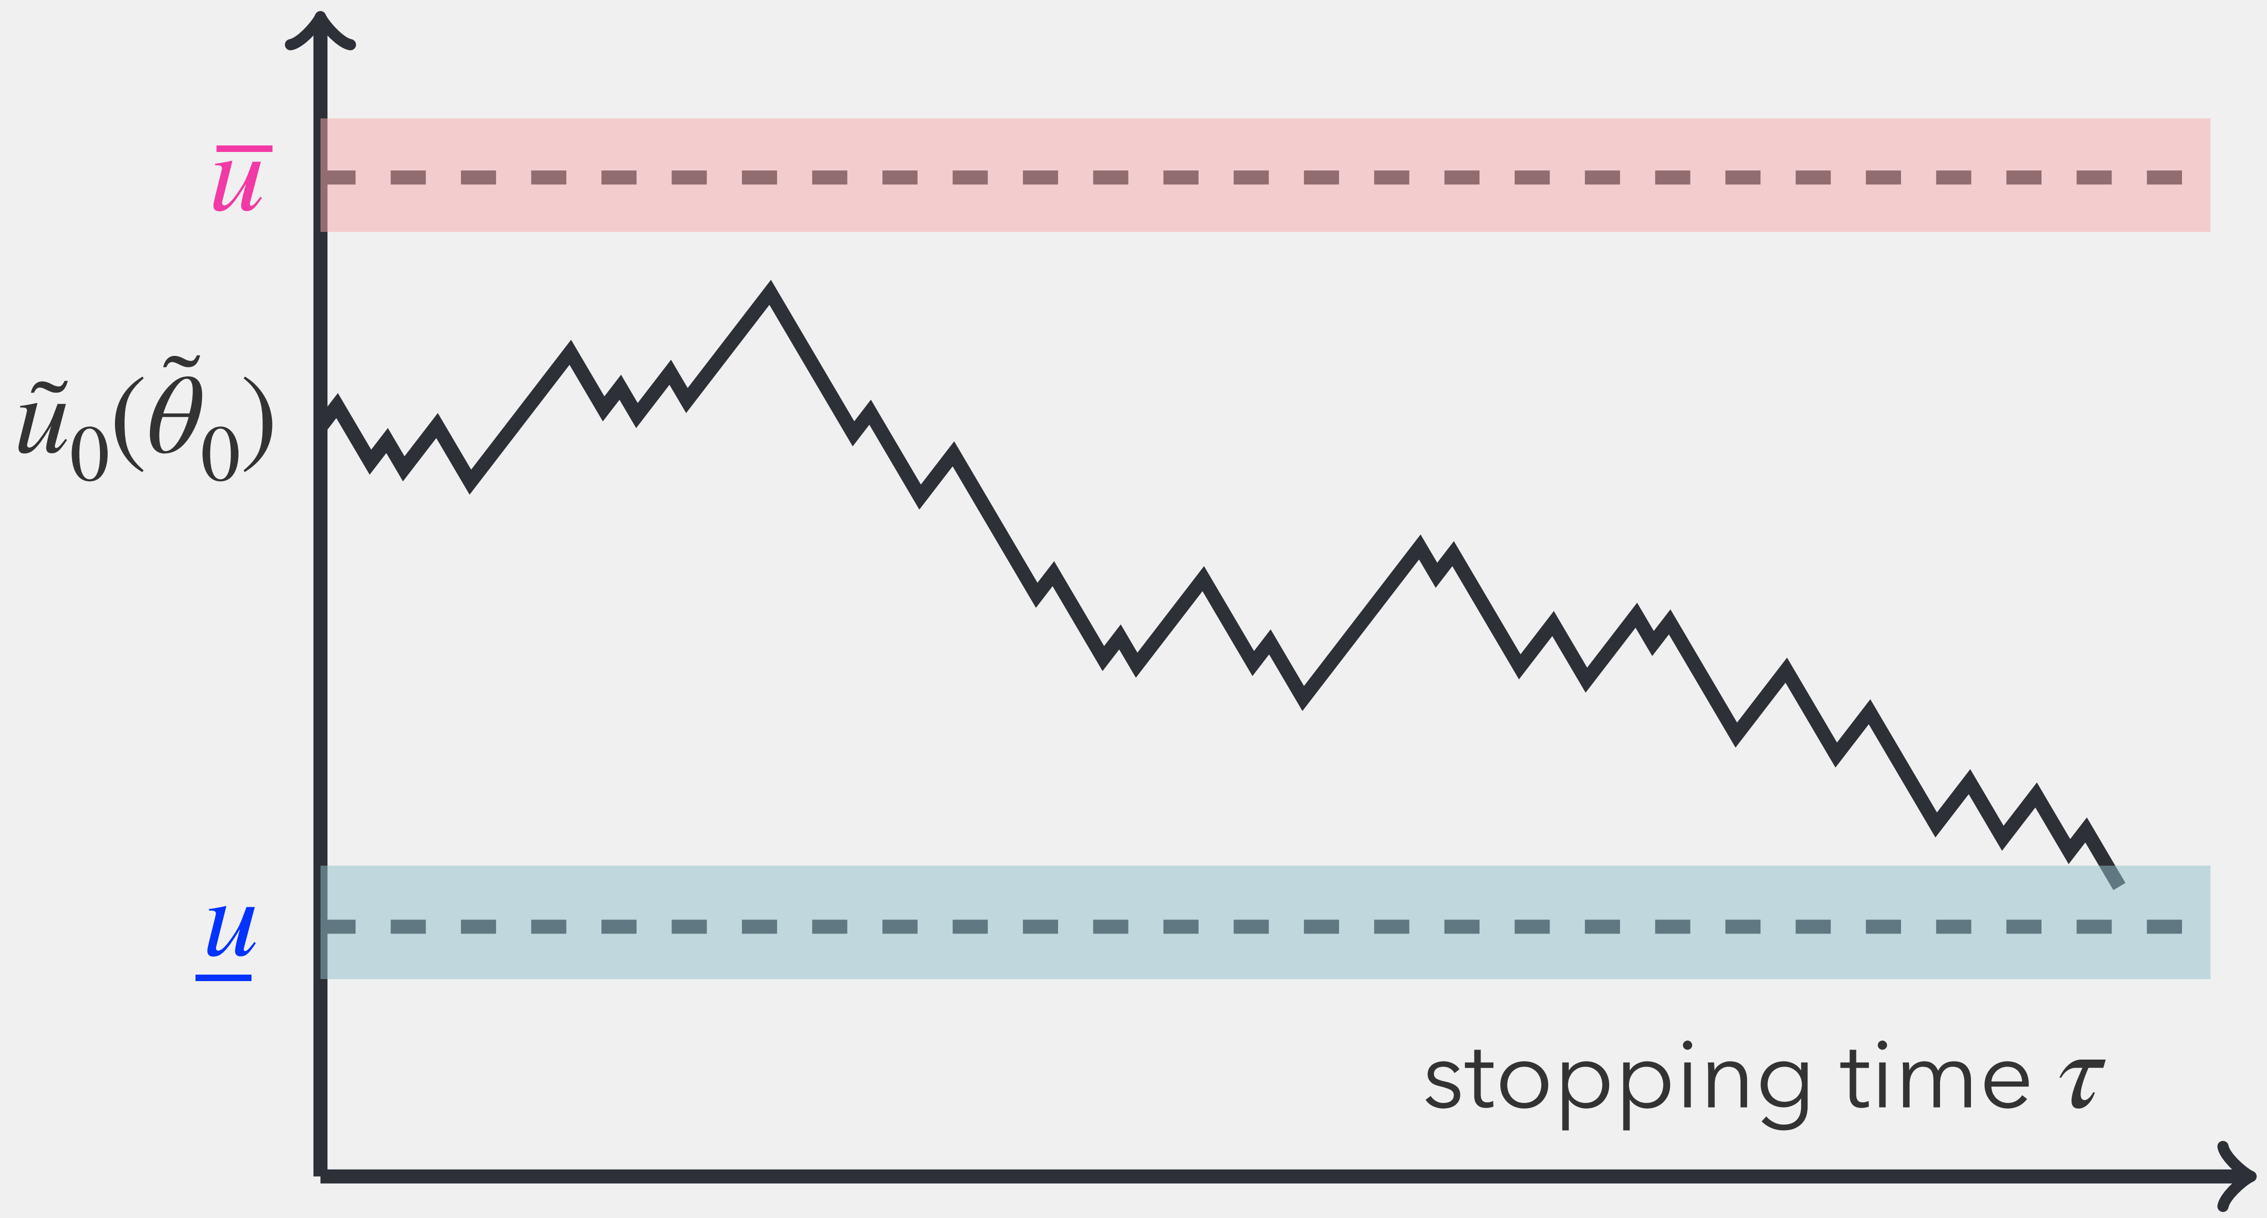

$$\begin{aligned} \mathbb{E}\big[u^\gamma(\mathbf{U}) \mid \theta_0\big] &= \mathbb{P}\big[\mathbf{U} = \overline{u} \mid \theta_0\big]\, u^\gamma(\overline{u}) + \mathbb{P}\big[\mathbf{U} = \underline{u} \mid \theta_0\big]\, u^\gamma(\underline{u}) \\ \approx \mathbb{E}\big[u^\gamma(\tilde{\mathbf{U}}) \mid \tilde{\theta}_0\big] &= \mathbb{P}\big[\tilde{\mathbf{U}} \approx \overline{u} \mid \tilde{\theta}_0\big]\, u^\gamma(\overline{u}) + \mathbb{P}\big[\tilde{\mathbf{U}} \approx \underline{u} \mid \tilde{\theta}_0\big]\, u^\gamma(\underline{u}) \end{aligned}$$

$$\longrightarrow \quad \mathbf{U}\,|_{\theta_0} \approx \tilde{\mathbf{U}}\,|_{\tilde{\theta}_0}$$

modified Kupper–Schachermayer–Skorokhod Embedding Theorem

$$\mathbf{U}\,|_{\theta_0} \approx \tilde{\mathbf{U}}\,|_{\theta_0}$$

$$\mathbb{V}\big[\mathbf{U}\big] \approx \mathbb{V}\big[\tilde{\mathbf{U}}\big] = \mathbb{E}\big[\tilde{u}_0(\tilde{\boldsymbol{\theta}}_0)\big] \approx \mathbb{E}\Big[(u^\gamma)^{-1}\Big(\mathbb{E}\big[u^\gamma(\mathbf{U}) \mid \boldsymbol{\theta}_0\big]\Big)\Big]$$

martingale-ish

certainty-equivalent of CARA utility

$$\mathbb{E}\left[\frac{1}{\gamma}\log \mathbb{E}\big[e^{\gamma \mathbf{U}} \mid \boldsymbol{\theta}_0\big]\right]$$

modified Kupper–Schachermayer–Skorokhod Embedding Theorem

$$\mathbf{U}\,|_{\theta_0} \approx \tilde{\mathbf{U}}\,|_{\theta_0}$$

$$\mathbb{V}[\mathbf{U}] \approx \mathbb{V}[\tilde{\mathbf{U}}] = \mathbb{E}\left[\tilde{u}_0(\tilde{\boldsymbol{\theta}}_0)\right] \approx \mathbb{E}\left[(u^\gamma)^{-1}\left(\mathbb{E}\left[u^\gamma(\mathbf{U}) \mid \boldsymbol{\theta}_0\right]\right)\right]$$

martingale-ish

certainty-equivalent of CARA utility

equal at limit $\epsilon \to 0$

$$\mathbb{E}\left[\frac{1}{\gamma}\log \mathbb{E}\left[e^{\gamma \mathbf{U}} \mid \boldsymbol{\theta}_0\right]\right]$$