Protocols for Verifying Smooth Strategies in Bandits and Games

Jonathan Shafer
MIT

November 2025

Miranda Christ

Daniel Reichman

Verification in Games

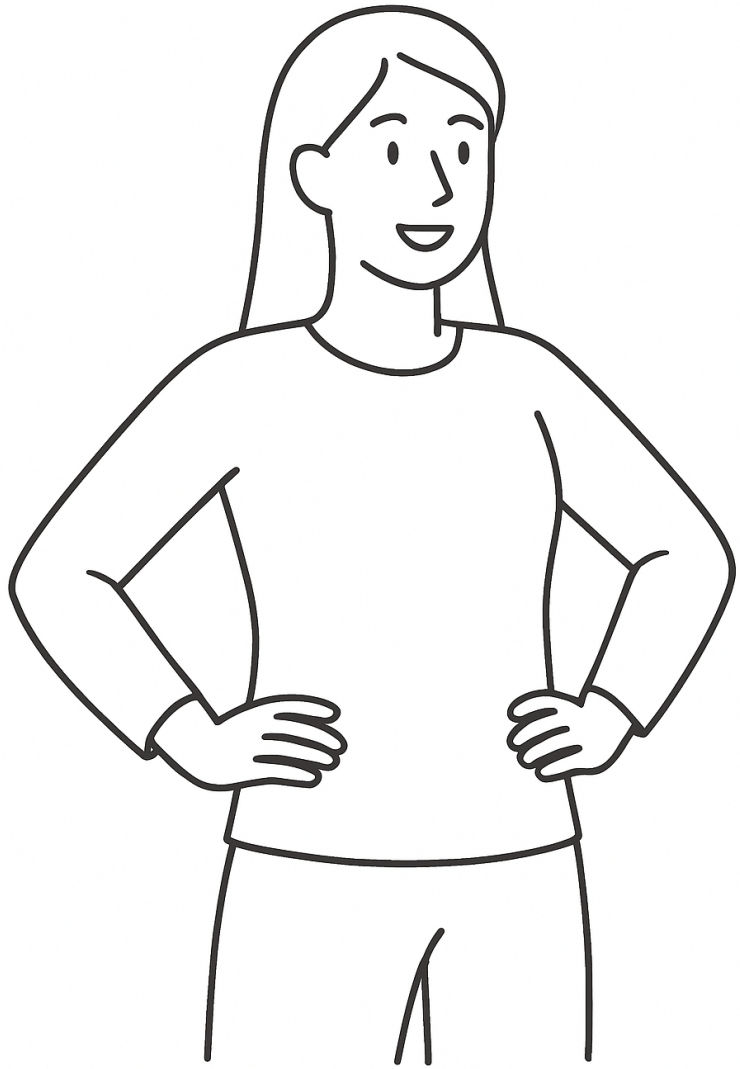

Alice

Game

Normal Form Game

k players

Normal Form Game

k players

n possible actions per player

Normal Form Game

k players

n possible actions per player, mixed strategies

Normal Form Game

k players

n possible actions per player, mixed strategies

Payoff function $U: \ [n]^k \to [0,1]^k$

Normal Form Game

k players

n possible actions per player, mixed strategies

Payoff function $U: [n]^k \to [0,1]^k$

unknown, players have query access

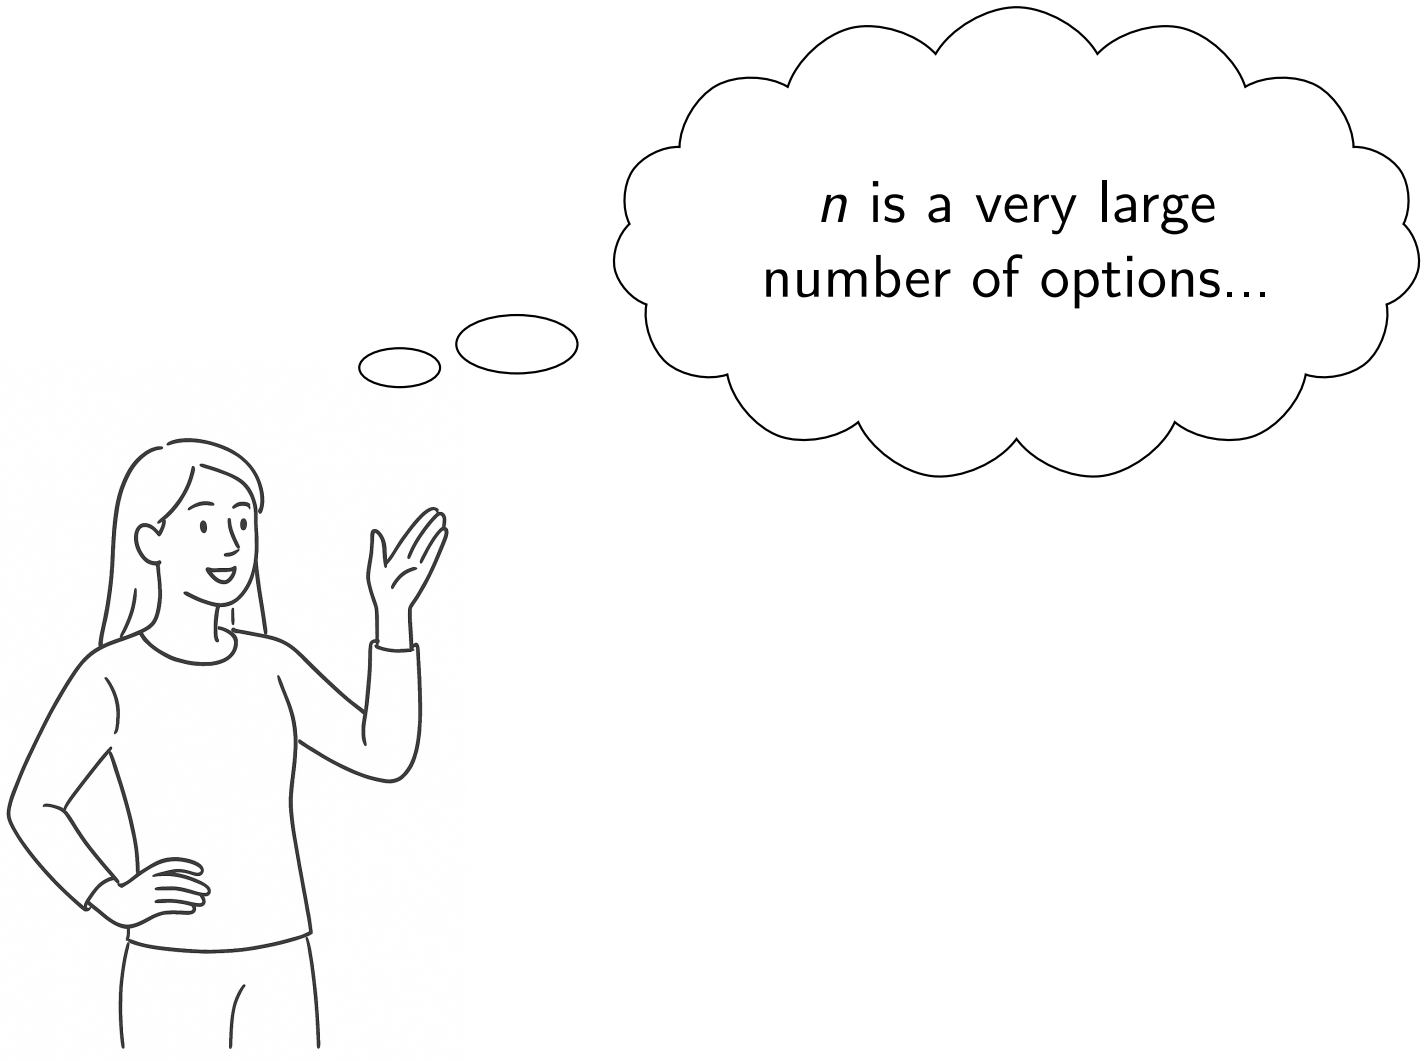
n is a very large
number of options...

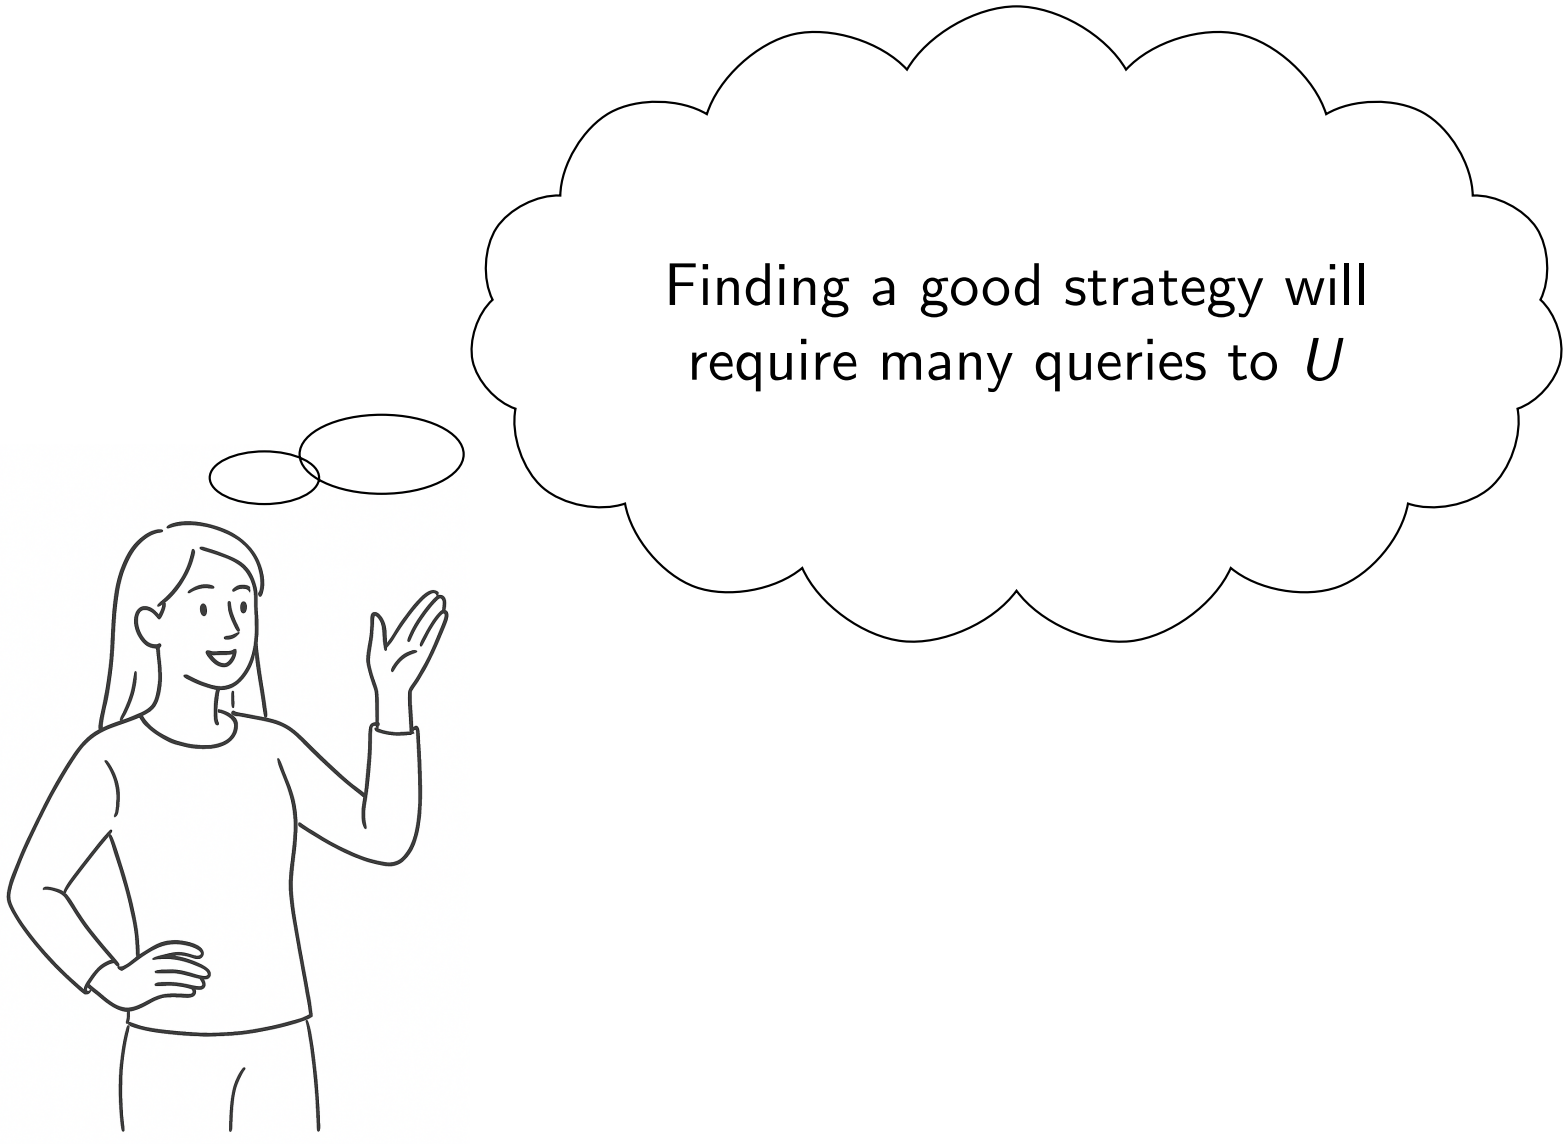
Finding a good strategy will require many queries to U

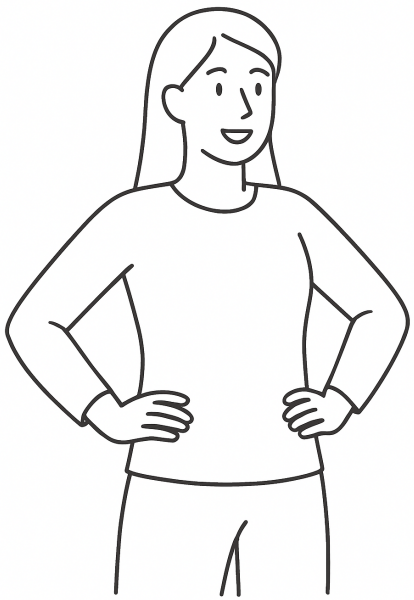

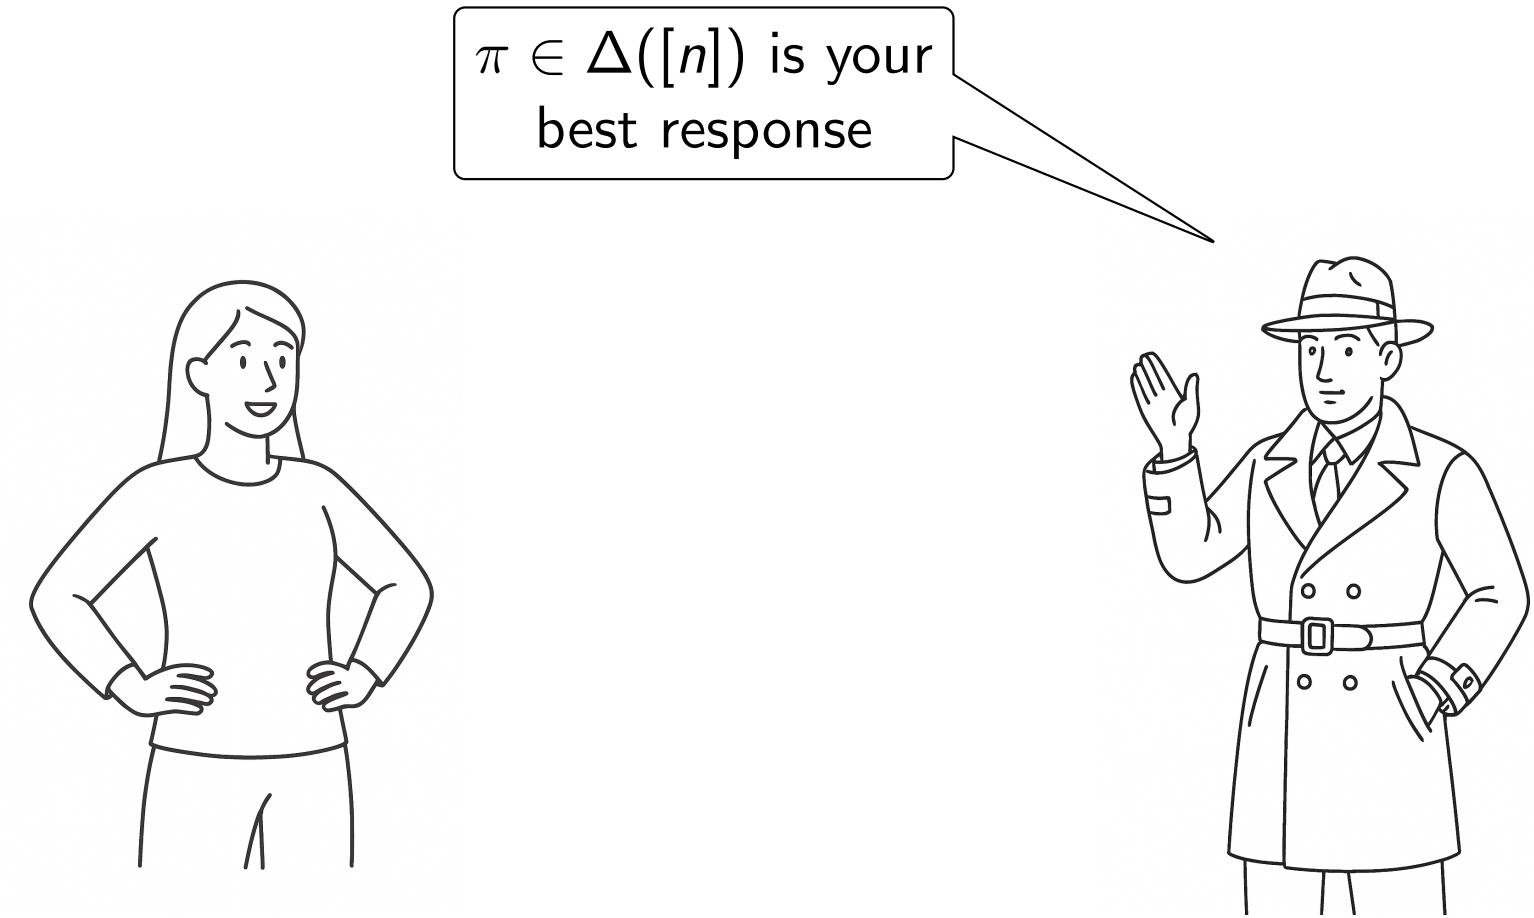
$\pi \in \Delta([n])$ is your best response

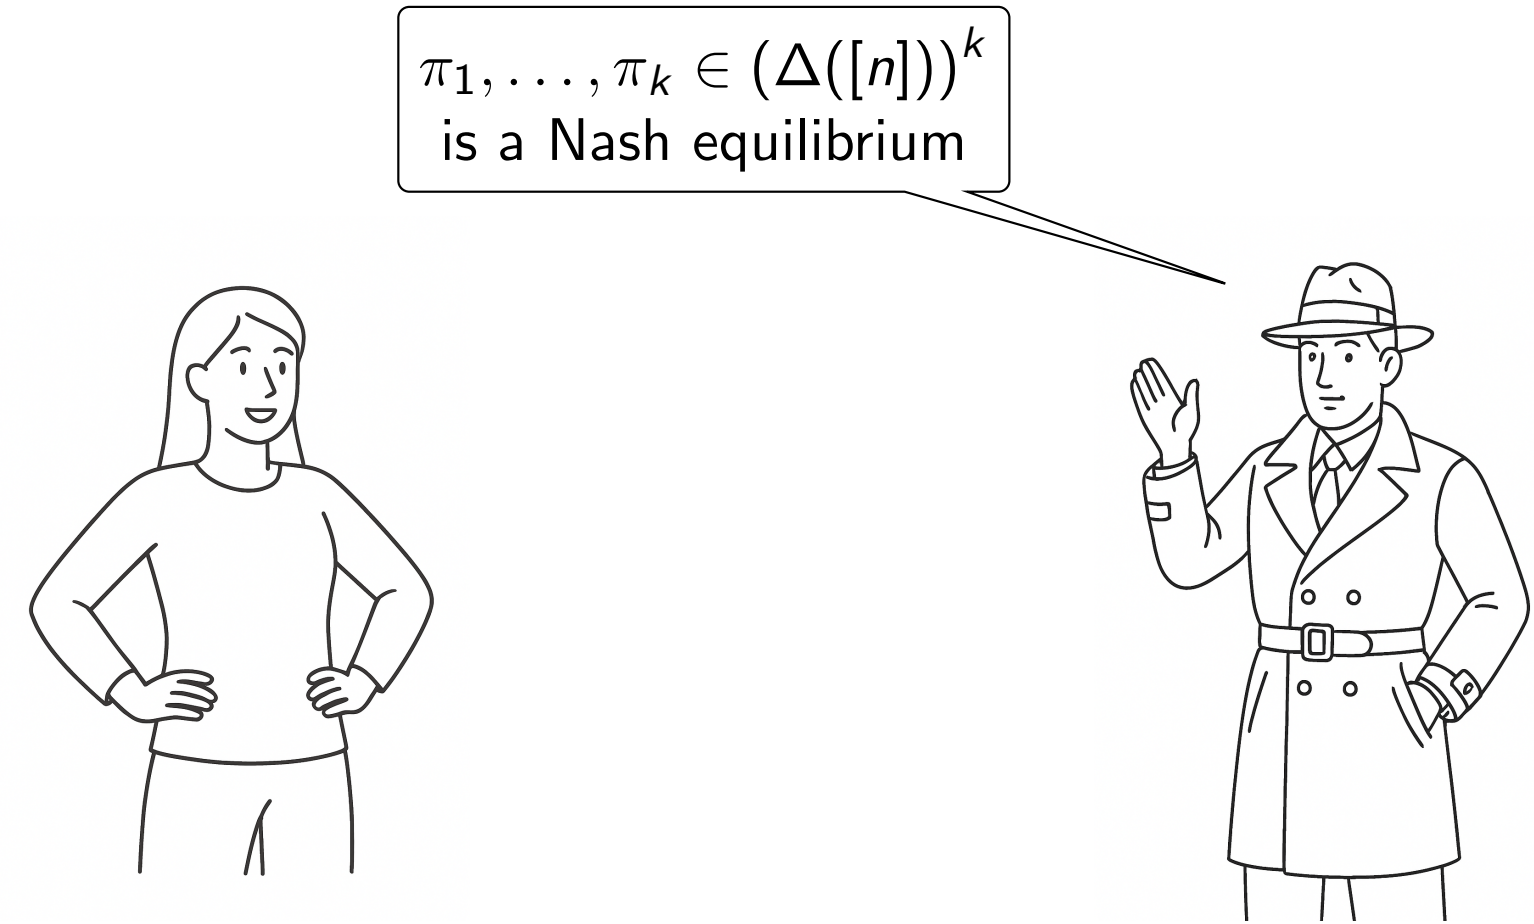
$\pi_1, \ldots, \pi_k \in (\Delta([n]))^k$
is a Nash equilibrium

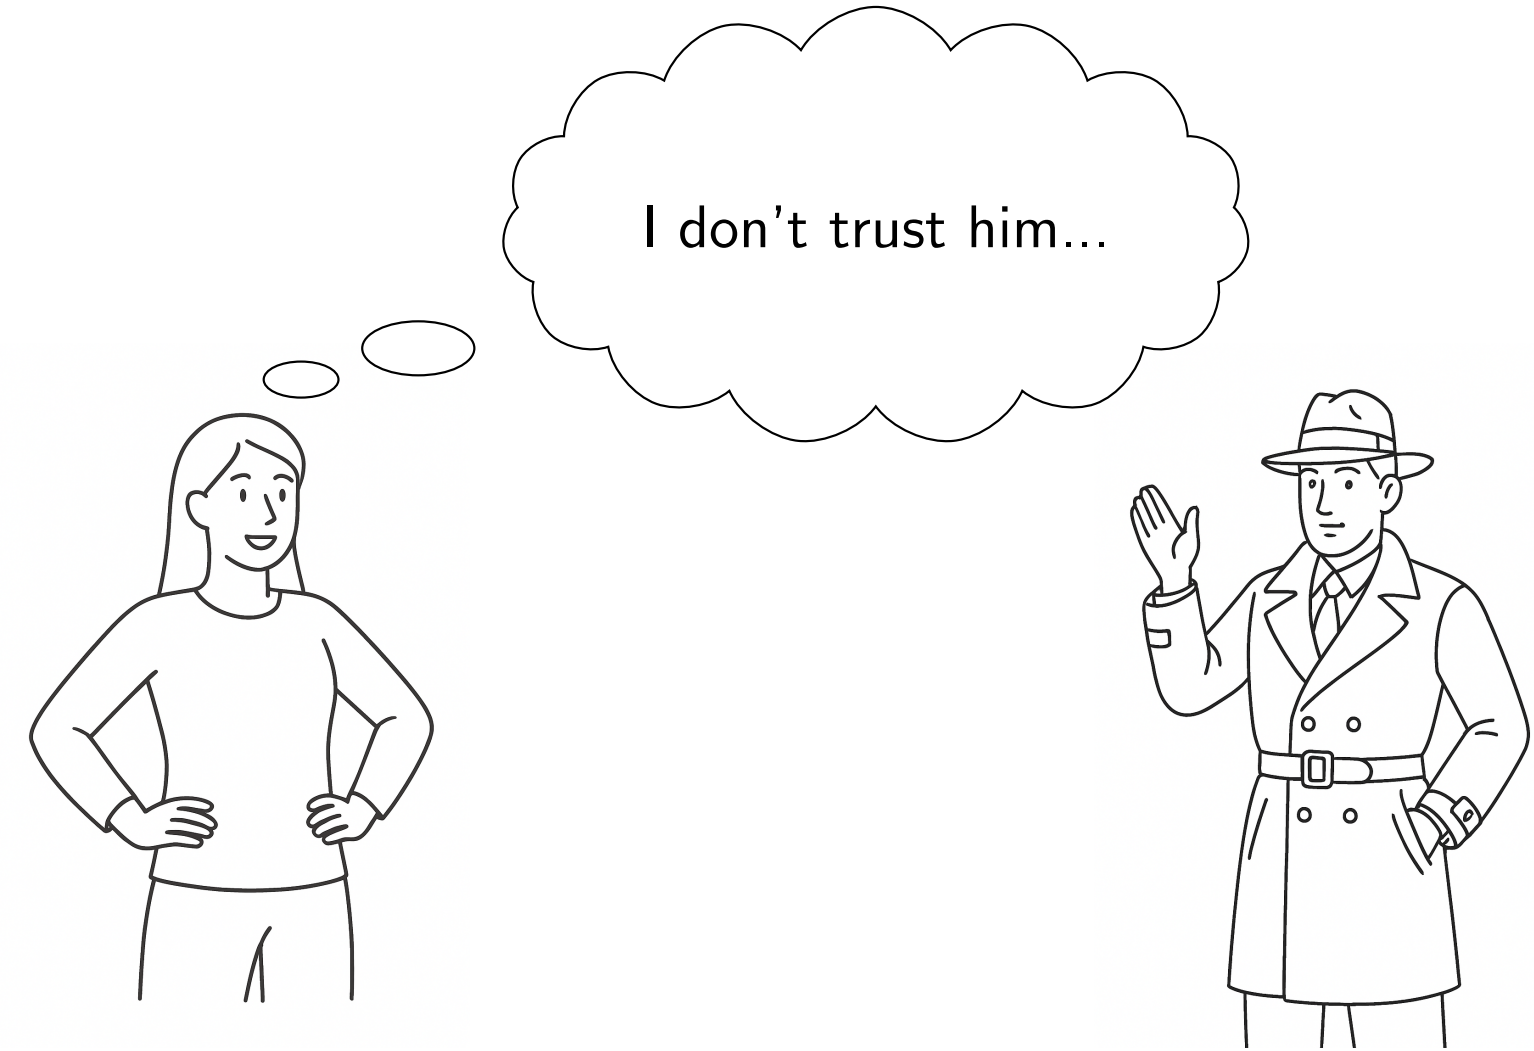
I don't trust him...

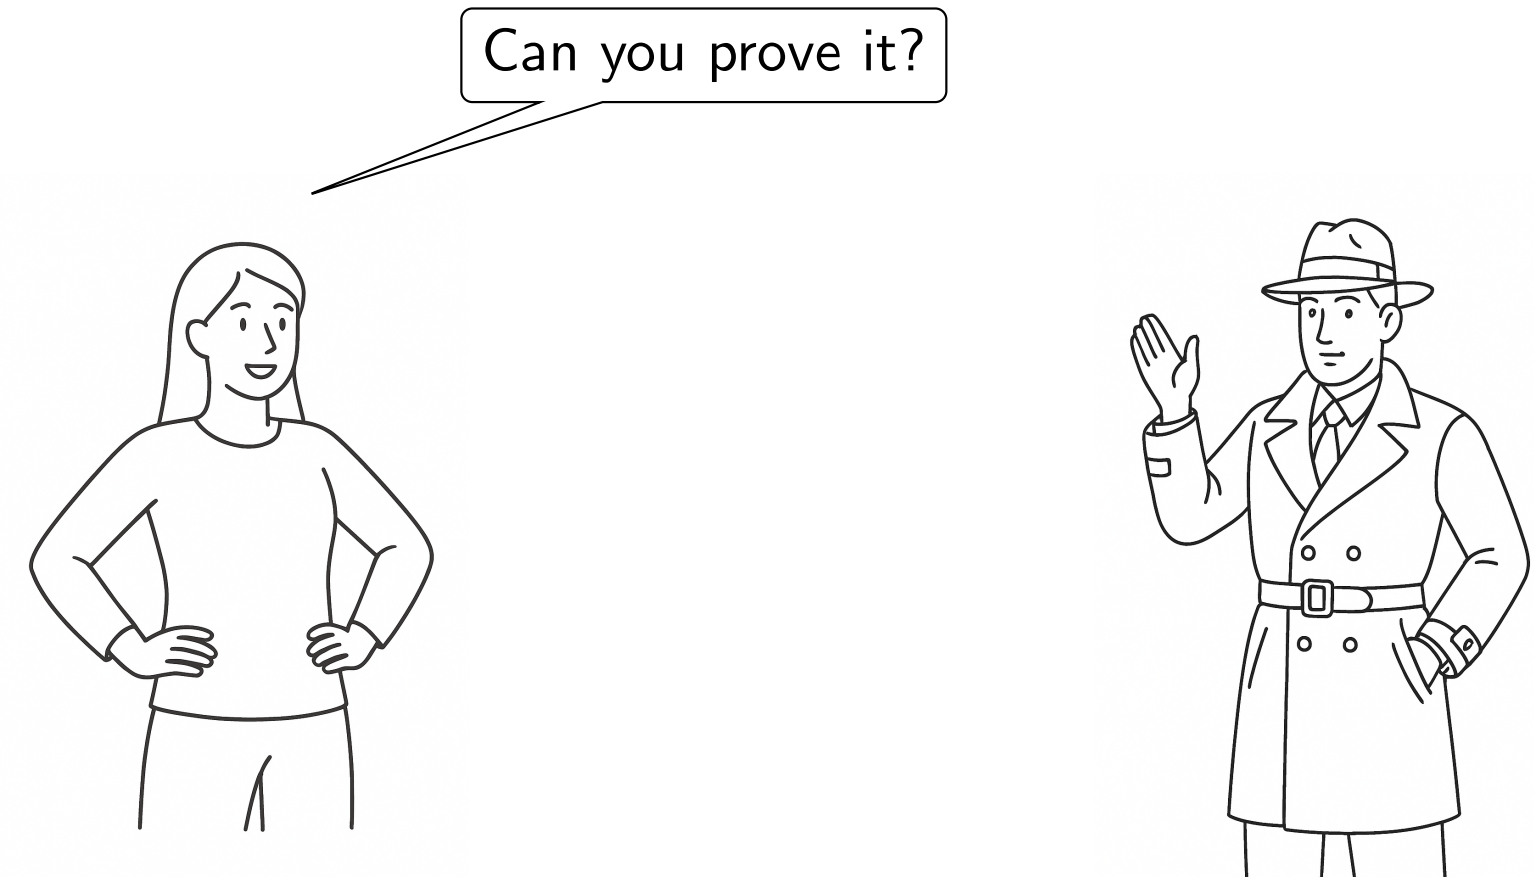
Can you prove it?

Q: Can verifying untrusted advice be cheaper than learning without advice?

Smoothness is Necessary

NO: In general, efficient verification protocol not possible

Smoothness is Necessary

NO: In general, efficient verification protocol not possible

To ensure that no better action exists,

Alice must query U on all options

Smoothness is Necessary

NO: In general, efficient verification protocol not possible

To ensure that no better action exists,

Alice must query U on all options

Def. A strategy $\pi \in \Delta([n])$ is σ-smooth if $\pi(i) \leq \sigma$ for all $i \in [n]$

Smoothness is Necessary

NO: In general, efficient verification protocol not possible

To ensure that no better action exists,

Alice must query U on all options

Def. A strategy $\pi \in \Delta([n])$ is σ-smooth if $\pi(i) \leq \sigma$ for all $i \in [n]$

(σ-smooth Nash equilibrium always exists)

Smoothness is Necessary

NO: In general, efficient verification protocol not possible

To ensure that no better action exists,

Alice must query U on all options

Def. A strategy $\pi \in \Delta([n])$ is σ-smooth if $\pi(i) \leq \sigma$ for all $i \in [n]$

(σ-smooth Nash equilibrium always exists)

YES: Efficiently verifying smooth strategies is possible!

Results

Theorem. $\exists$ protocol for verifying

ε-approximate σ-smooth Nash equilibria

where verifier uses $\tilde{O}\left(\frac{kn\sigma}{\varepsilon^2}\right)$ queries to U

Results

Theorem. $\exists$ protocol for verifying

ε-approximate σ-smooth Nash equilibria

where verifier uses $\tilde{O}\left(\frac{kn\sigma}{\varepsilon^2}\right)$ queries to U

Example:

Results

Theorem. $\exists$ protocol for verifying

ε-approximate σ-smooth Nash equilibria

where verifier uses $\tilde{O}\left(\frac{kn\sigma}{\varepsilon^2}\right)$ queries to U

Example: $\sigma = \frac{1}{\sqrt{n}}$

Results

Theorem. $\exists$ protocol for verifying

ε-approximate σ-smooth Nash equilibria

where verifier uses $\tilde{O}\left(\frac{k \cdot \sqrt{n}}{\varepsilon^2}\right)$ queries to U

Example: $\sigma = \frac{1}{\sqrt{n}}$

Results

Theorem. $\exists$ protocol for verifying

ε-approximate σ-smooth Nash equilibria

where verifier uses $\tilde{O}\left(\frac{k \cdot \sqrt{n}}{\varepsilon^2}\right)$ queries to U

Example: $\sigma = \frac{1}{\sqrt{n}}$

Compare — without prover: $\Omega(n)$

Results

Theorem. $\exists$ protocol for verifying

ε-approximate σ-smooth Nash equilibria

where verifier uses $\tilde{O}\left(\frac{k \cdot \sqrt{n}}{\varepsilon^2}\right)$ queries to U

Example: $\sigma = \frac{1}{\sqrt{n}}$

Compare — without prover: $\Omega(n)$, $2^{\Omega(k)}$

Results

Results for bandits

Results

Results for bandits

Lower bounds for bandits and games

Results

Results for bandits

Lower bounds for bandits and games

Using cryptography to reduce communication

Results

Results for bandits

Lower bounds for bandits and games

Using cryptography to reduce communication

...

Thank You!