


# Causal Imitation Learning under Temporally Correlated Noise

Gokul Swamy




*(joint work with Sanjiban Choudhury, Drew Bagnell, Steven Wu)*



$$\pi_E \xleftrightarrow{f} \pi$$


*Offline*



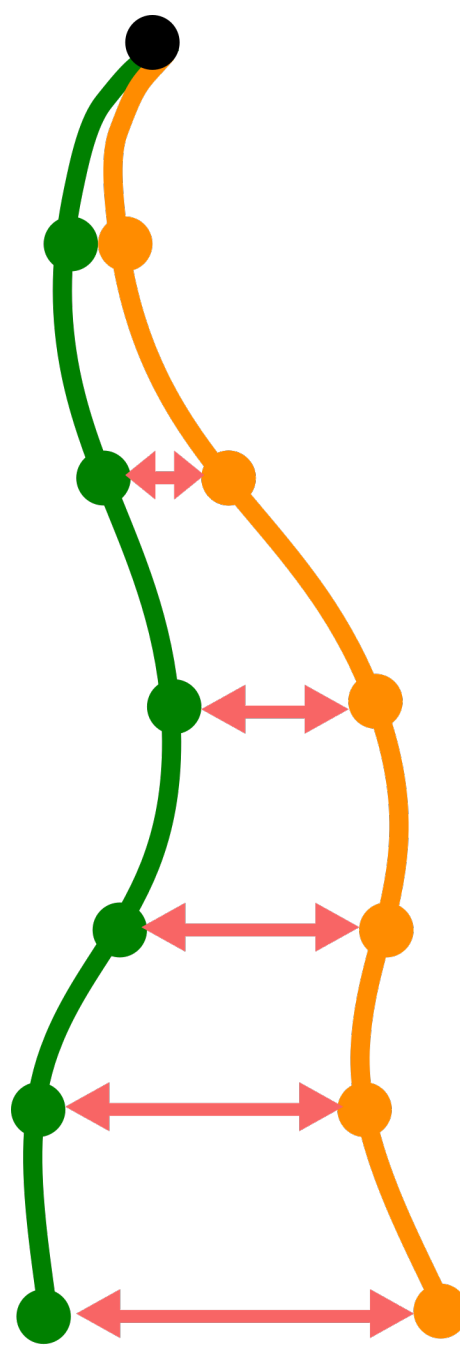
*Behavioral Cloning ...*

$$\pi_E \xleftrightarrow{f} \pi$$

*Offline*



*Online*




*Behavioral Cloning ...*


*GAIL, MaxEnt IRL ...*

$$\pi_E \xleftrightarrow{f} \pi$$


*Offline*



*Online*



*Interactive*



*Behavioral Cloning ...*


*GAIL, MaxEnt IRL ...*

*Dagger ...*


$$\pi_E \xleftrightarrow{f} \pi$$

[Swamy et al, 2021]


Offline



Online



Interactive



$$J(\pi_E) - J(\pi) \leq O(\epsilon T^2)$$

$$J(\pi_E) - J(\pi) \leq O(\epsilon T)$$

$$J(\pi_E) - J(\pi) \leq O(\epsilon HT)$$

Behavioral Cloning ...

GAIL, MaxEnt IRL ...

Dagger ...



*“Hence, a system trained with multiple frames would merely predict a steering angle equal to the current rate of turn as observed through the camera. This would lead to catastrophic behavior in test mode. **The robot would simply turn in circles.**”*

*— Muller et al., 2006*

*“Hence, a system trained with multiple frames would merely predict a steering angle equal to the current rate of turn as observed through the camera. This would lead to catastrophic behavior in test mode. **The robot would simply turn in circles.**”*

*— Muller et al., 2006*

*“Actually, since we were fitting a model to a time-series, **samples tend to be correlated in time** [...] Thus, when leaving out a sample in cross validation, **we actually left out a large window (16 seconds) of data** around that sample, to diminish this bias.”*


*— Ng et al., 2003*









$\pi E$



$\pi_E$



$\pi_{BC}$

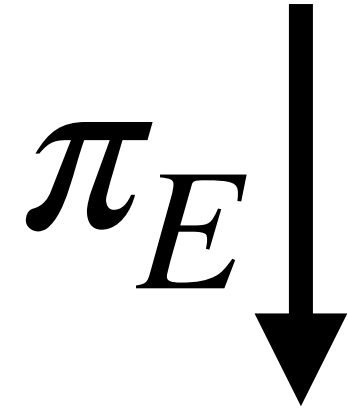






$$S_{t-2}$$





$S_{t-2}$



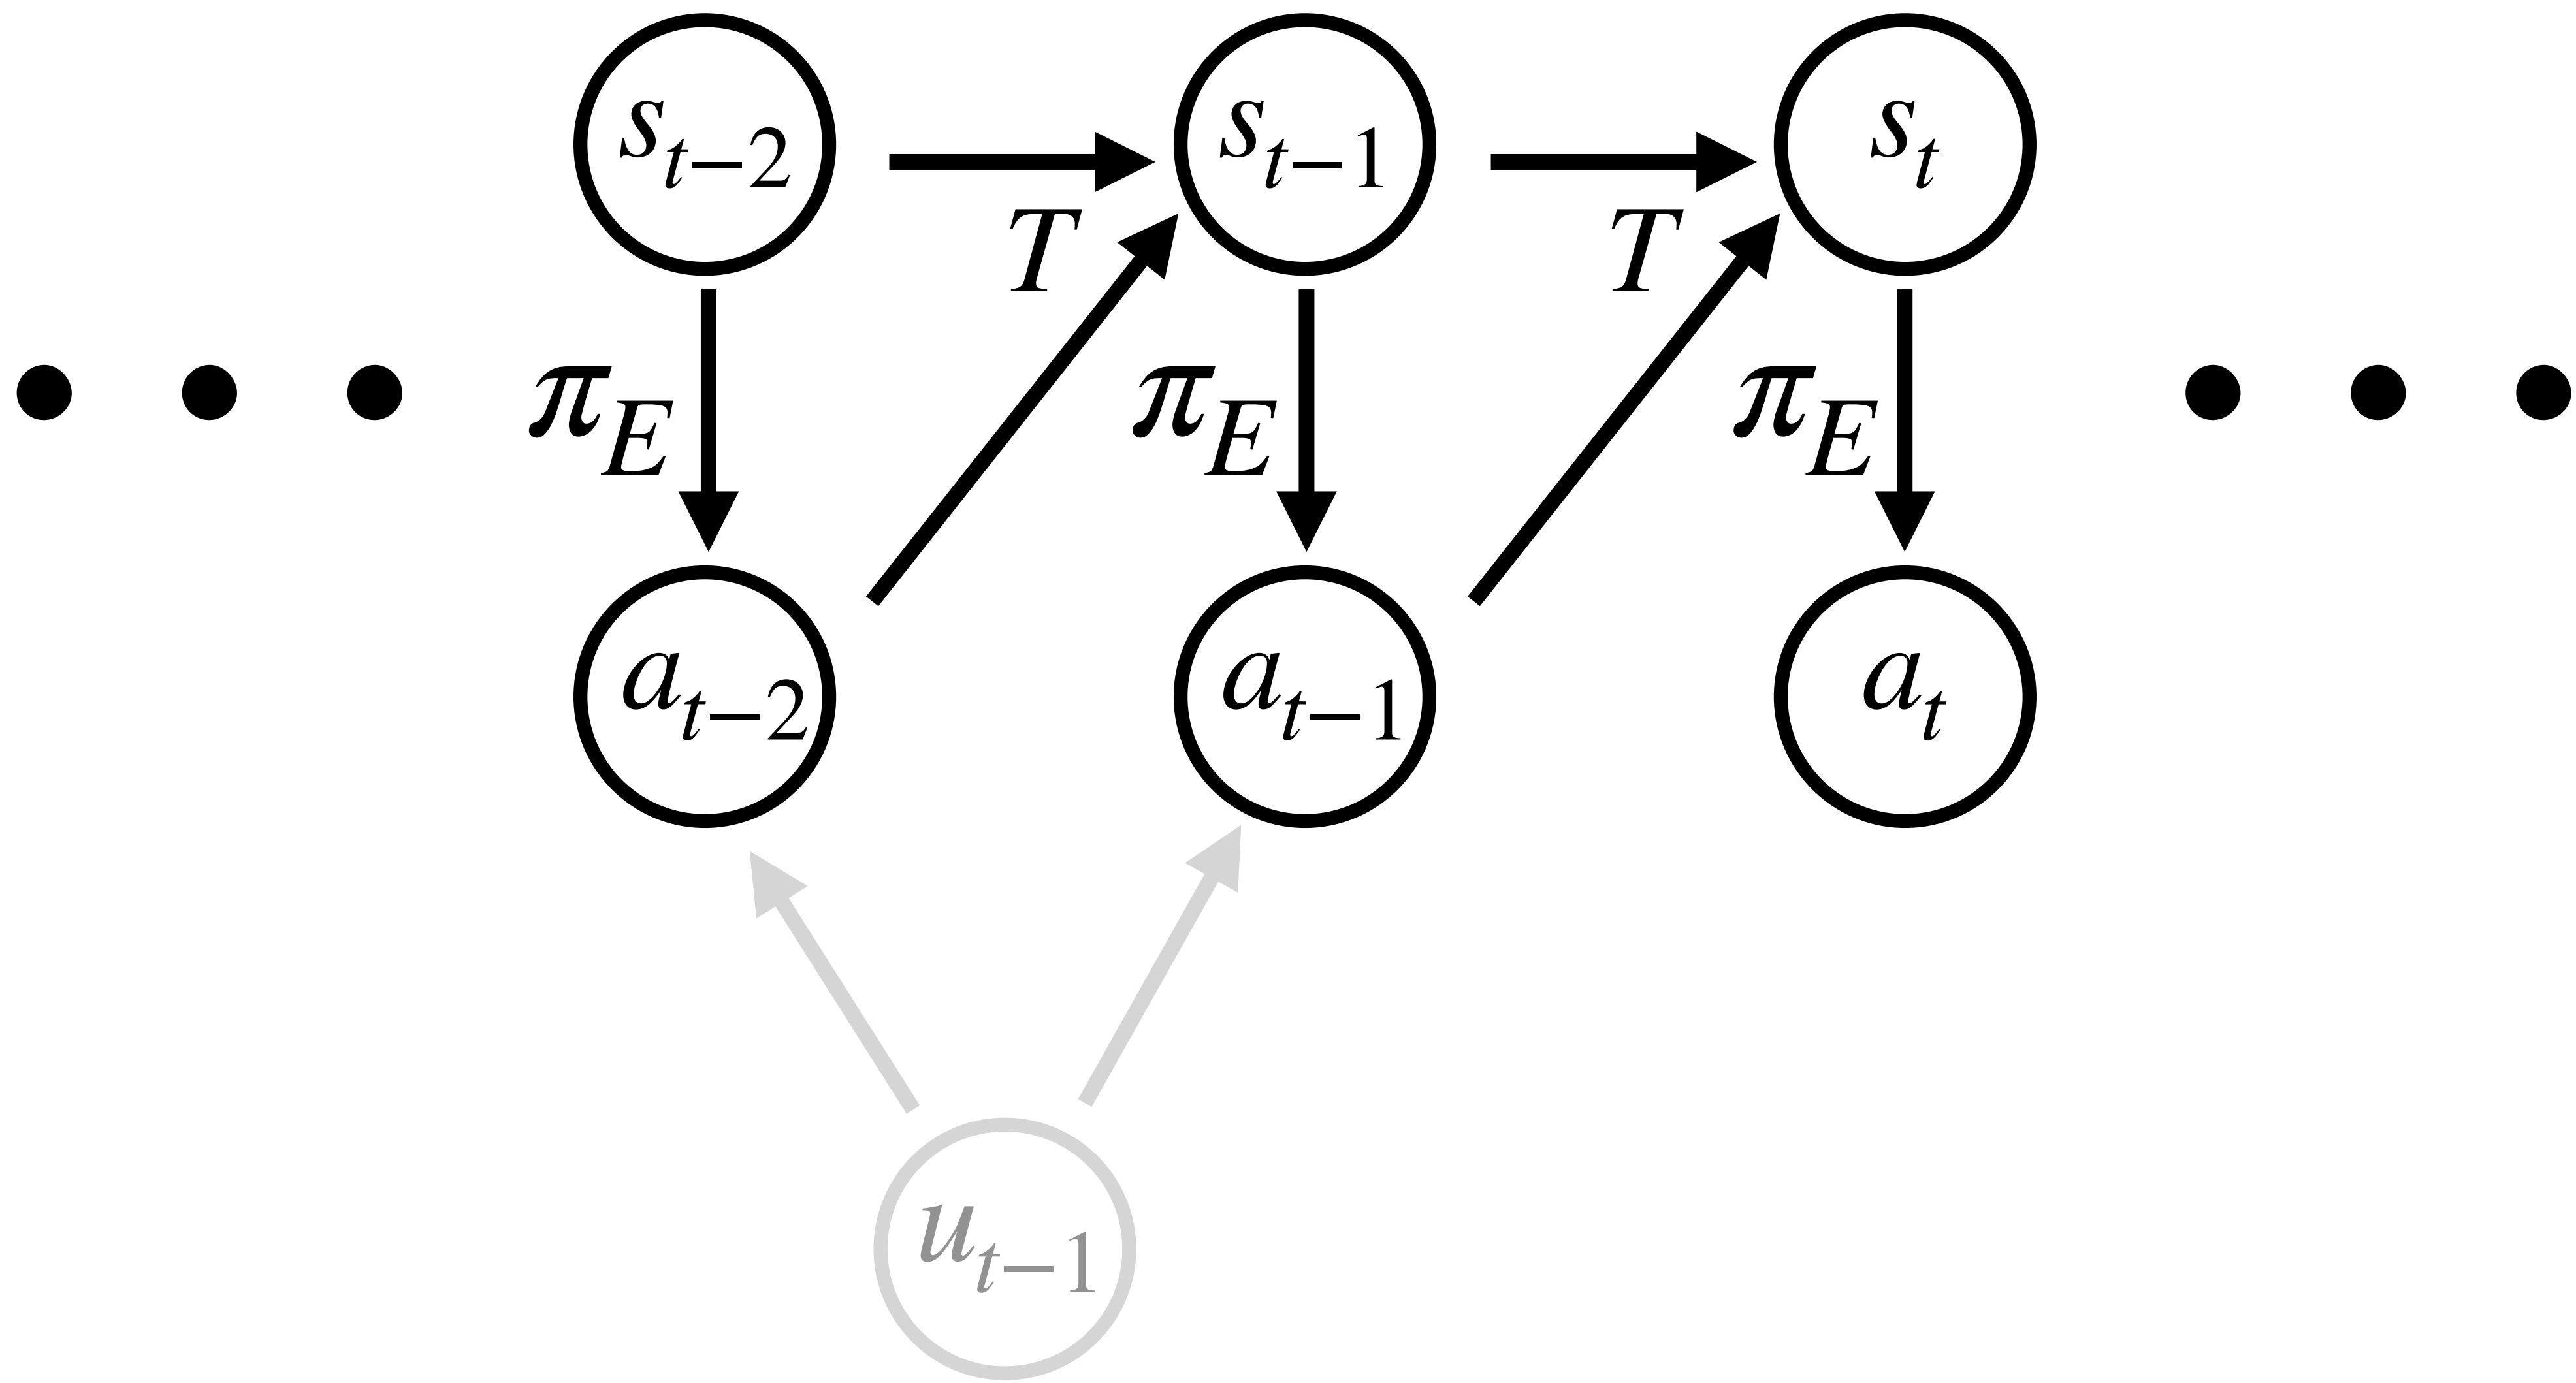




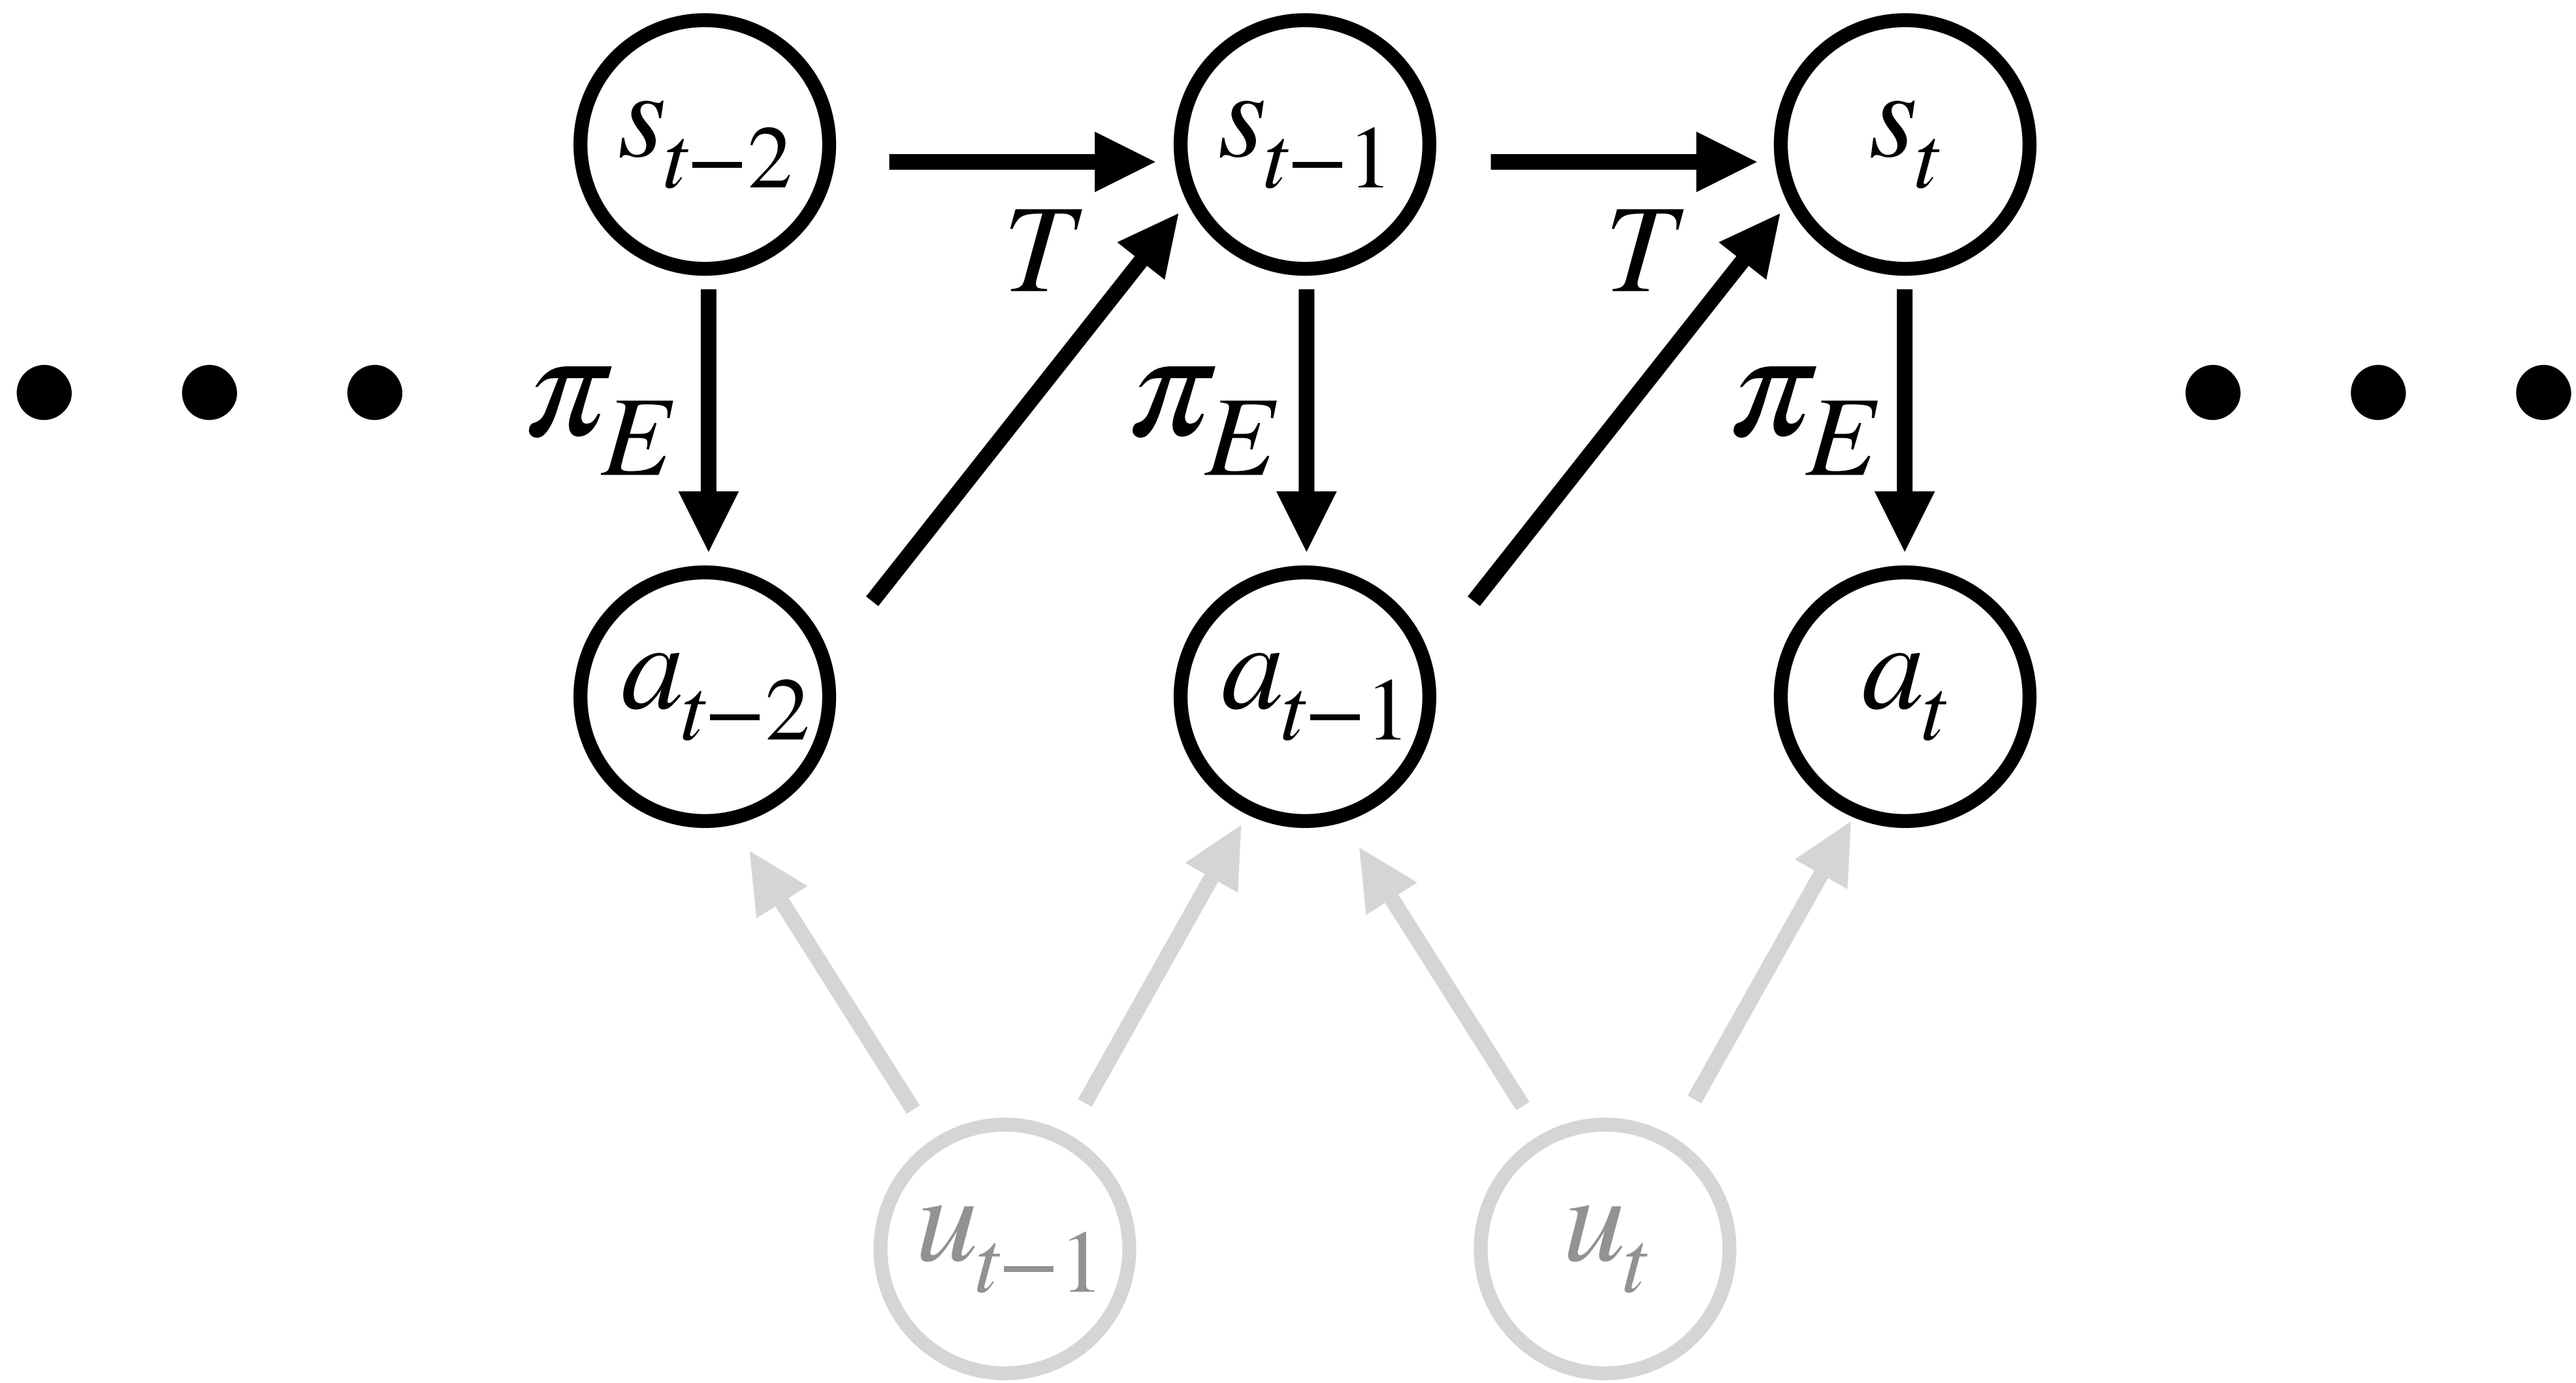





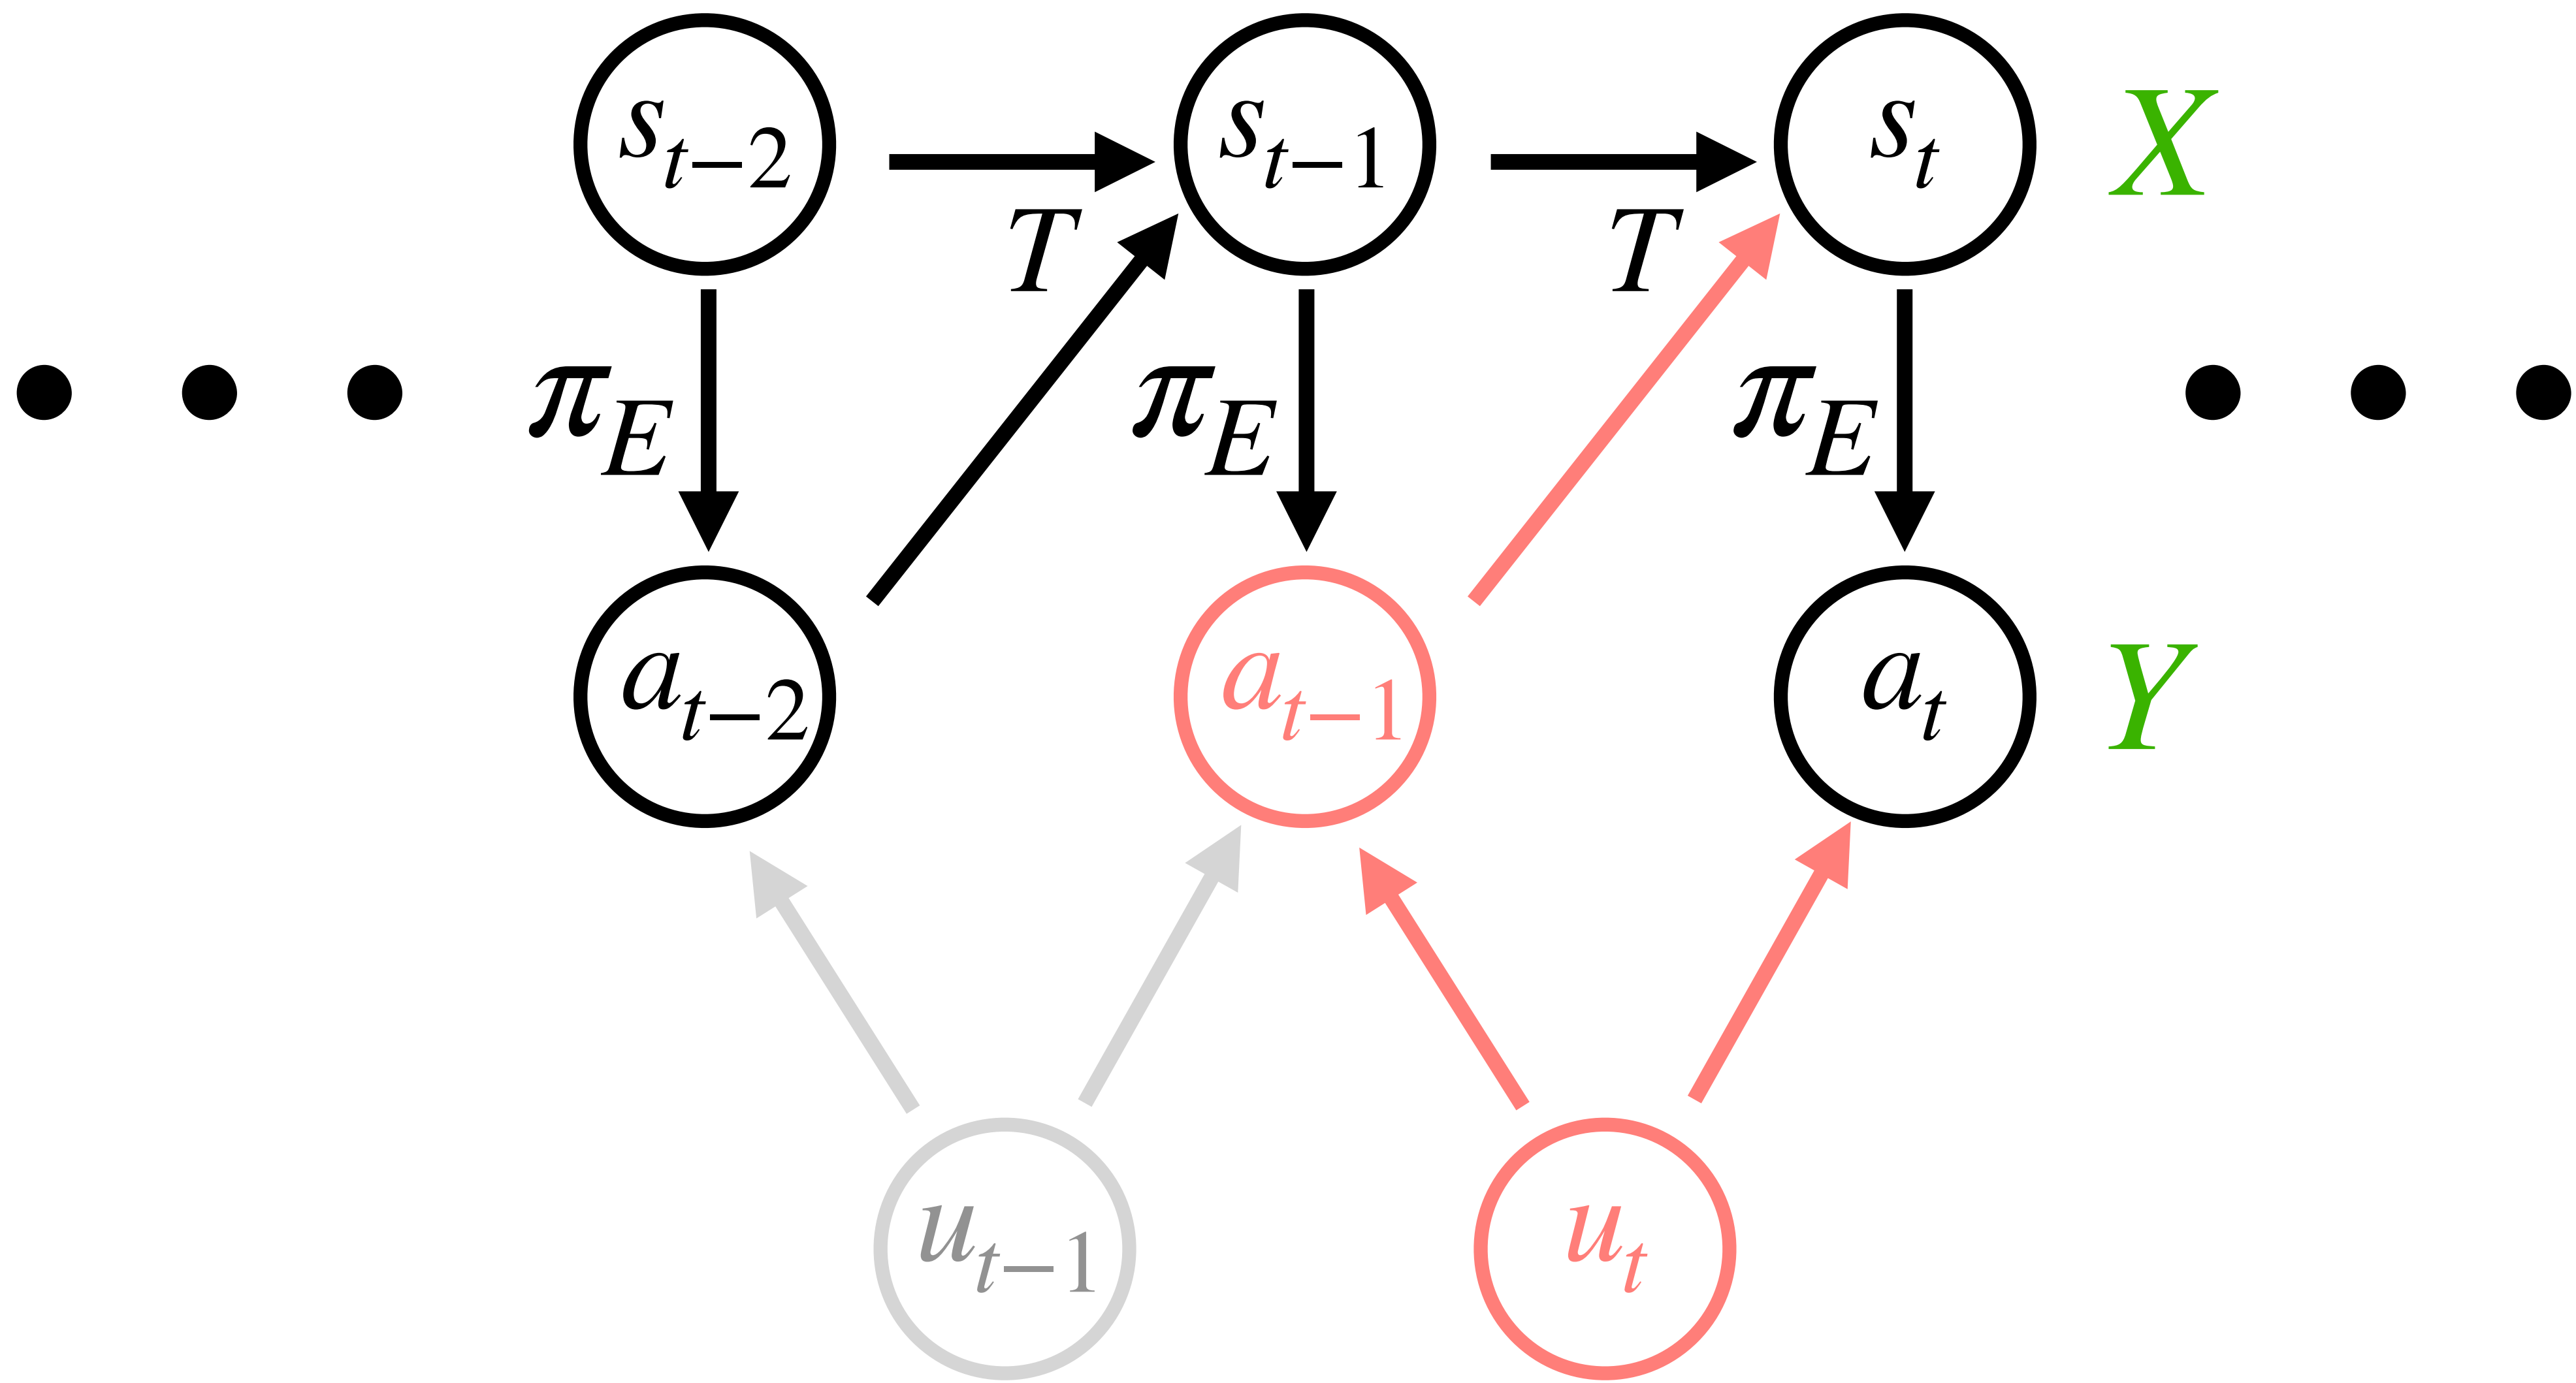
$u_{t-1}$














$\pi_E$




$\pi_{BC}$




$\pi_E$



$\pi_{BC}$




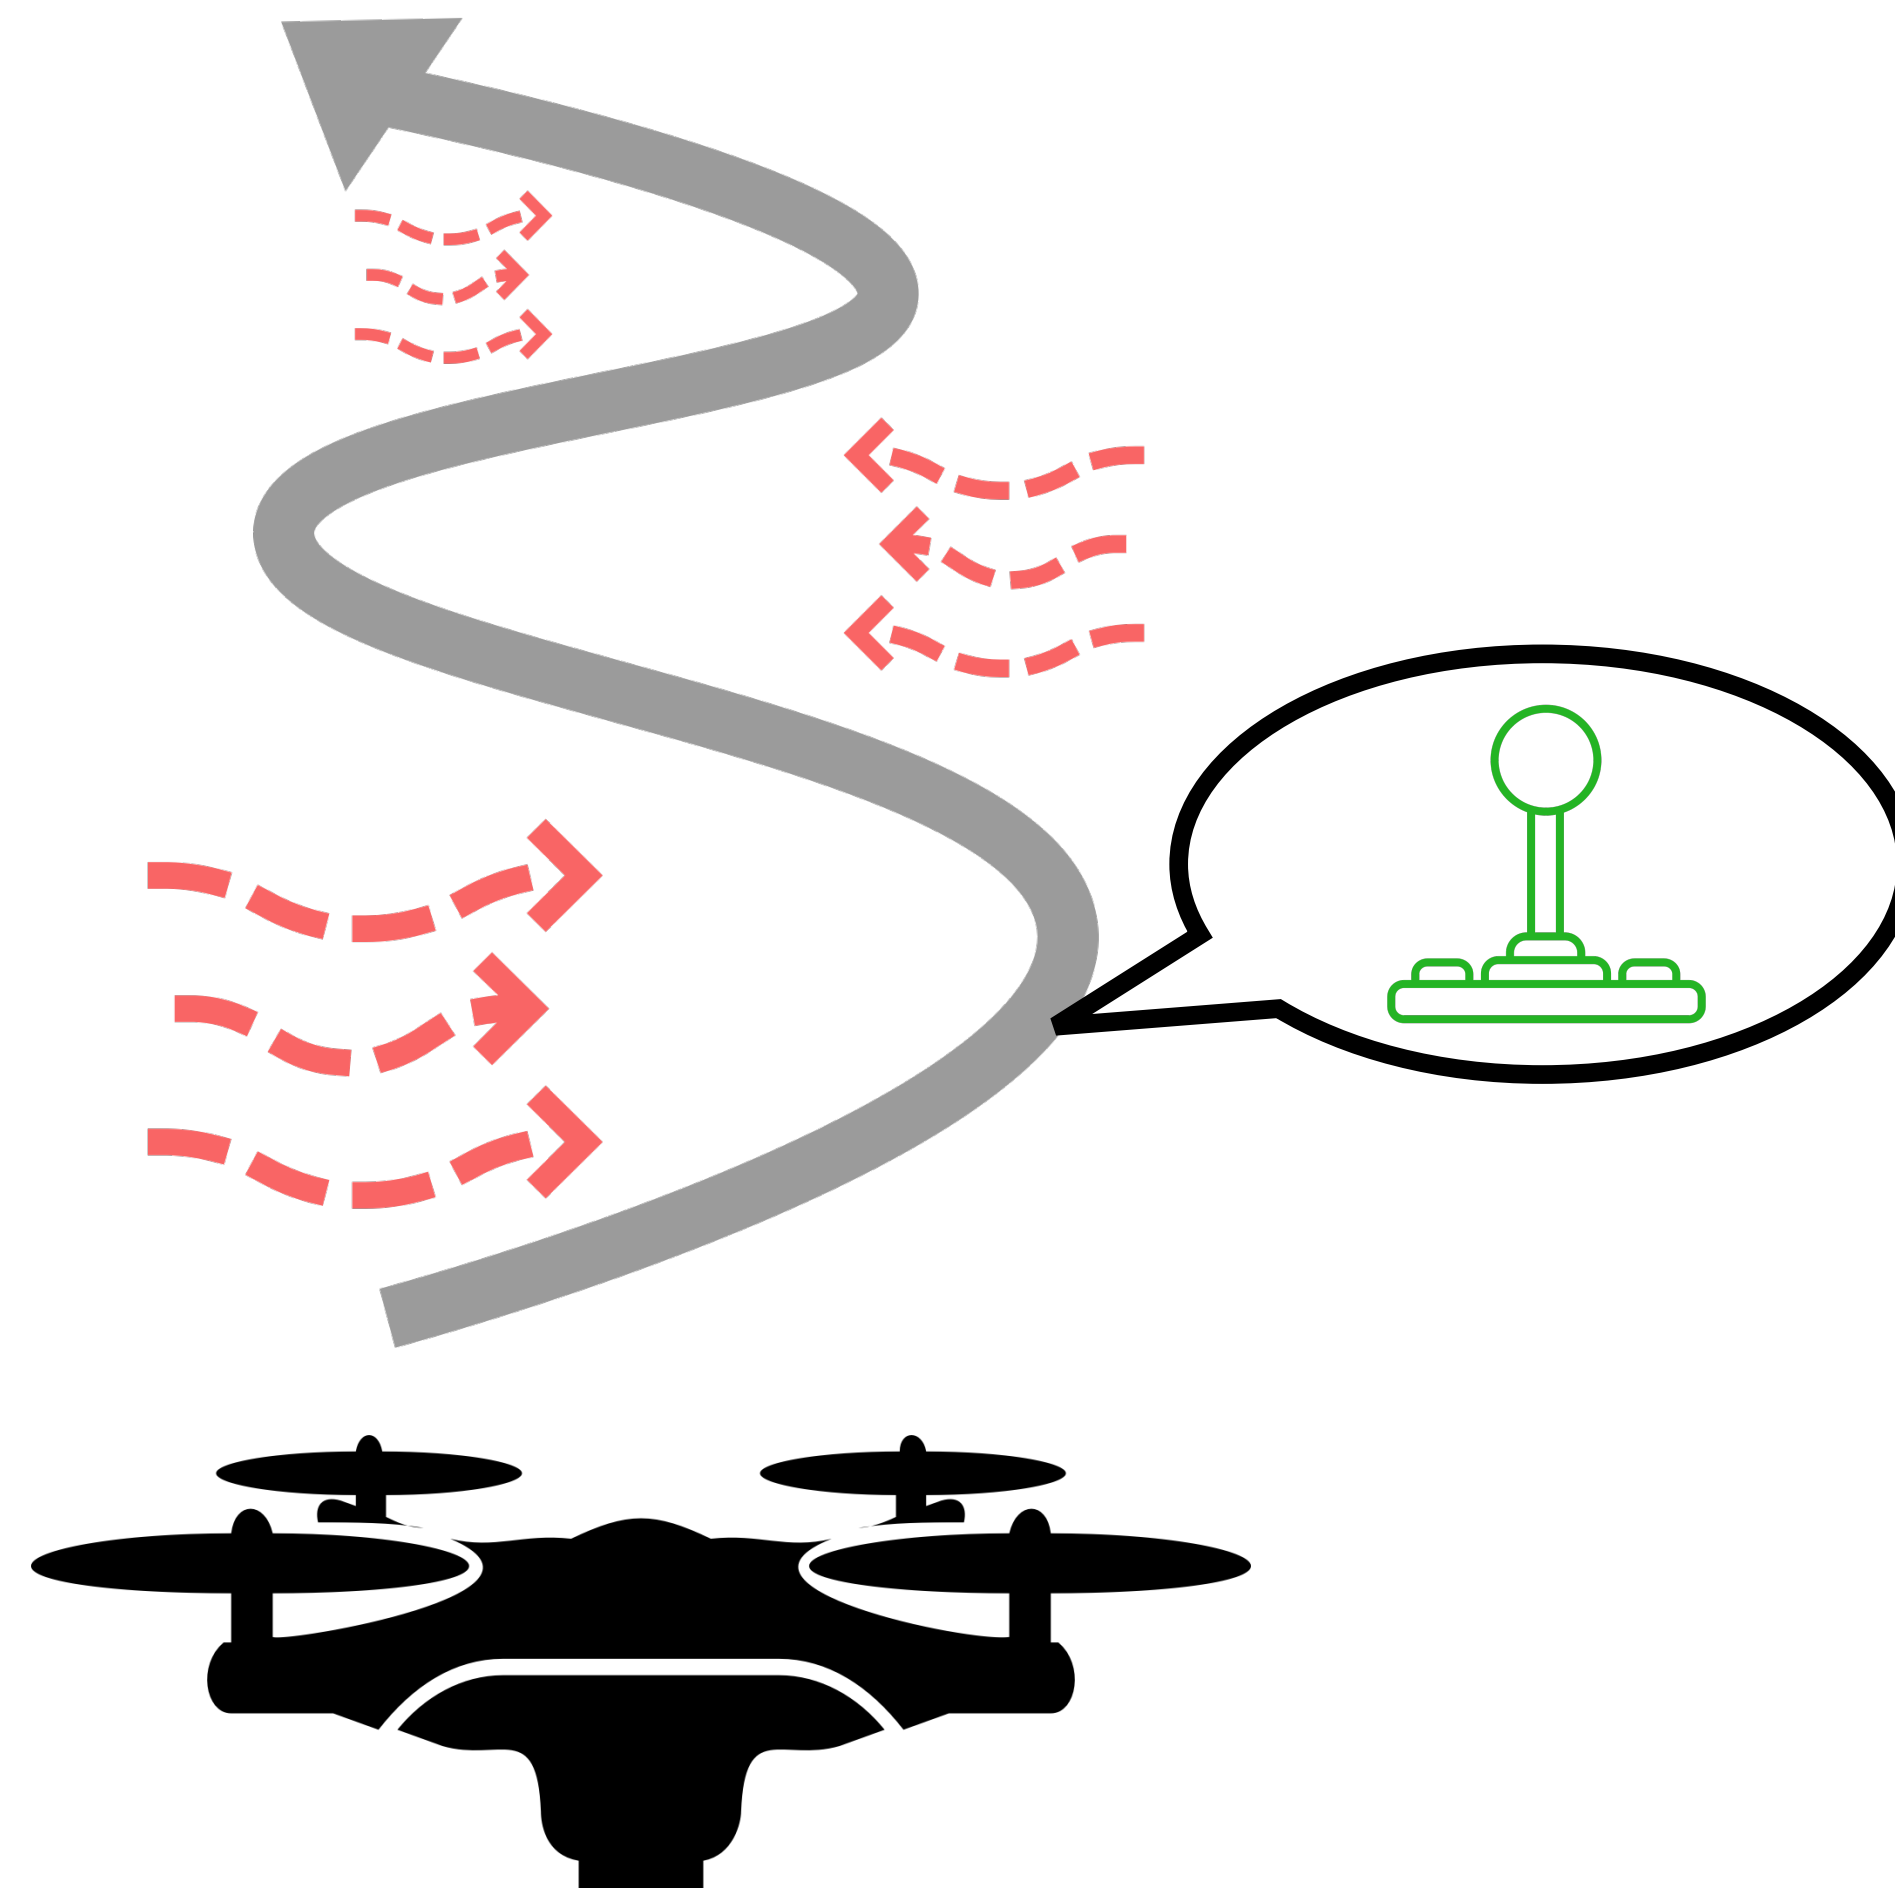
$\pi$











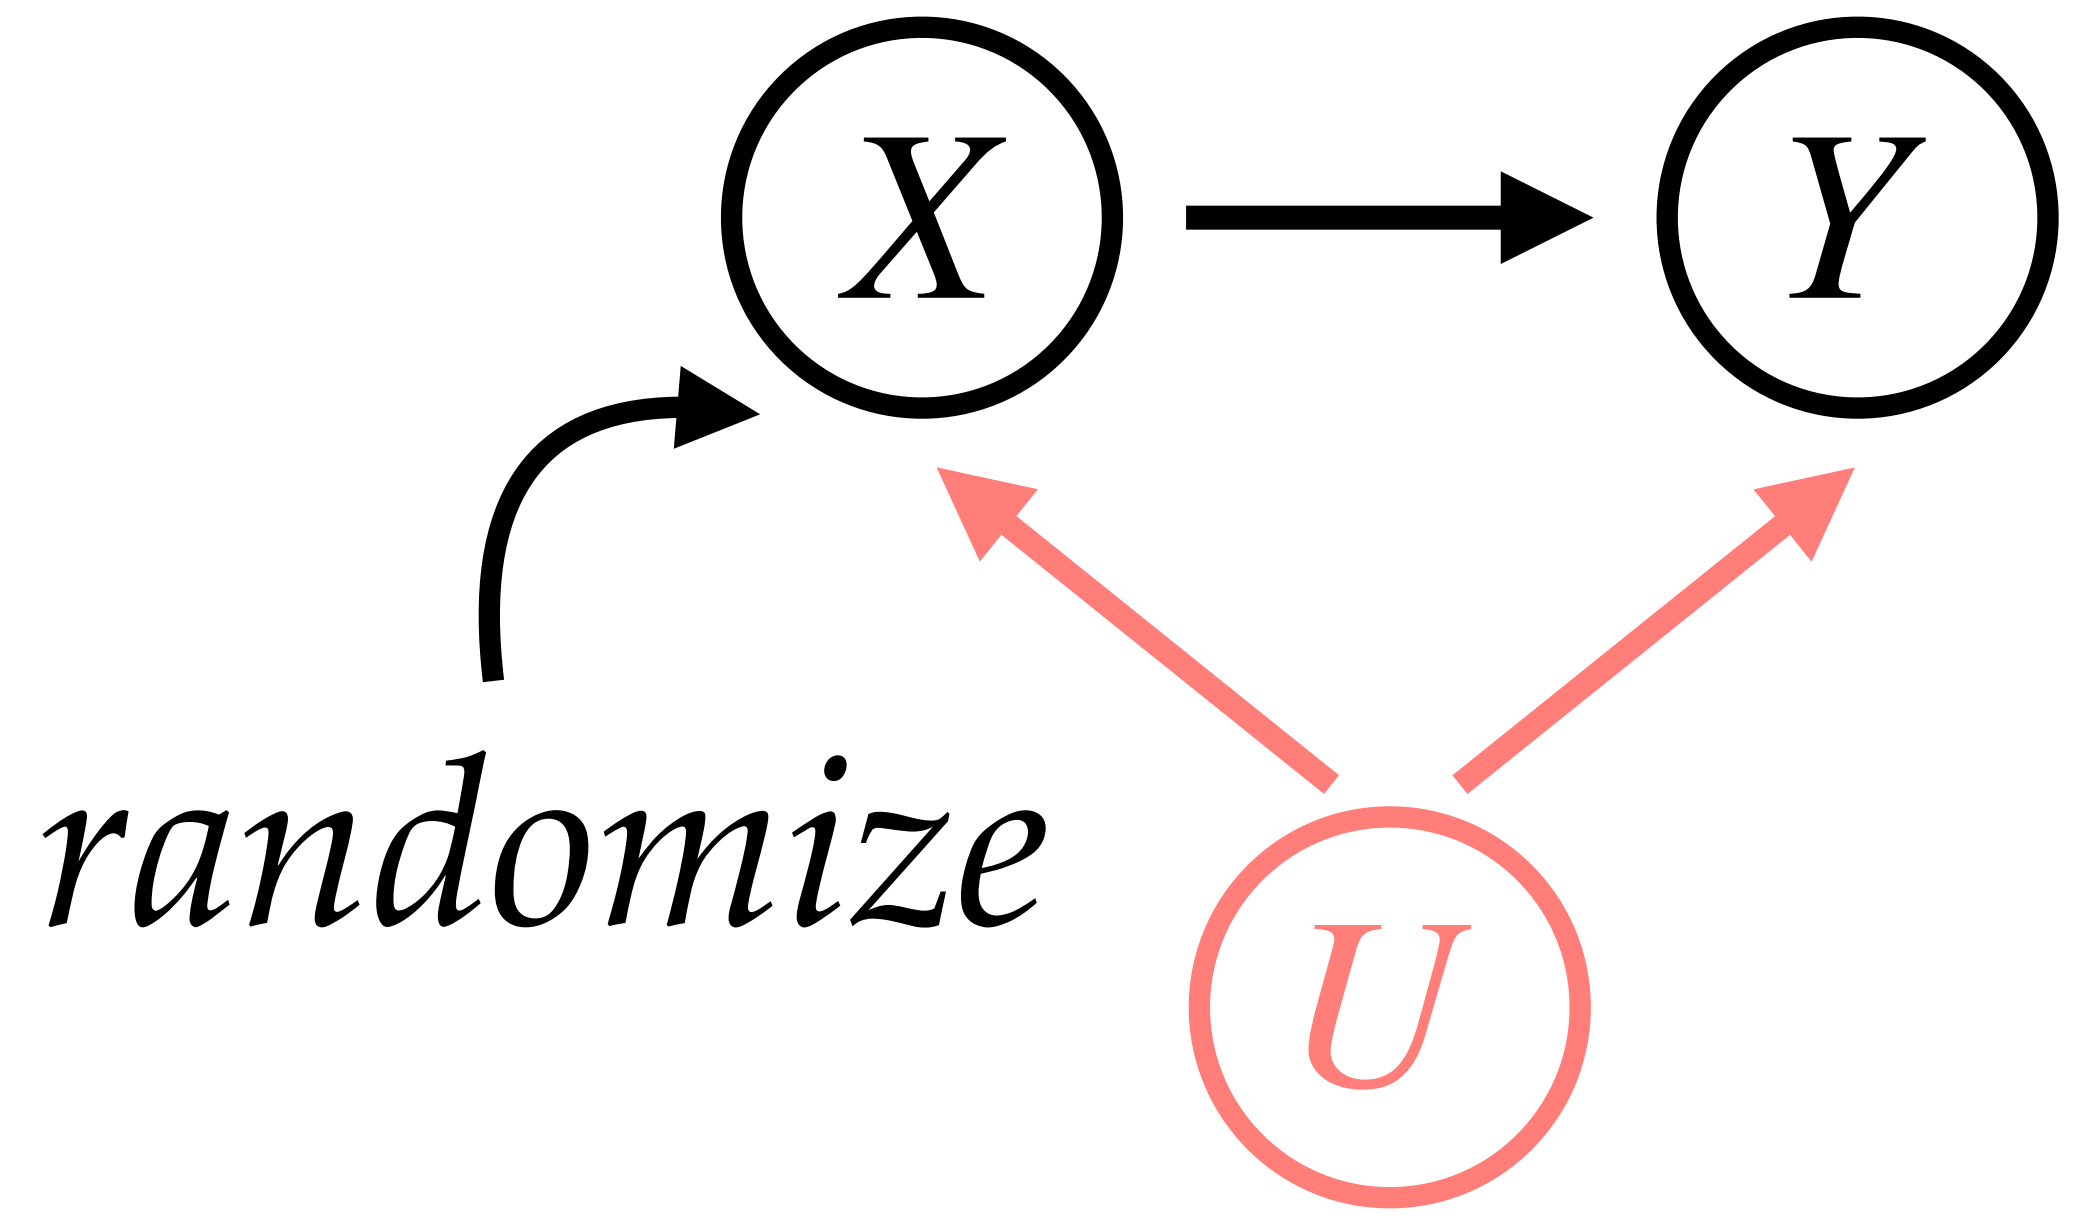


























*Key Idea: We can condition on instrument  $Z$  to counter the effect of confounder  $U$  on  $X$ .*






$$X = g(Z, U)$$



$$X = g(Z, U)$$


$$Y = h(X) + U$$



$$X = g(Z, U)$$

$$Y = h(X) + U$$

$$\mathbb{E}[U] = 0$$




$$X = g(Z, U)$$

$$Y = h(X) + U$$

$$\mathbb{E}[U] = 0$$

$$0 = \mathbb{E}[U] = \mathbb{E}[U | z] = \mathbb{E}[Y - h(X) | z]$$



$$X = g(Z, U)$$


$$Y = h(X) + U$$

$$\mathbb{E}[U] = 0$$

$$0 = \mathbb{E}[U] = \mathbb{E}[U | z] = \mathbb{E}[Y - h(X) | z]$$

$$\Rightarrow \mathbb{E}[Y | z] = \mathbb{E}[h(X) | z], \forall z$$





$$X = g(Z, U)$$


$$Y = h(X) + U$$

$$\mathbb{E}[U] = 0$$

$$0 = \mathbb{E}[U] = \mathbb{E}[U | z] = \mathbb{E}[Y - h(X) | z]$$

$$\Rightarrow \mathbb{E}[Y | z] = \mathbb{E}[h(X) | z], \forall z$$

$$\Rightarrow \min_h \mathbb{E}_z[(\mathbb{E}[Y | z] - \mathbb{E}[h(X) | z])^2]$$



$$X = g(Z, U)$$

$$Y = h(X) + U$$


$$\mathbb{E}[U] = 0$$


$$0 = \mathbb{E}[U] = \mathbb{E}[U | z] = \mathbb{E}[Y - h(X) | z]$$


$$\Rightarrow \mathbb{E}[Y | z] = \mathbb{E}[h(X) | z], \forall z$$


$$\Rightarrow \min_h \mathbb{E}_z[(\mathbb{E}[Y | z] - \mathbb{E}[h(X) | z])^2]$$


$$\Leftrightarrow \min_h \max_f \mathbb{E}_z[2(Y - h(X))f(Z) - f^2(Z)]$$

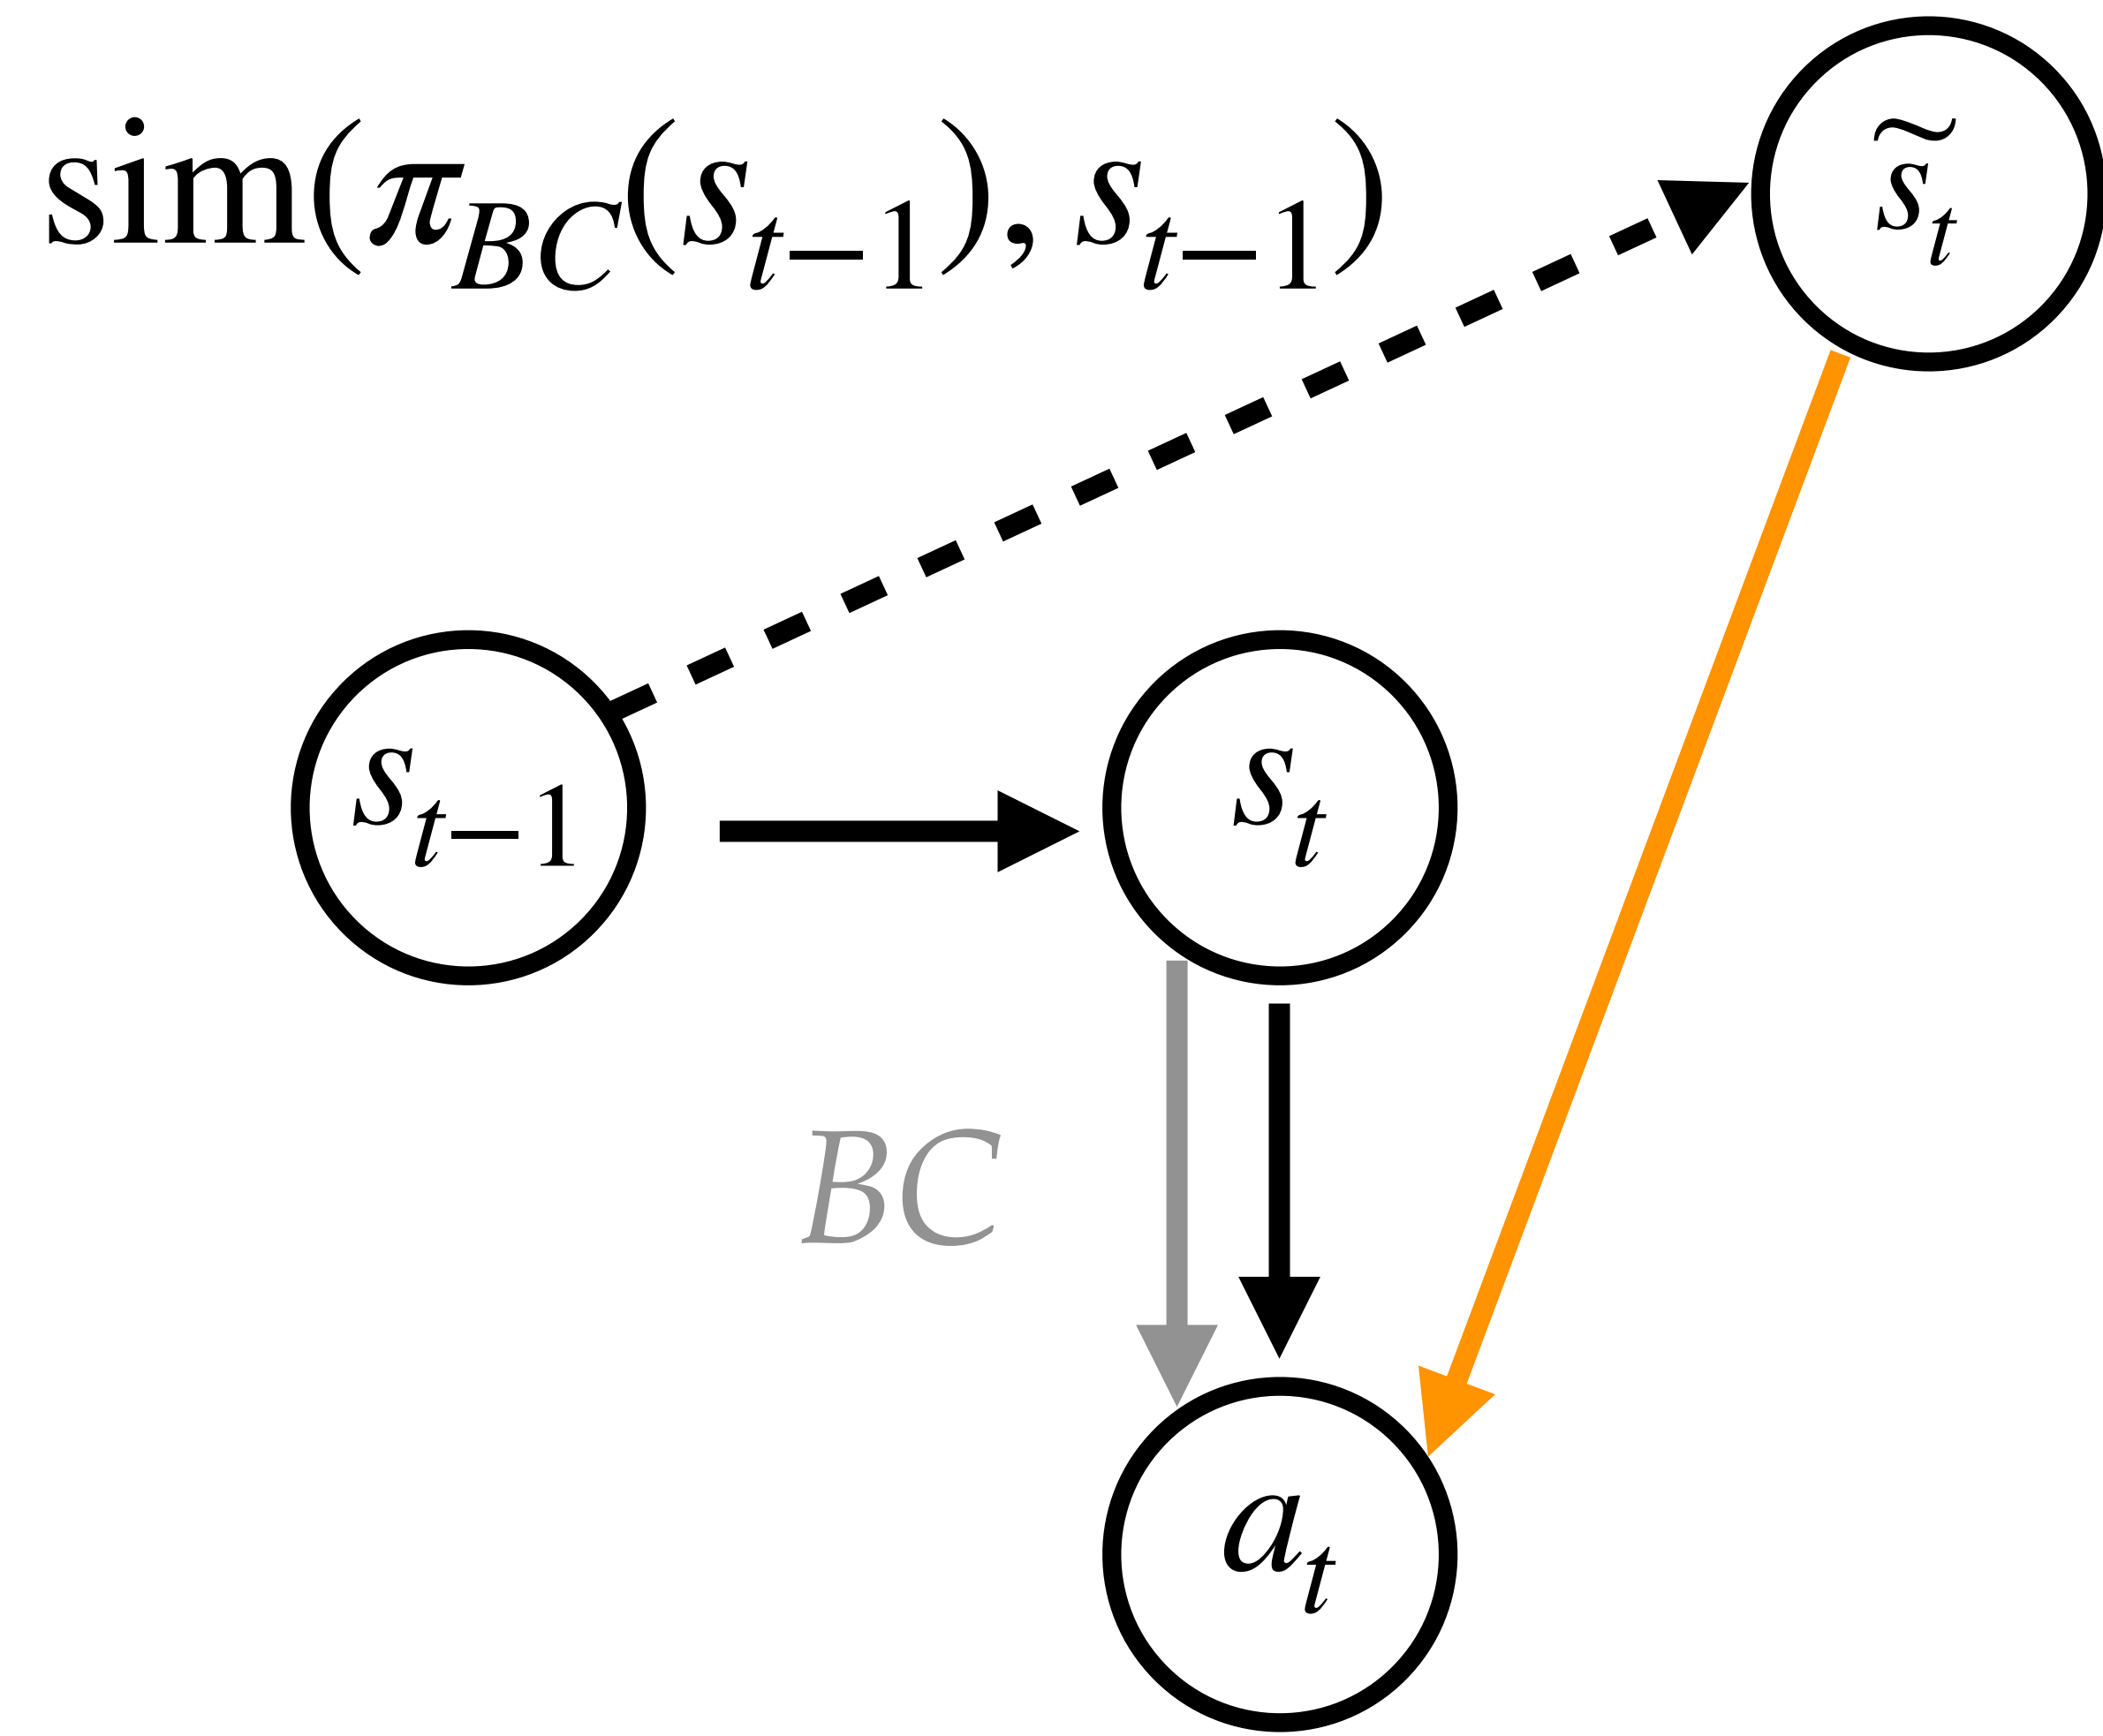




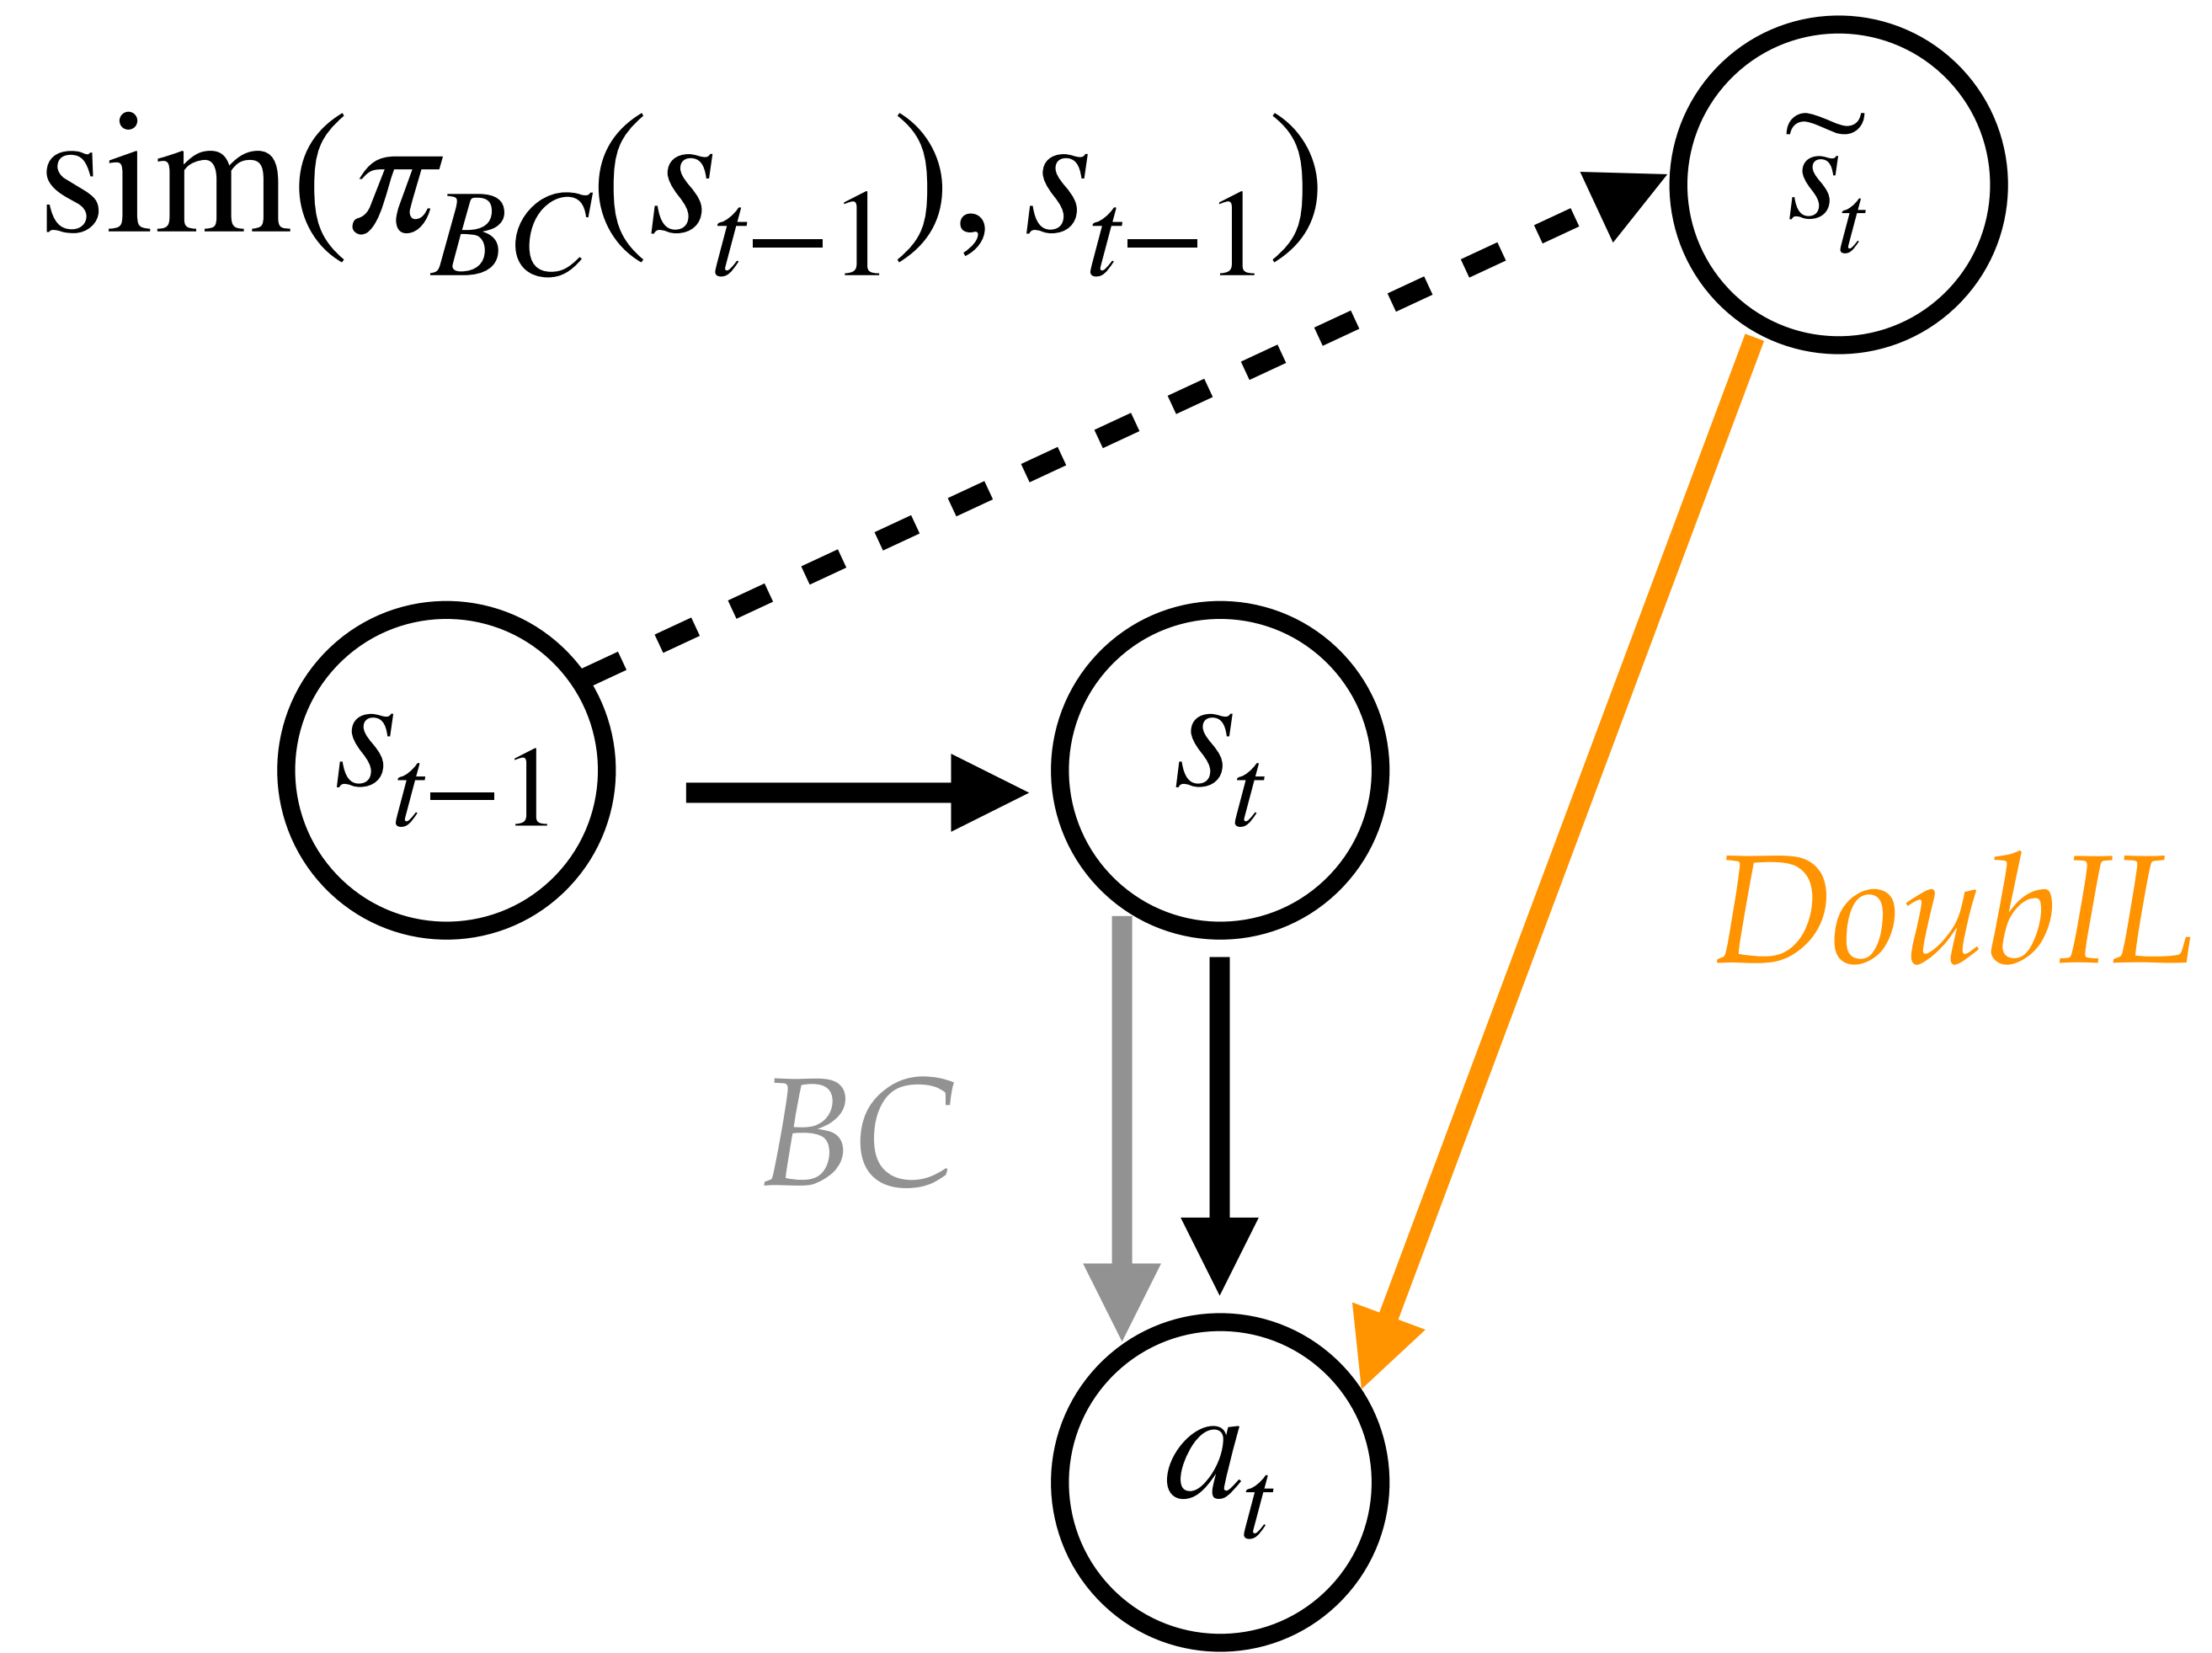


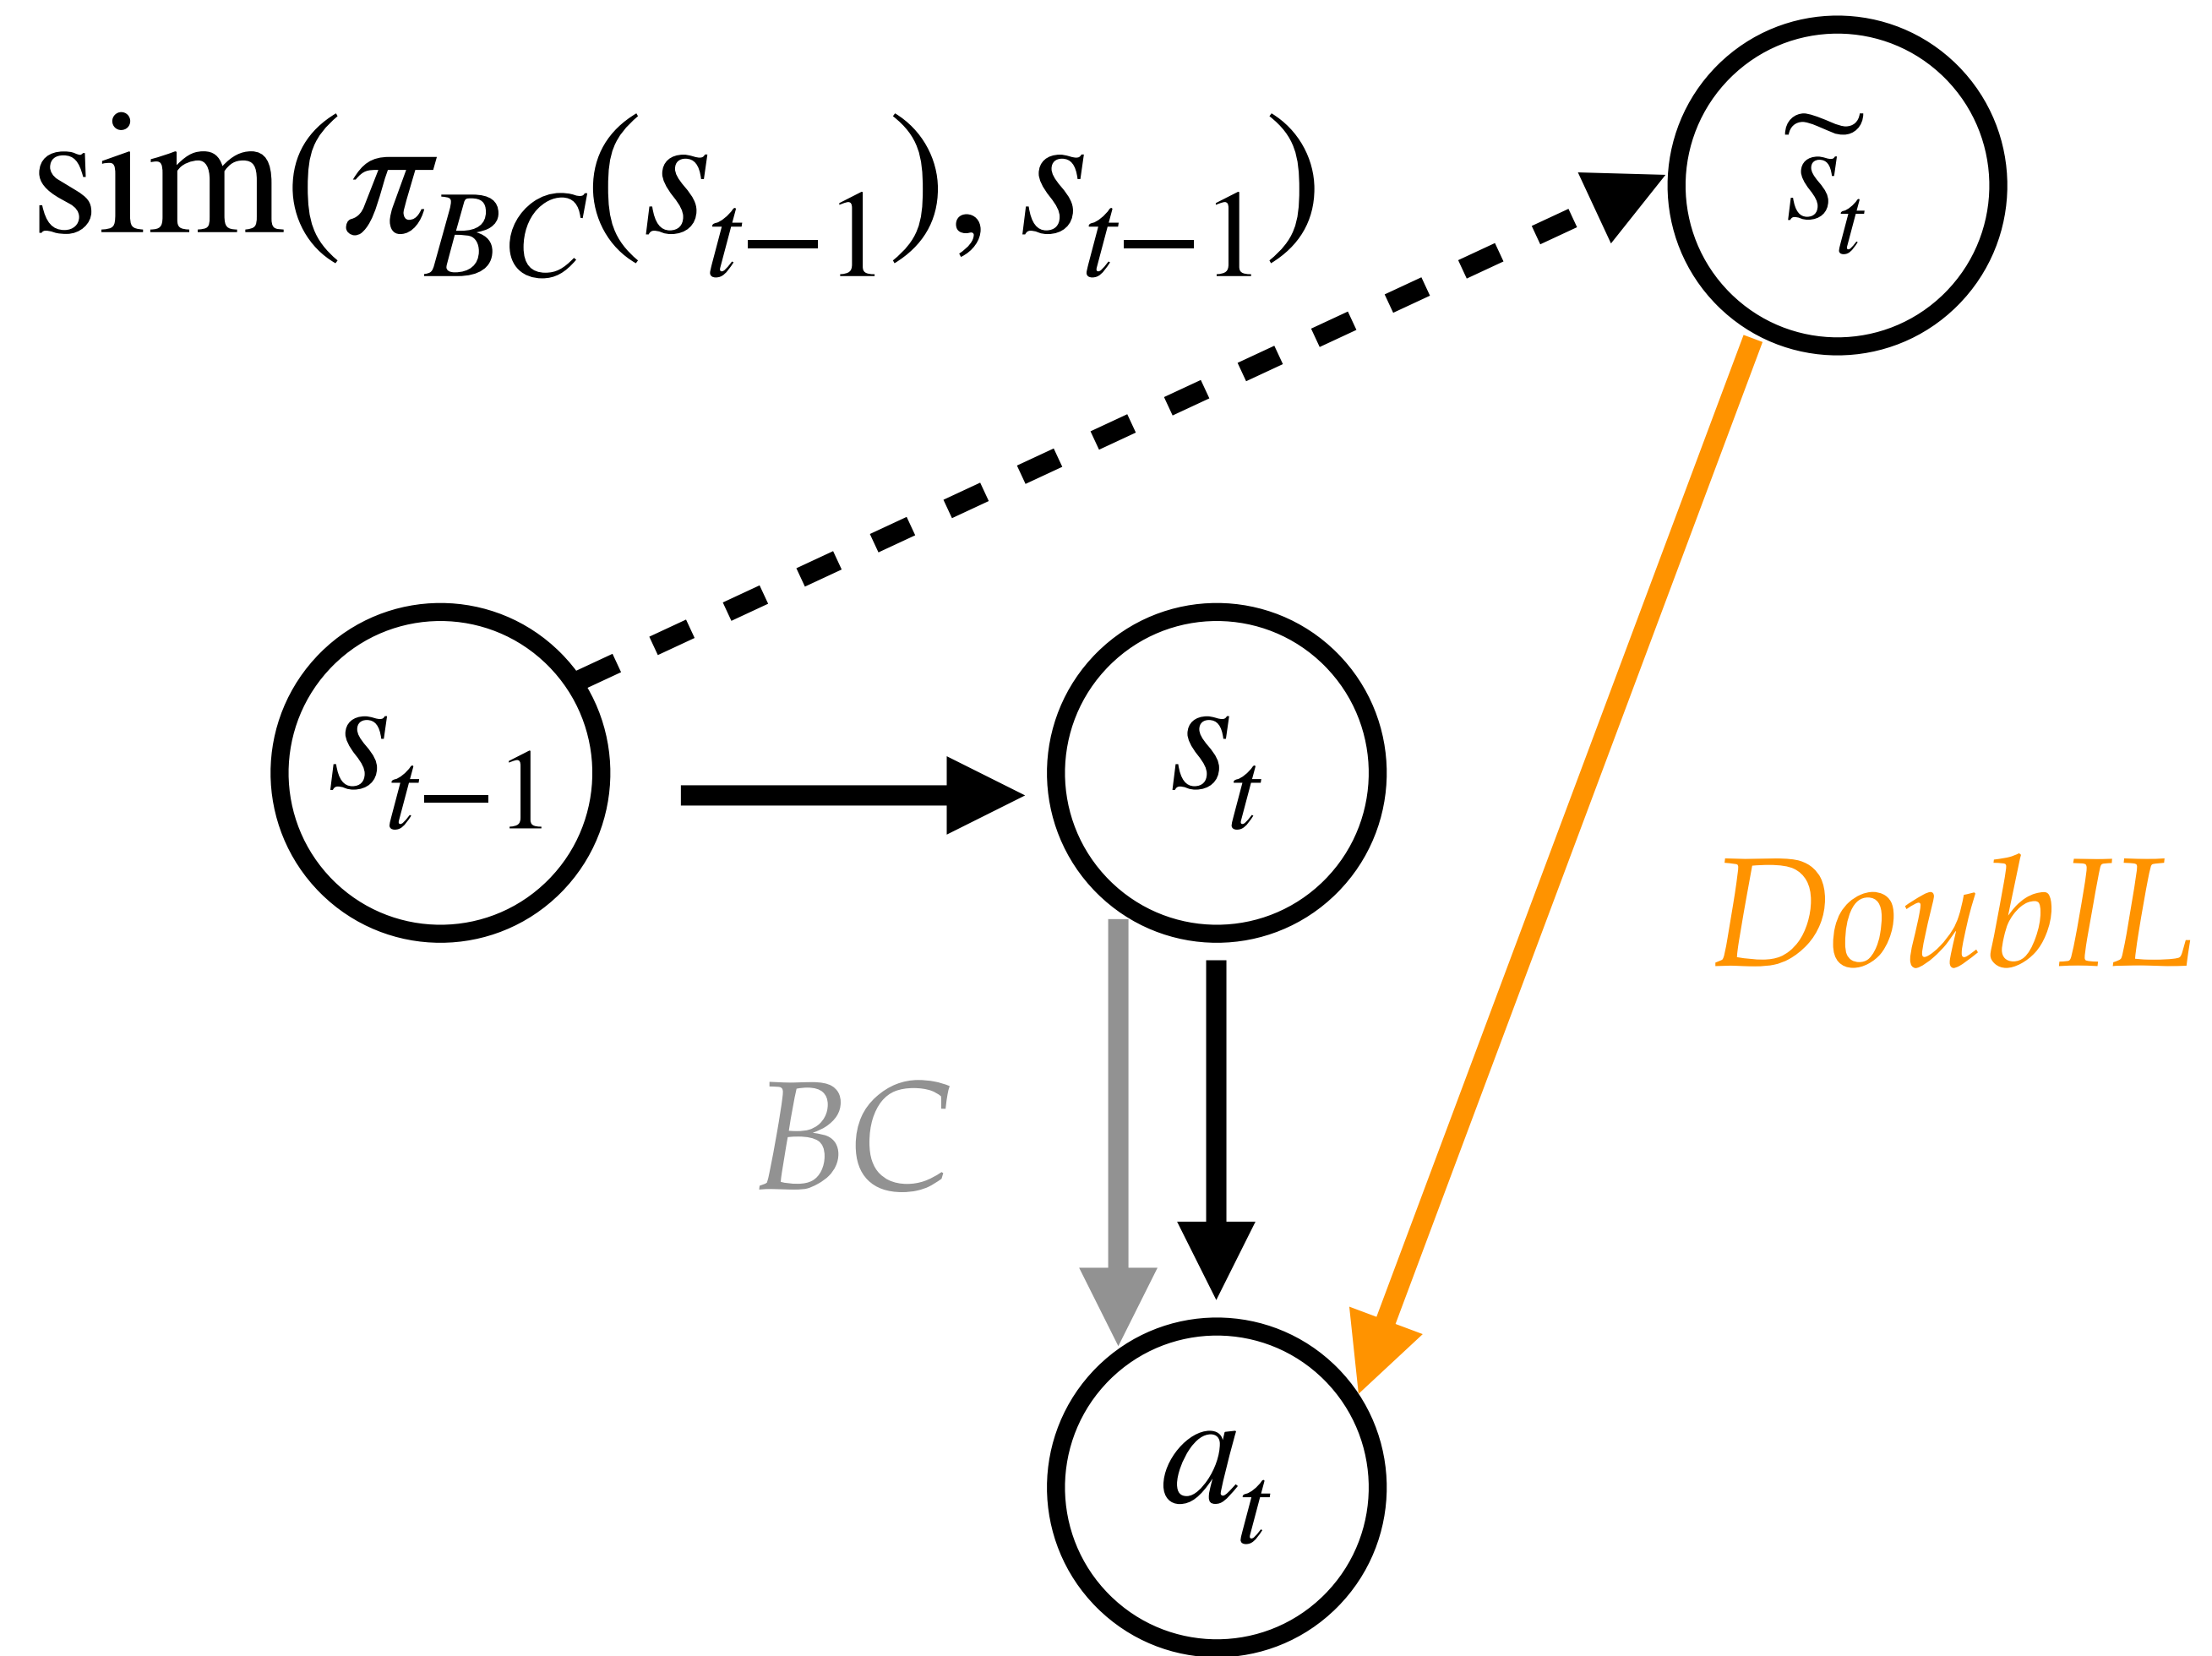












$$J(\pi_E) - J(\pi) \leq c(\sqrt{\epsilon} + \sqrt{\delta})\kappa(\Pi)T^2$$



$$\min_{\pi} \max_f \mathbb{E}[2(a_t - \pi(s_t))f(s_{t-1}) - f(s_{t-1})^2]$$

$$\min_{\pi} \max_f \mathbb{E}[2(a_t - \pi(s_t))f(s_{t-1}) - f(s_{t-1})^2]$$


$$J(\pi_E) - J(\pi) \leq c\sqrt{\epsilon}\kappa(\Pi)T^2$$

# Instrumental Variable Imitation Learning

*generative modeling*

*game-theoretic*

## DoubIL



$$J(\pi_E) - J(\pi) \leq c(\sqrt{\epsilon} + \sqrt{\delta})\kappa(\Pi)T^2$$

## ResiduIL


$$\min_{\pi} \max_f \mathbb{E}[2(a_t - \pi(s_t))f(s_{t-1}) - f(s_{t-1})^2]$$

$$J(\pi_E) - J(\pi) \leq c\sqrt{\epsilon}\kappa(\Pi)T^2$$


$$\pi_E \xleftrightarrow{f} \pi$$

[Swamy et al, 2021]


Offline



Online



Interactive




*Inconsistent,*  
*IVR Consistent*

*Consistent*

$$\pi_E \xleftrightarrow{f} \pi$$


[Swamy et al, 2021]

Offline




*Inconsistent,*  
*IVR Consistent*

Online







*Inconsistent,*  
*Hybrid?*


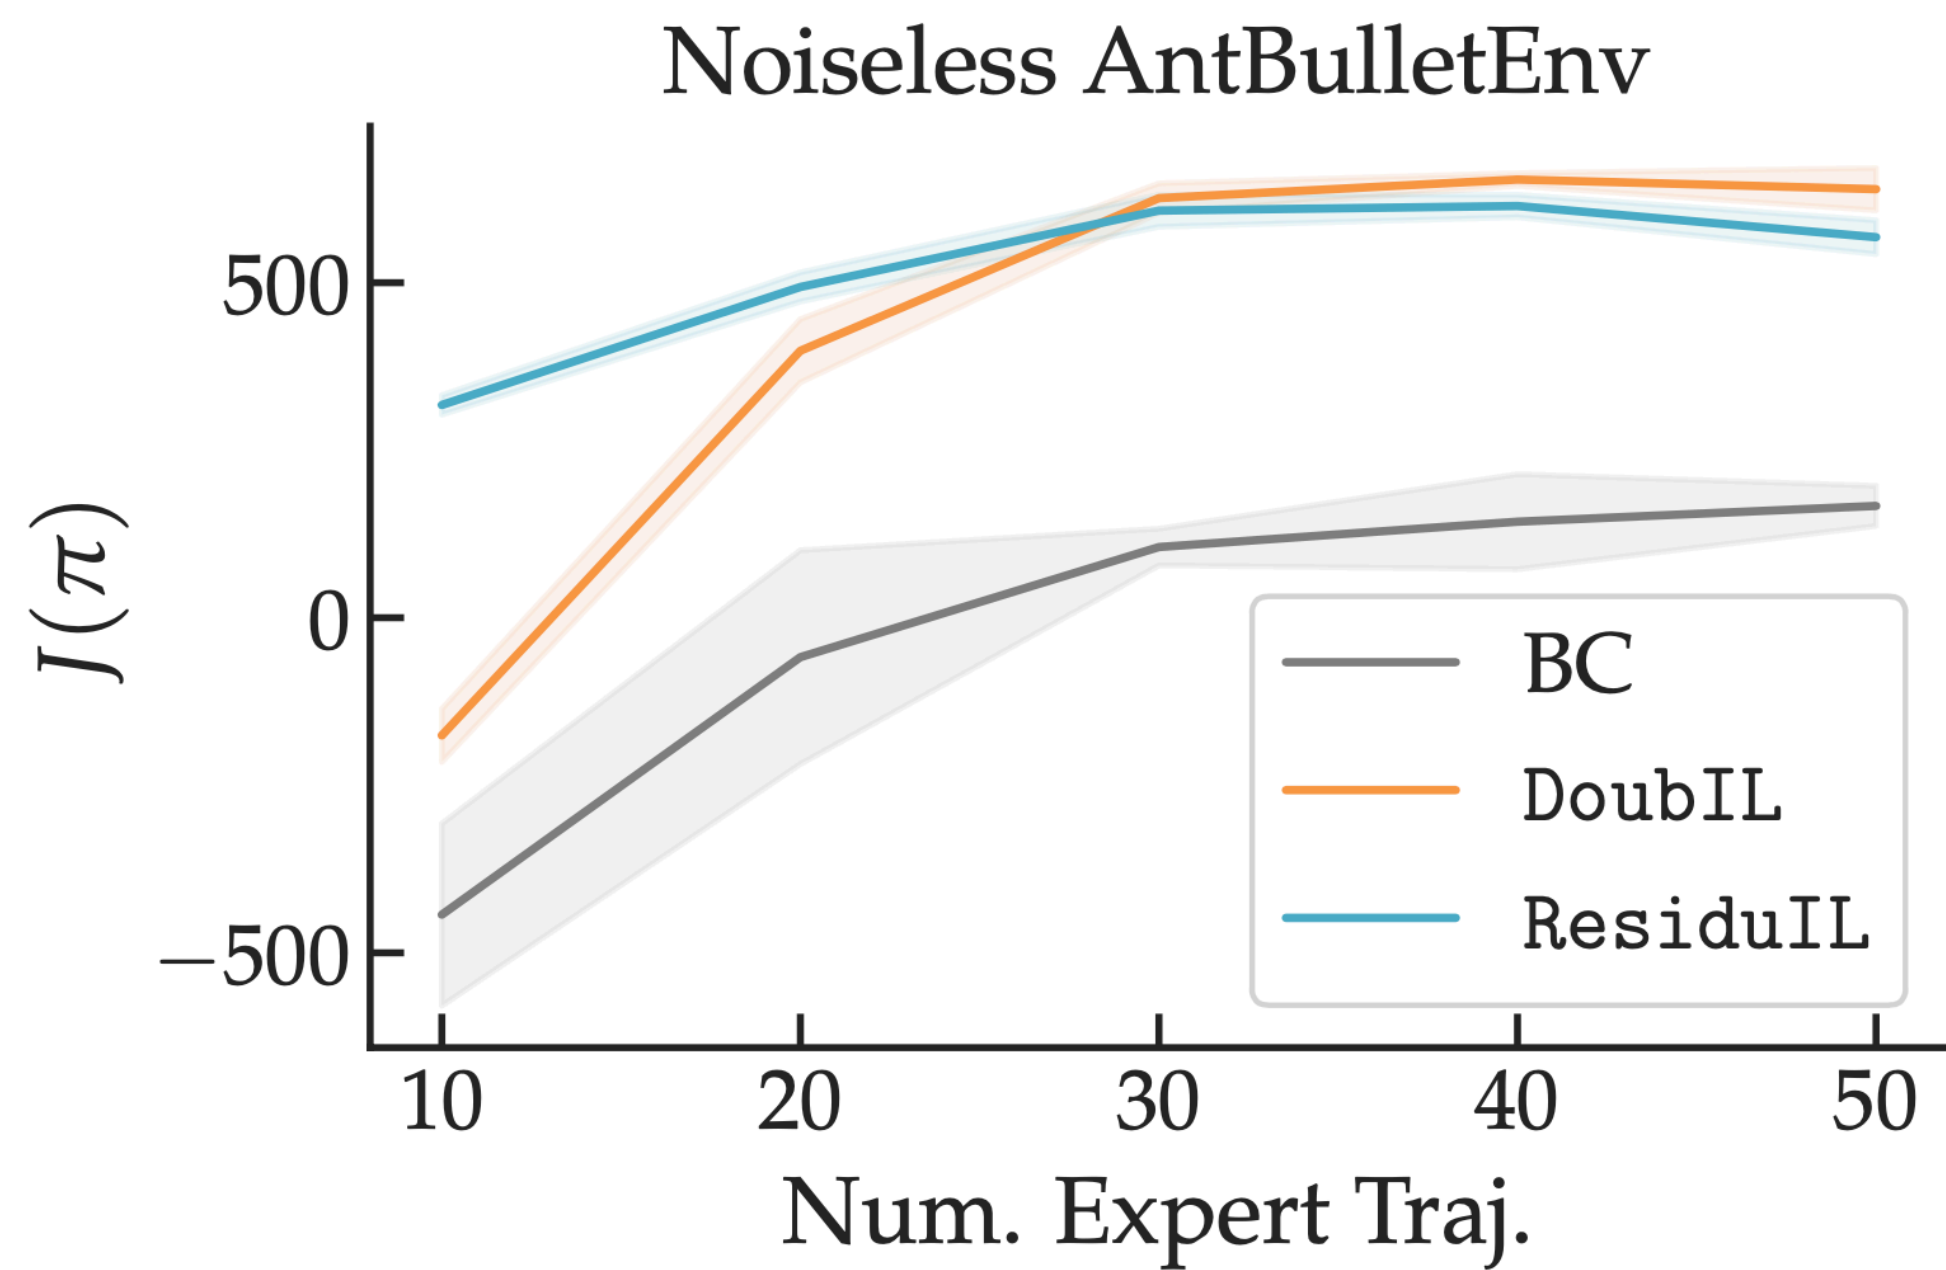


Interactive



*Consistent*

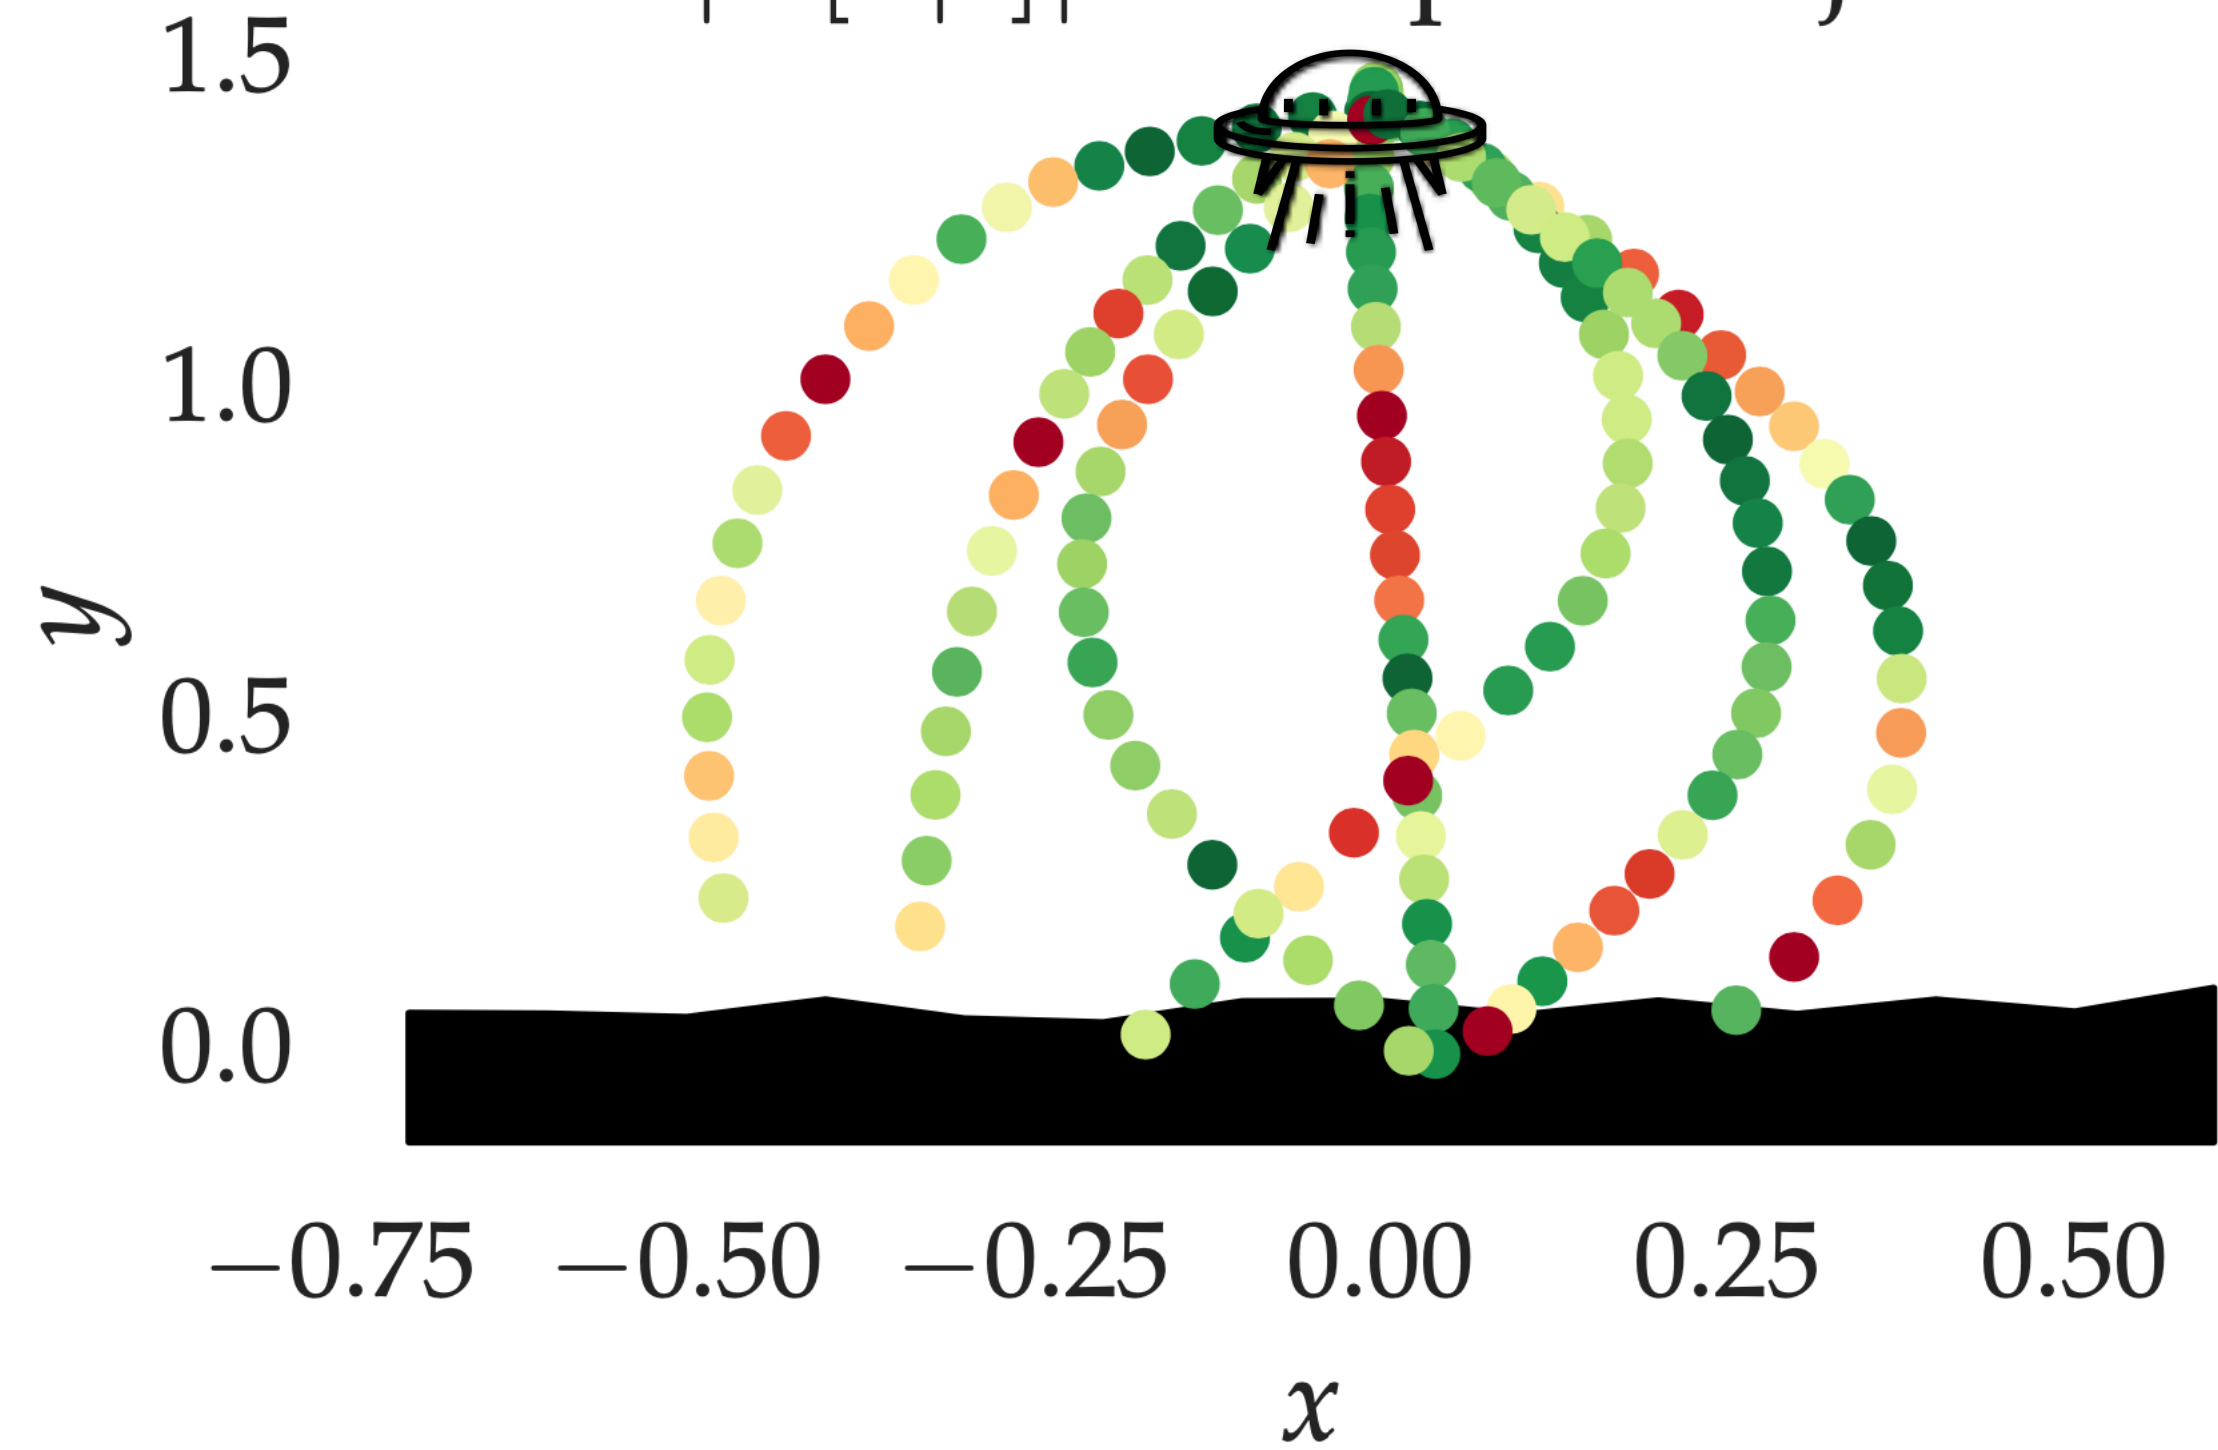


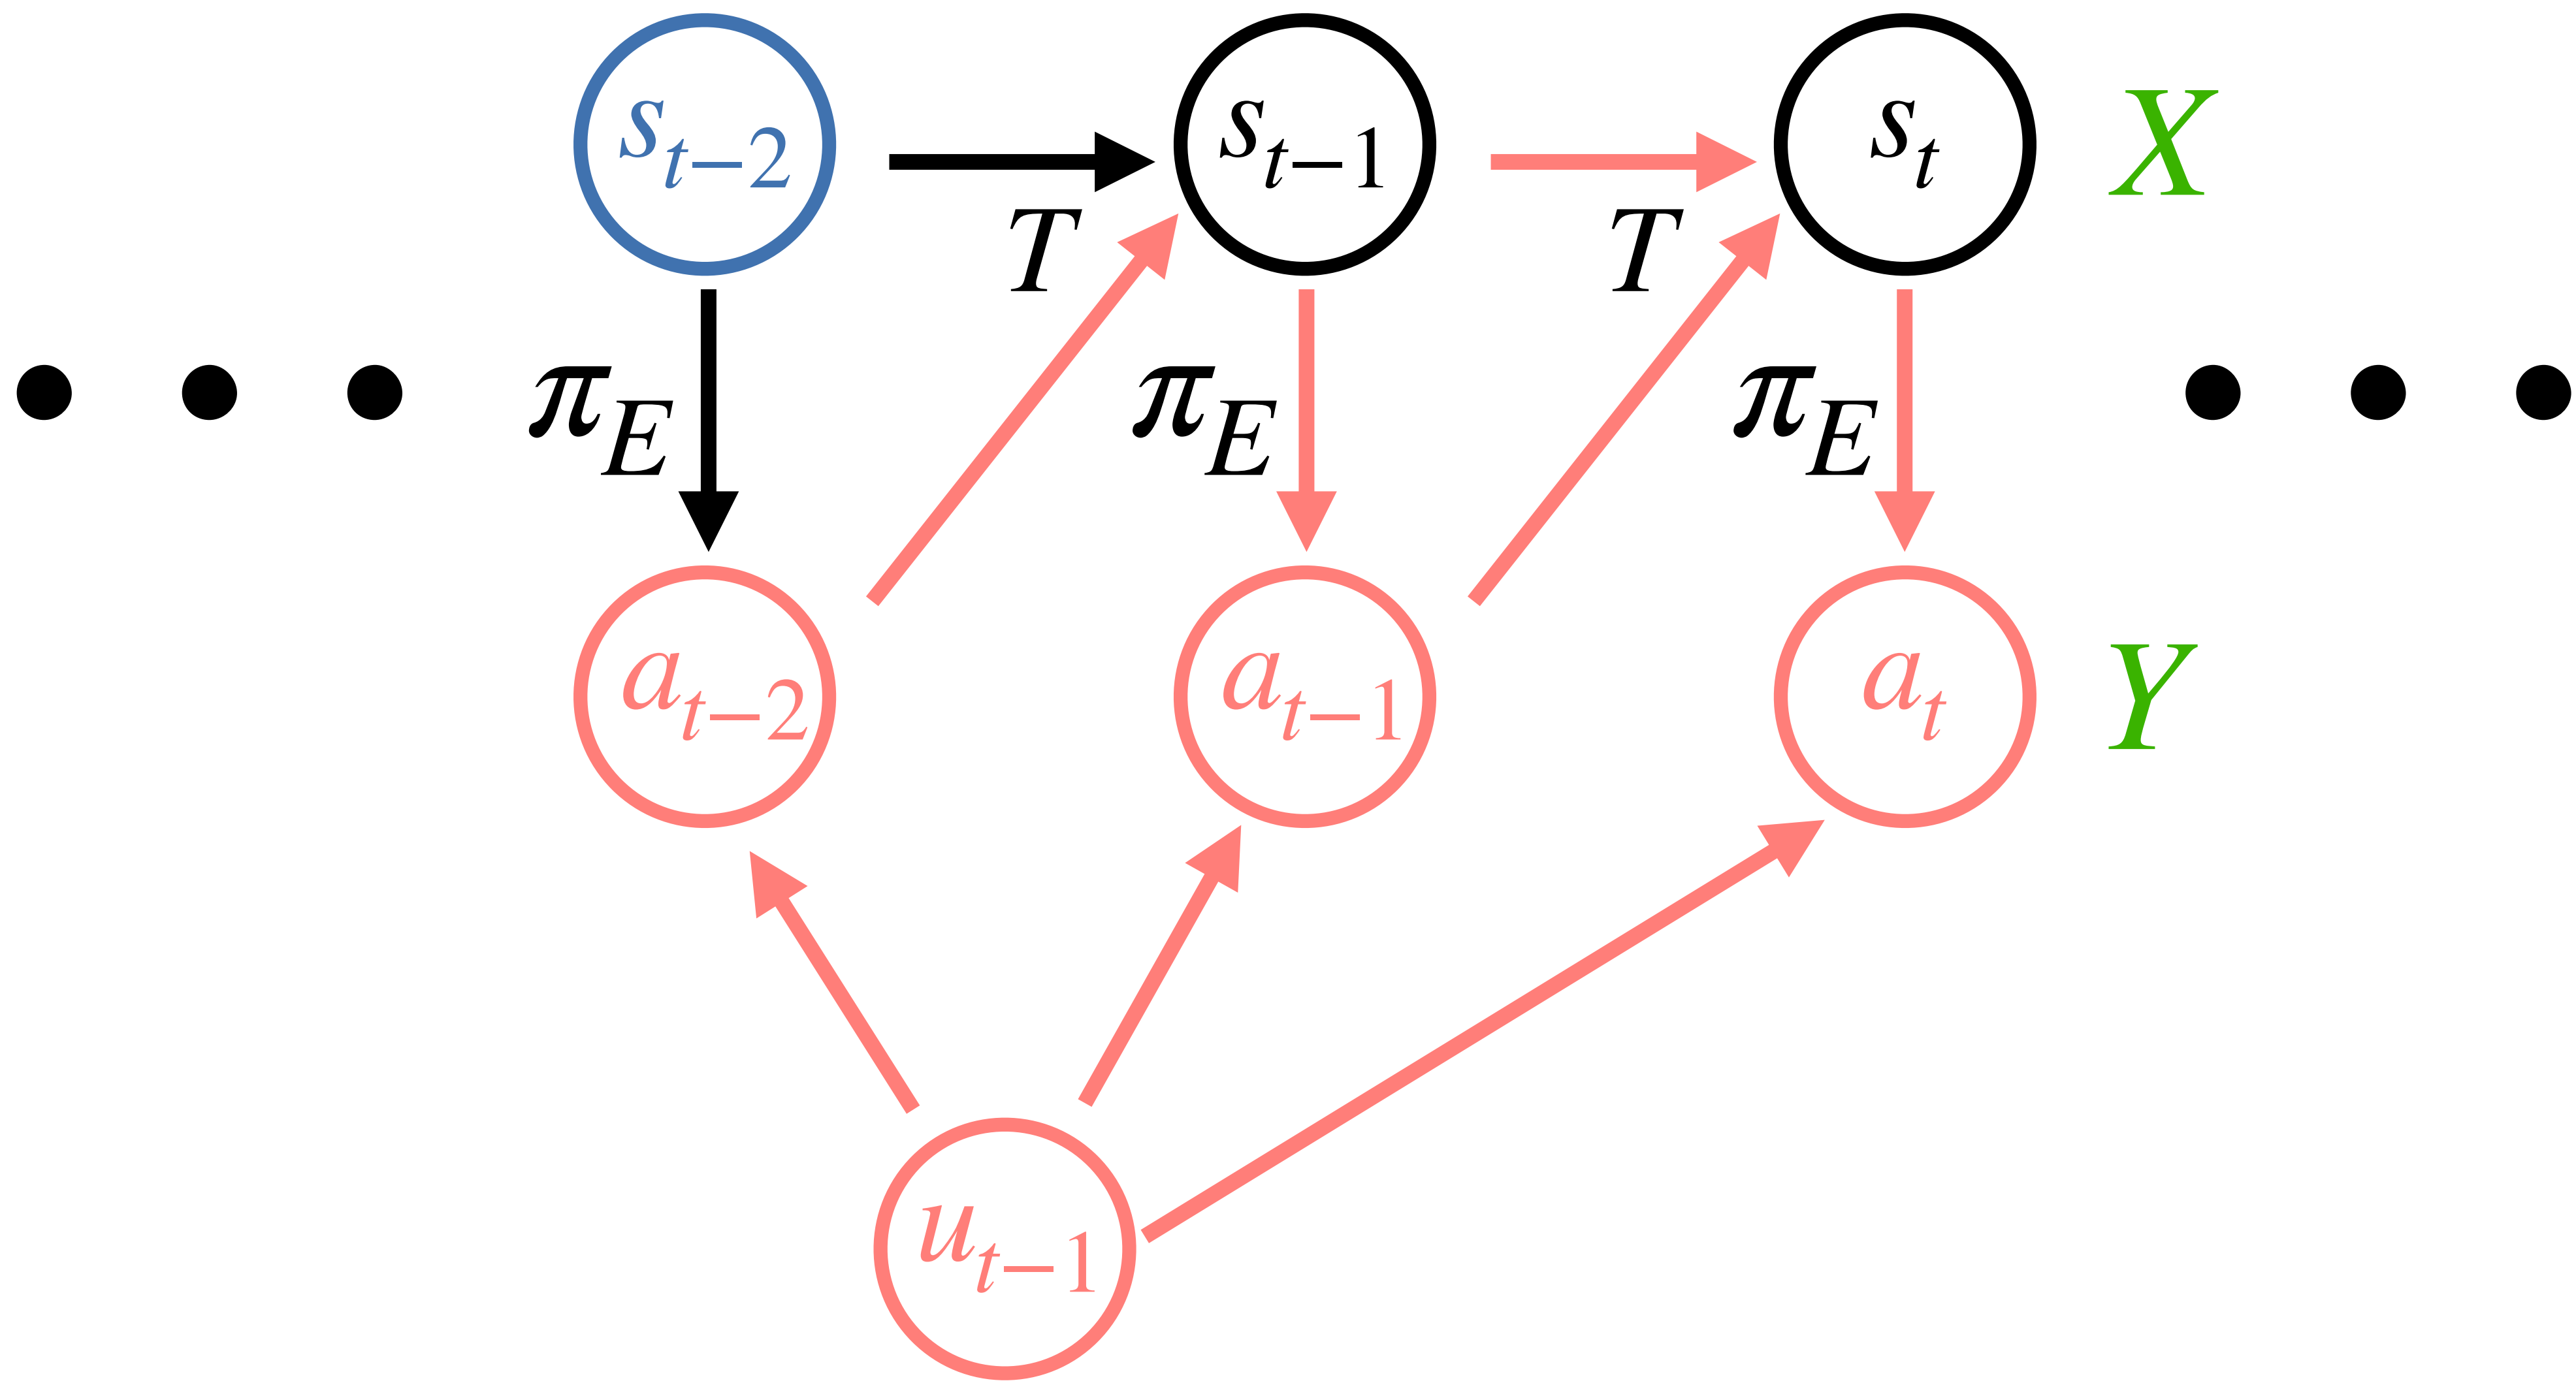


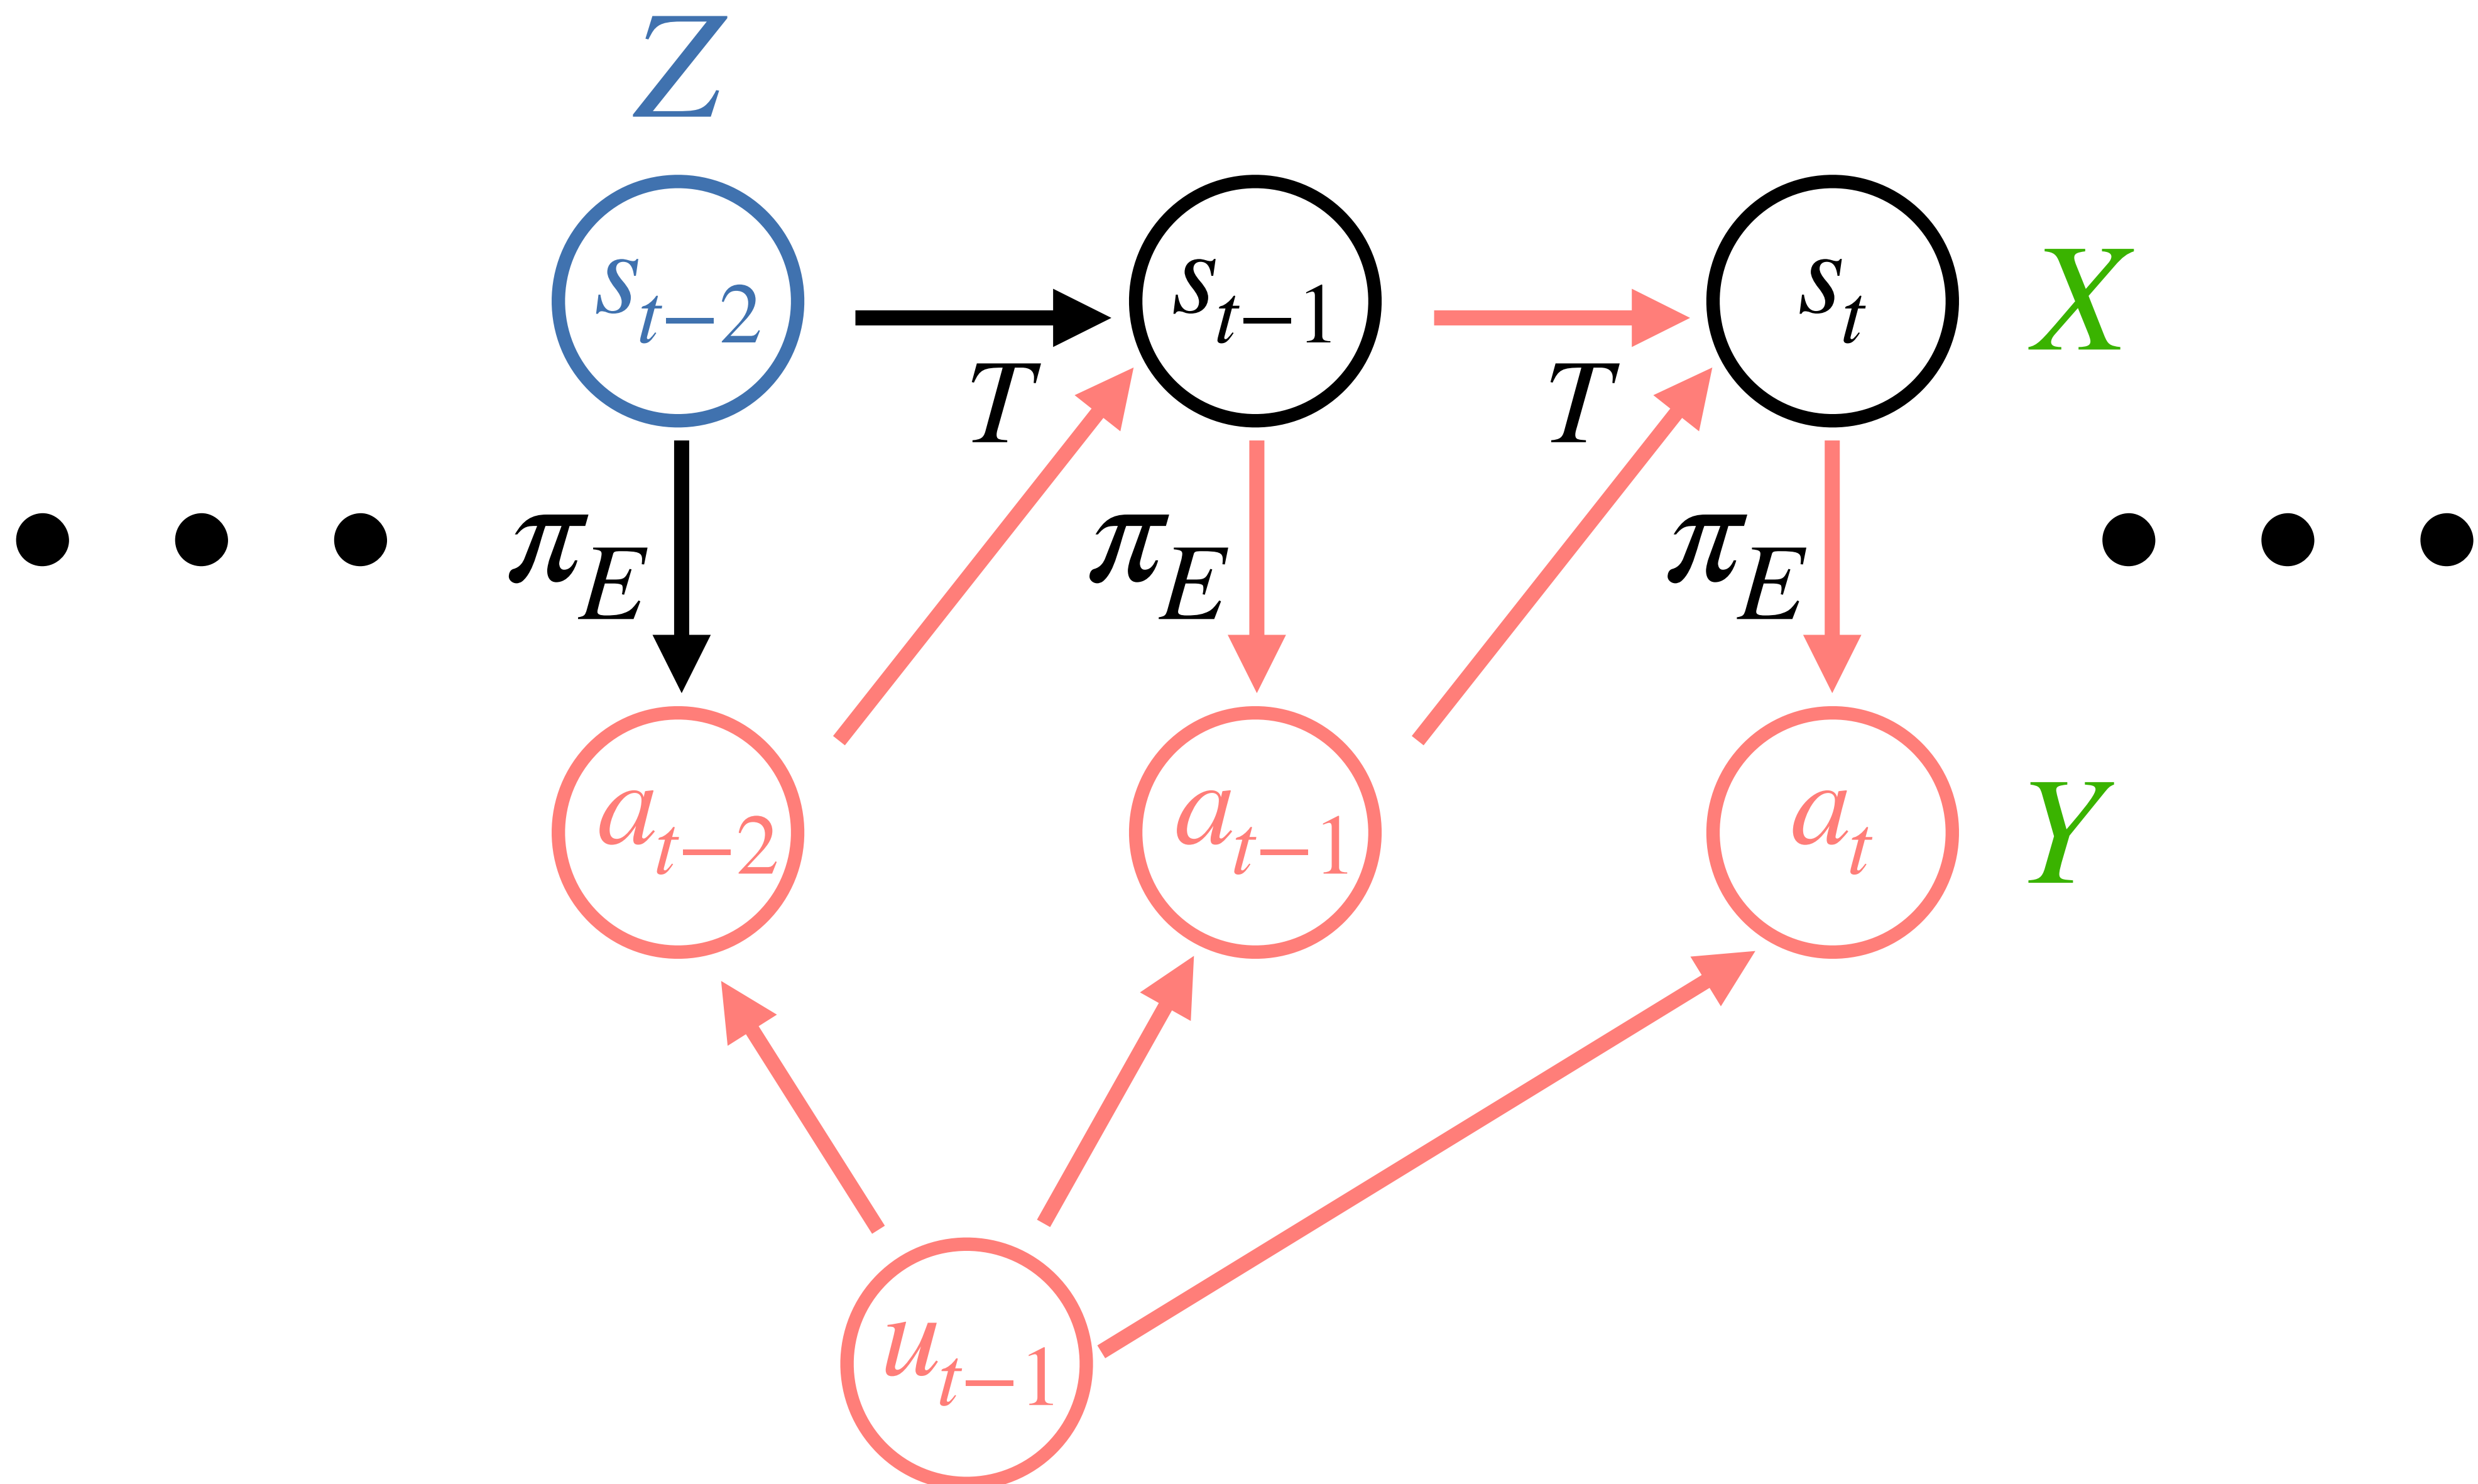




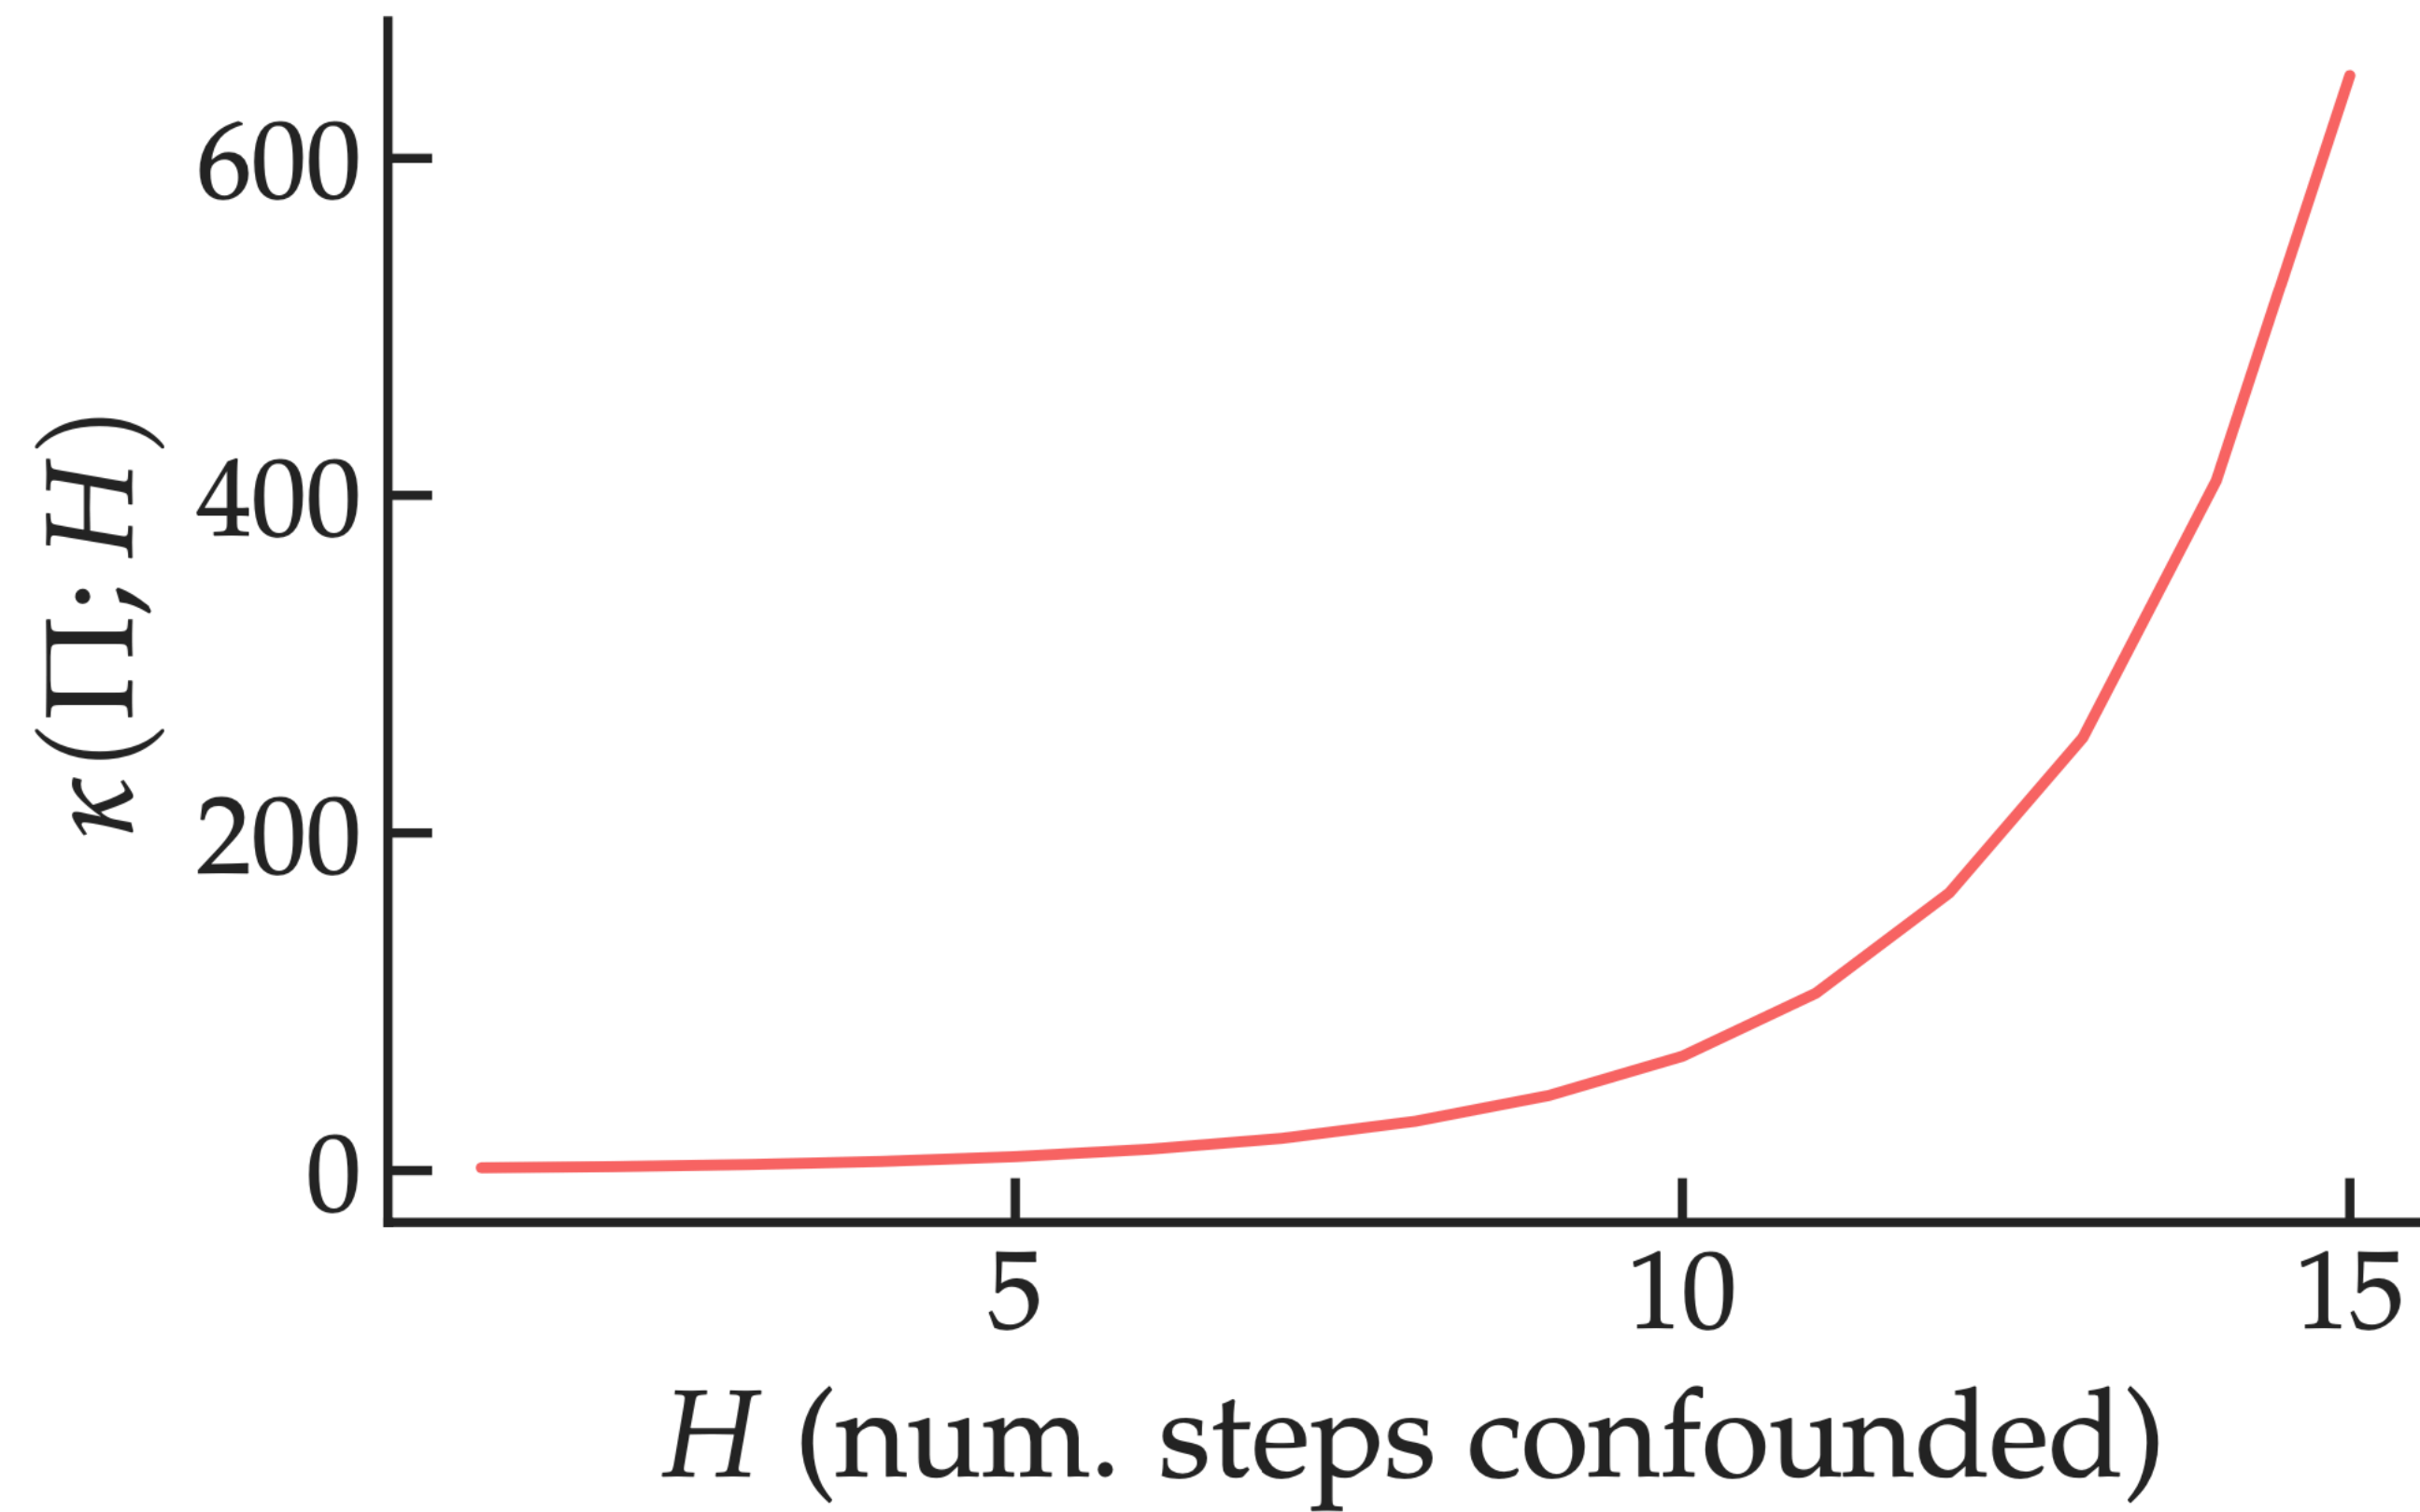
$|\mathbb{E}[u|s]|$  on Expert Trajs.













# Thanks!

<https://gokul.dev/causal/>  
[gswamy@cmu.edu](mailto:gswamy@cmu.edu)

