


Bounded-Loss Private Prediction Markets

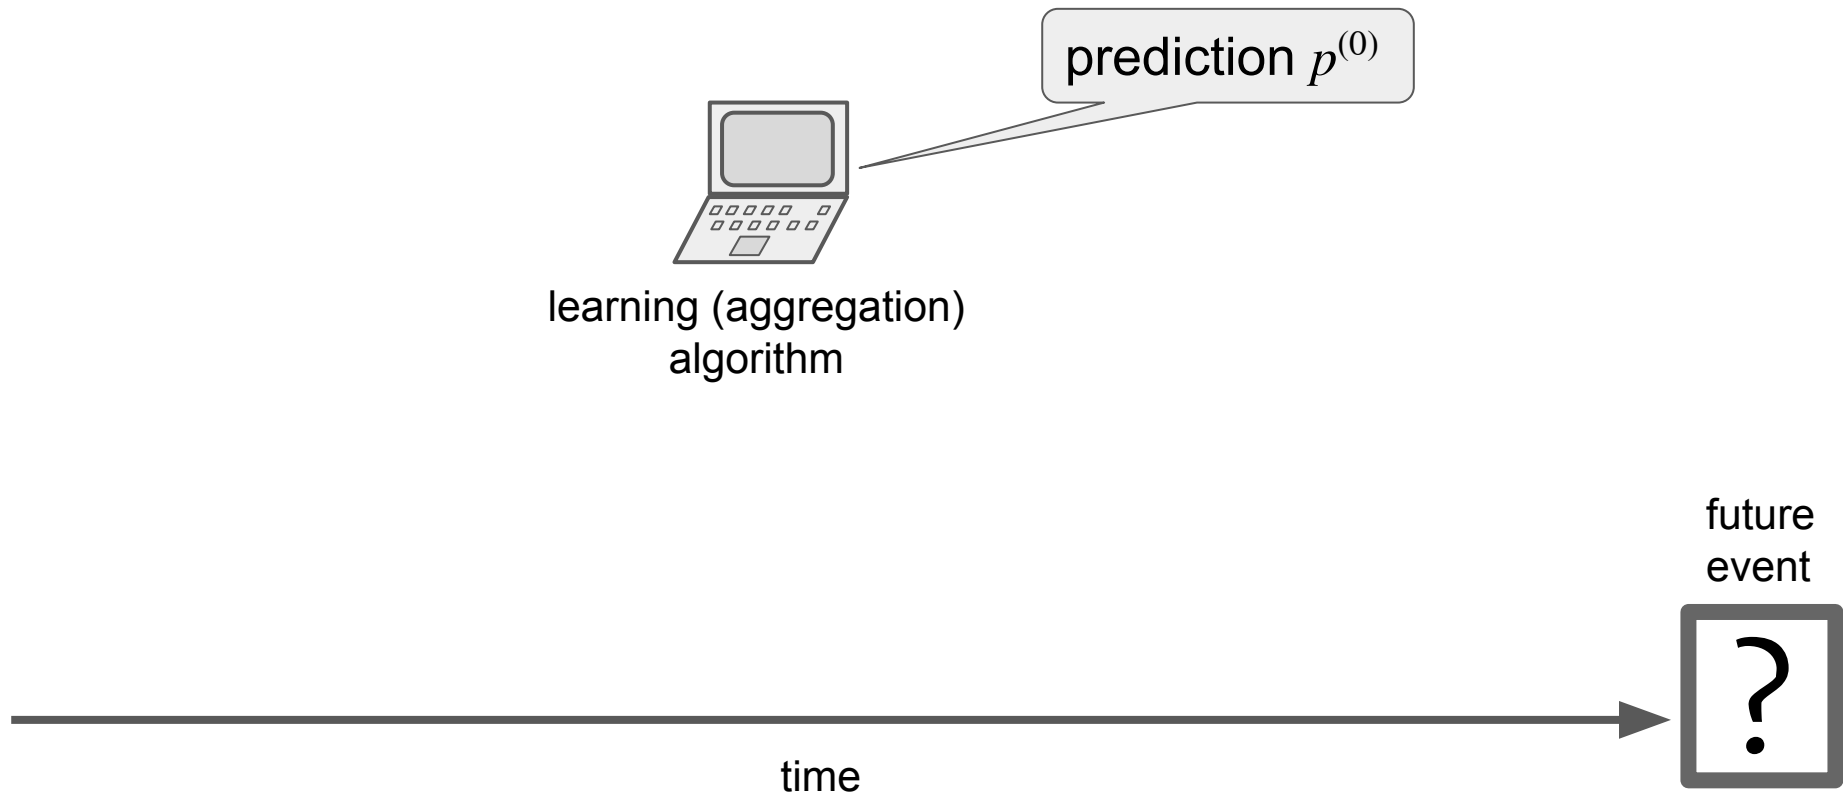


Rafael Frongillo
Bo Waggoner

University of Colorado, Boulder
Microsoft Research, NYC


*Neural Information
Processing Systems,
December 2018*

Prediction markets




Prediction markets


Menu of purchases
(gambles)



learning (aggregation)
algorithm

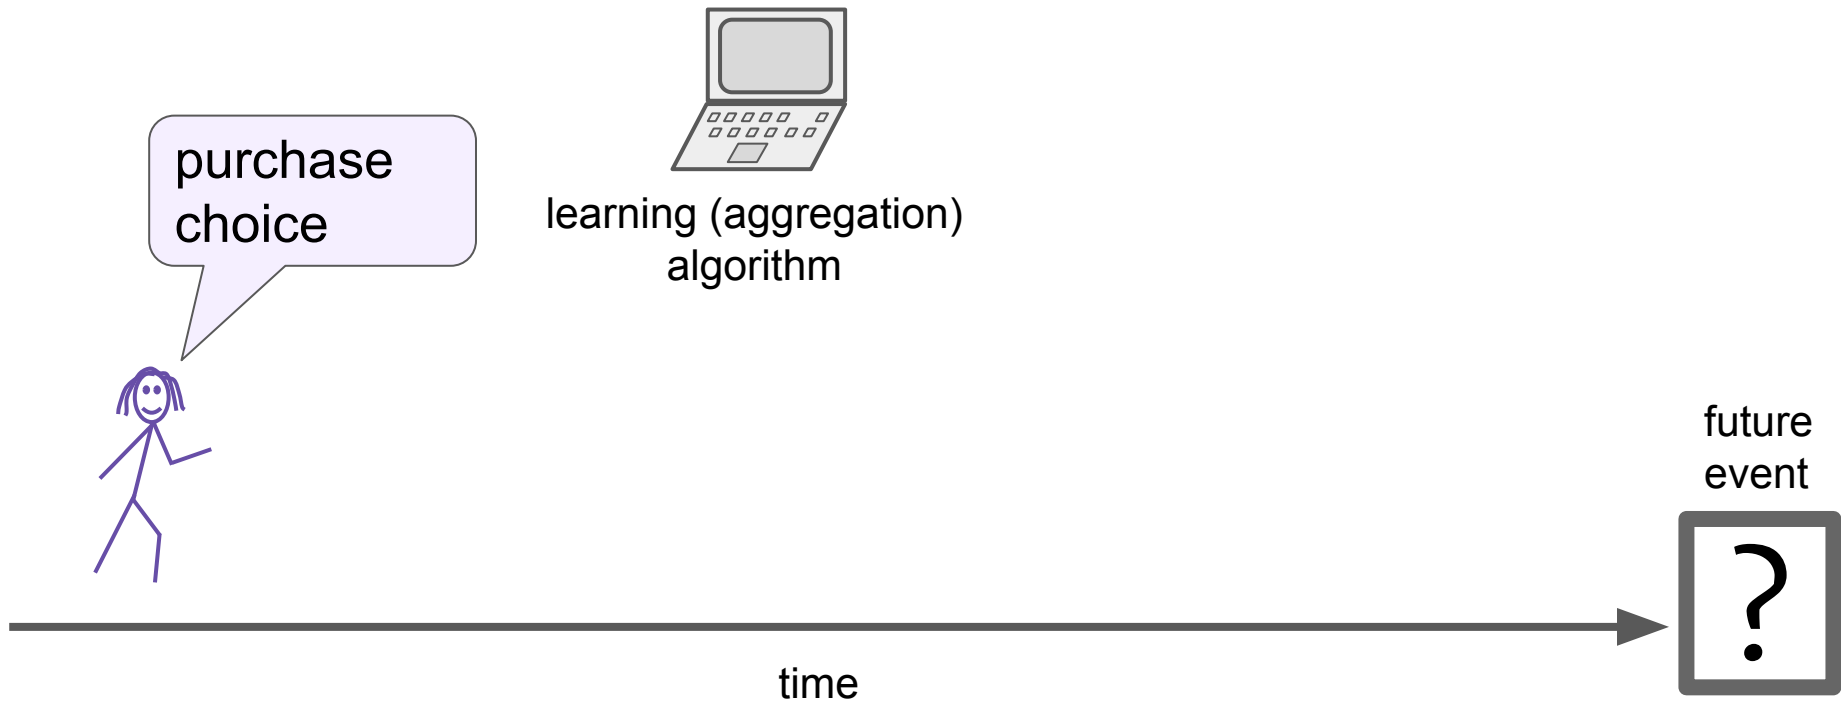


future
event

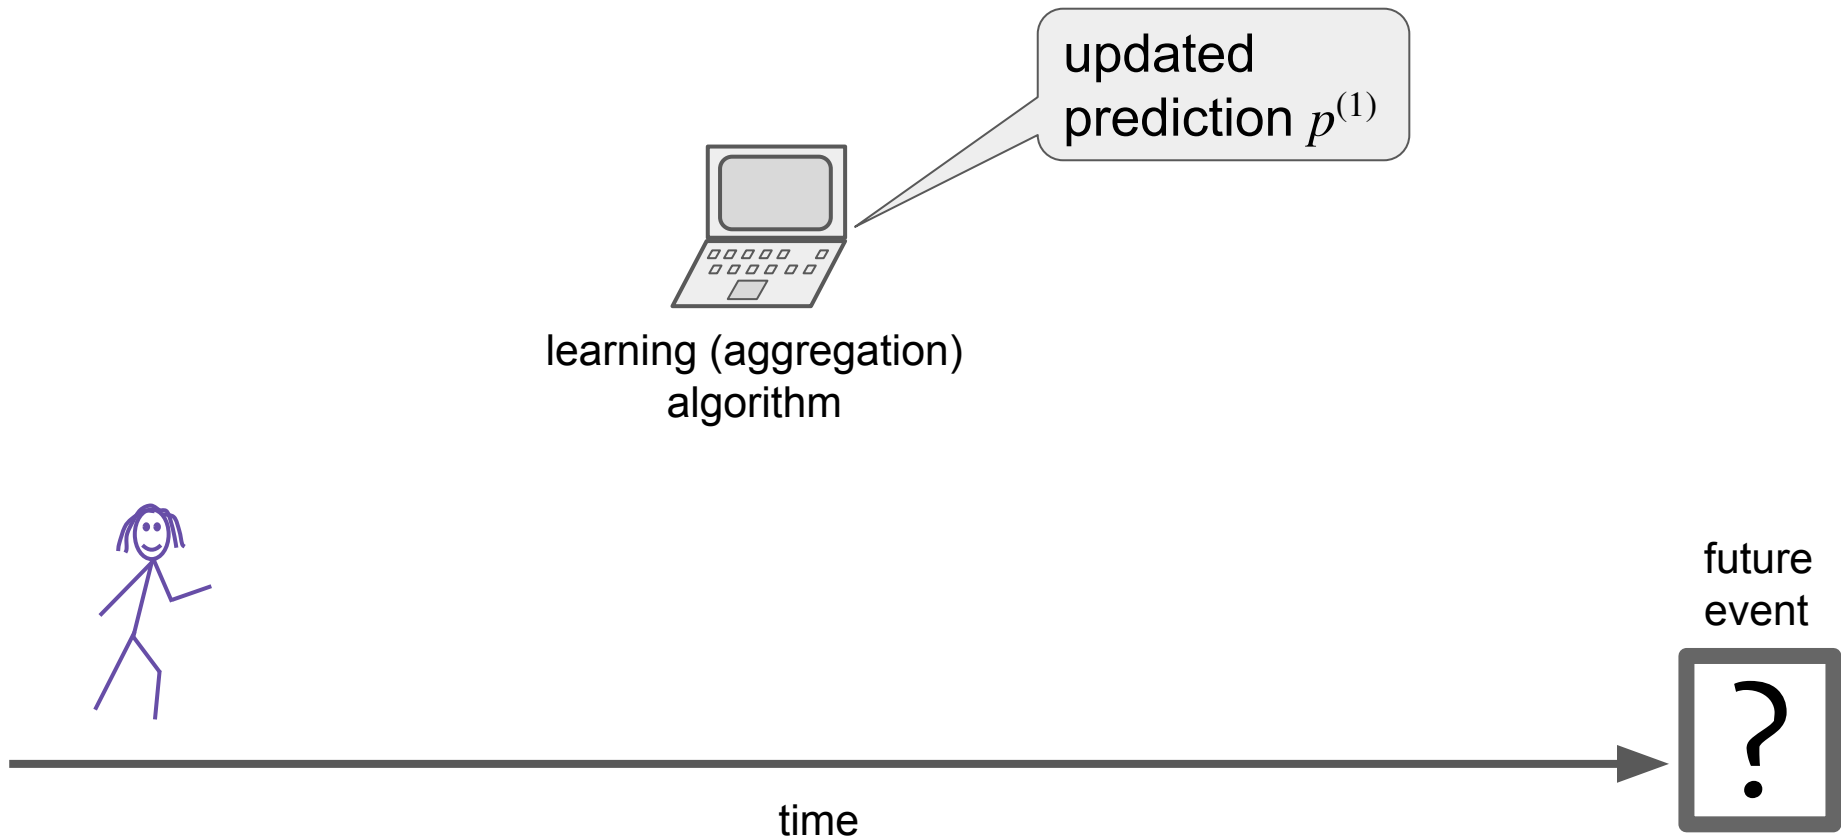


time

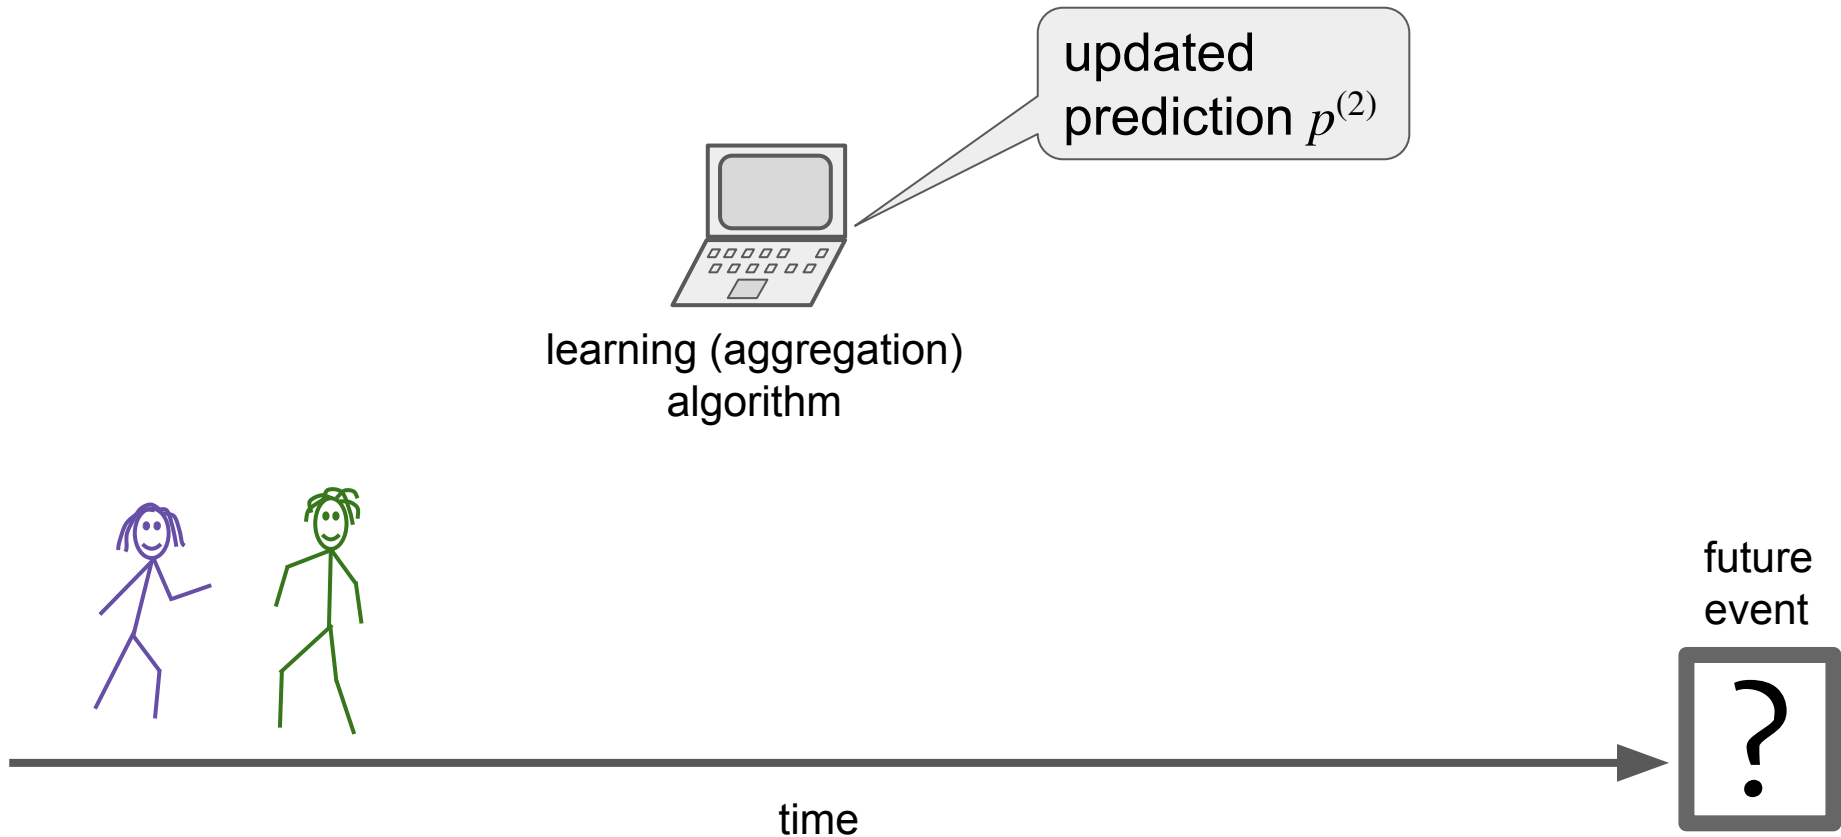
Prediction markets



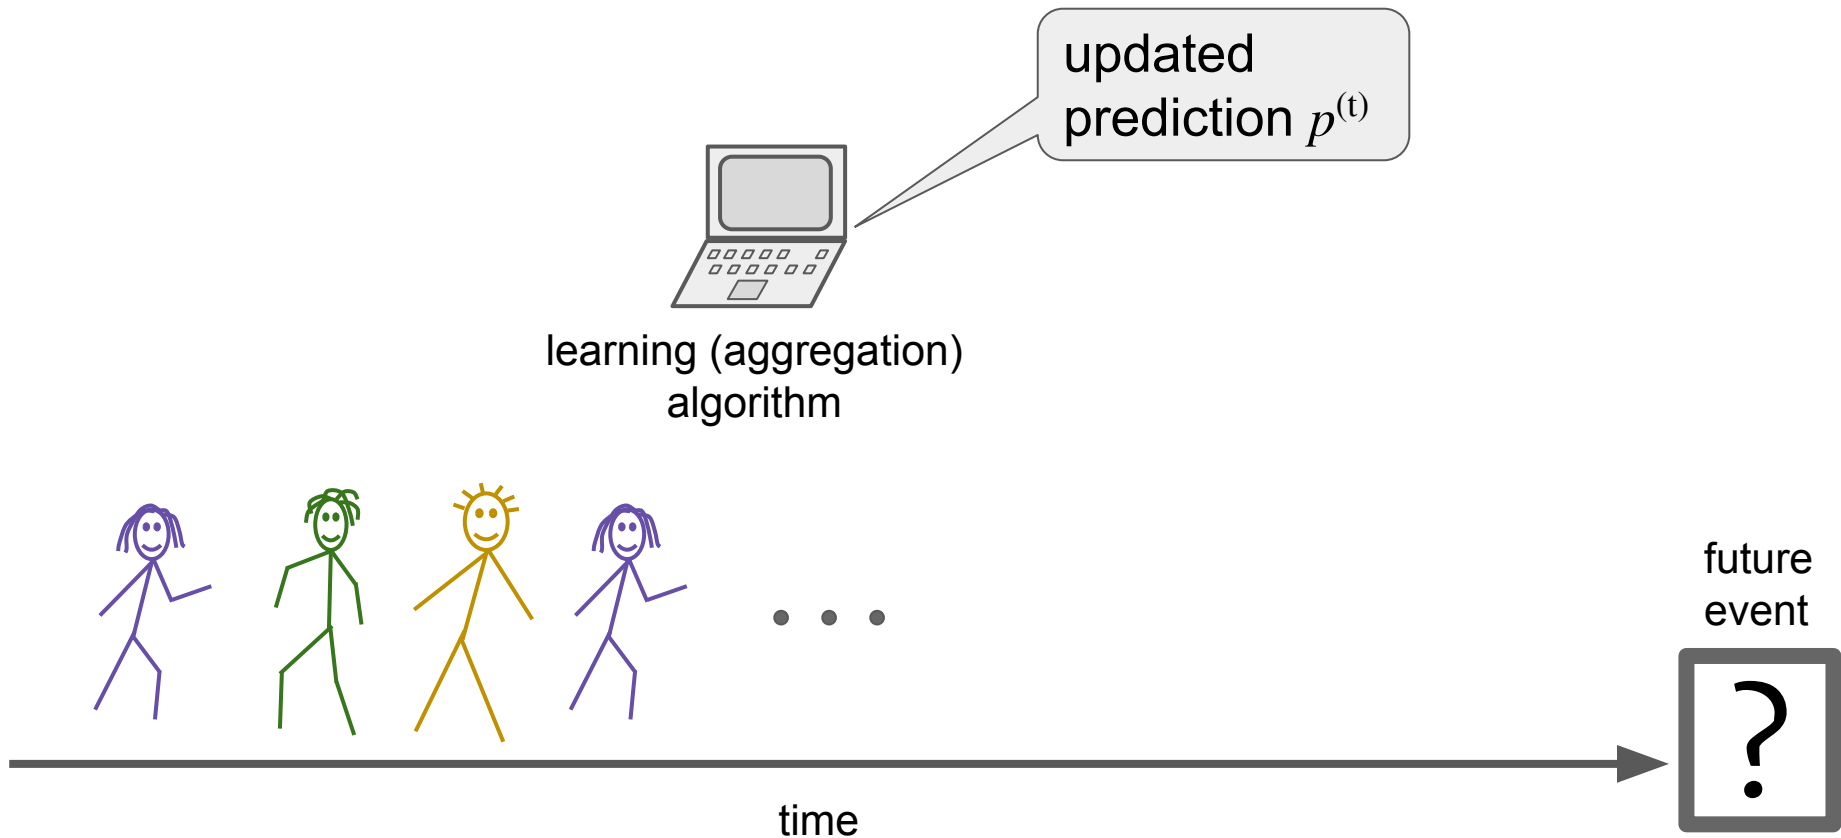
Prediction markets



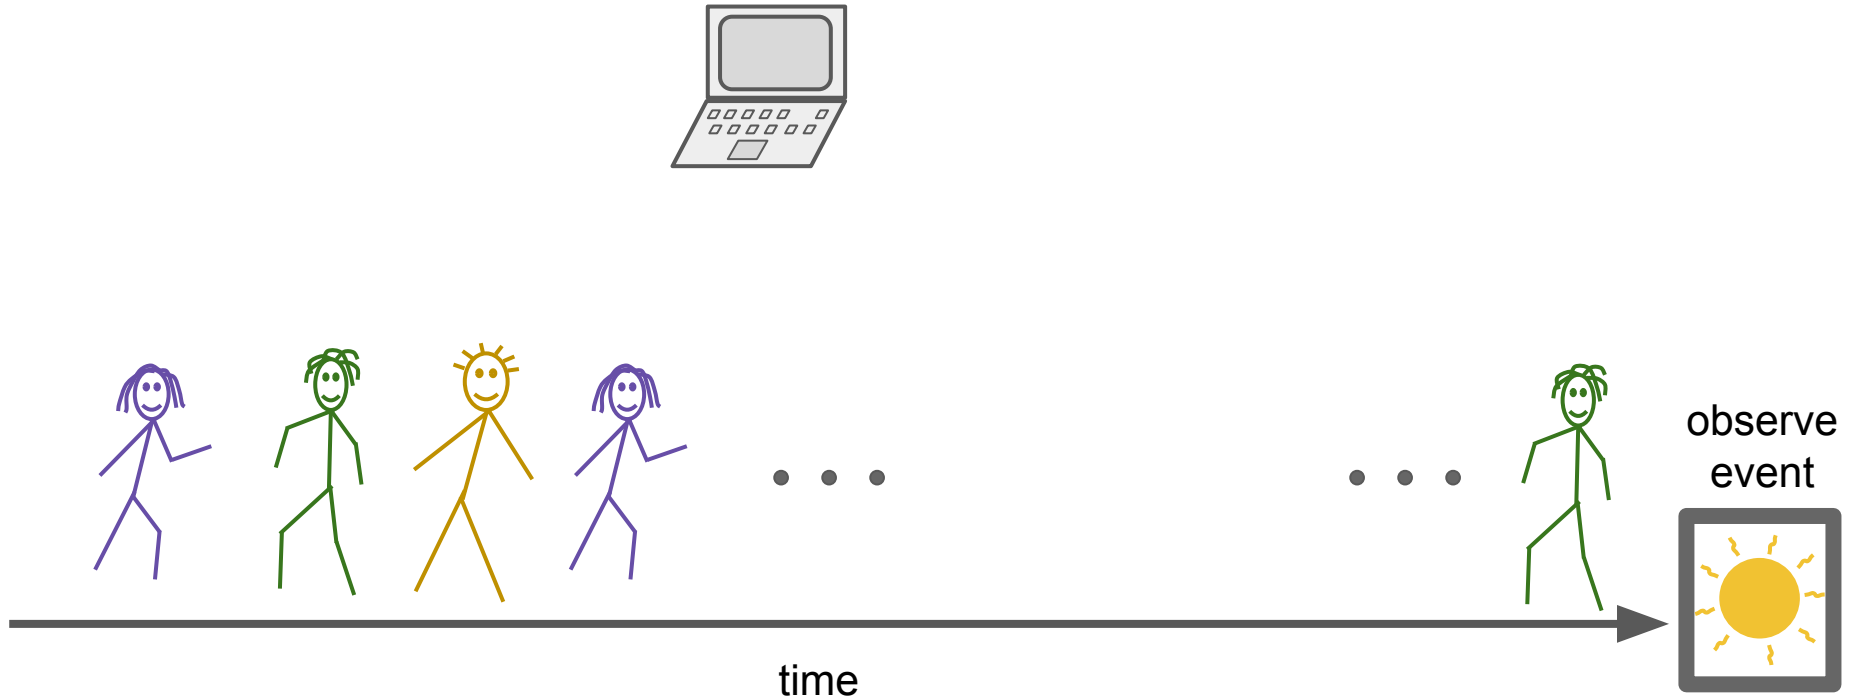
Prediction markets



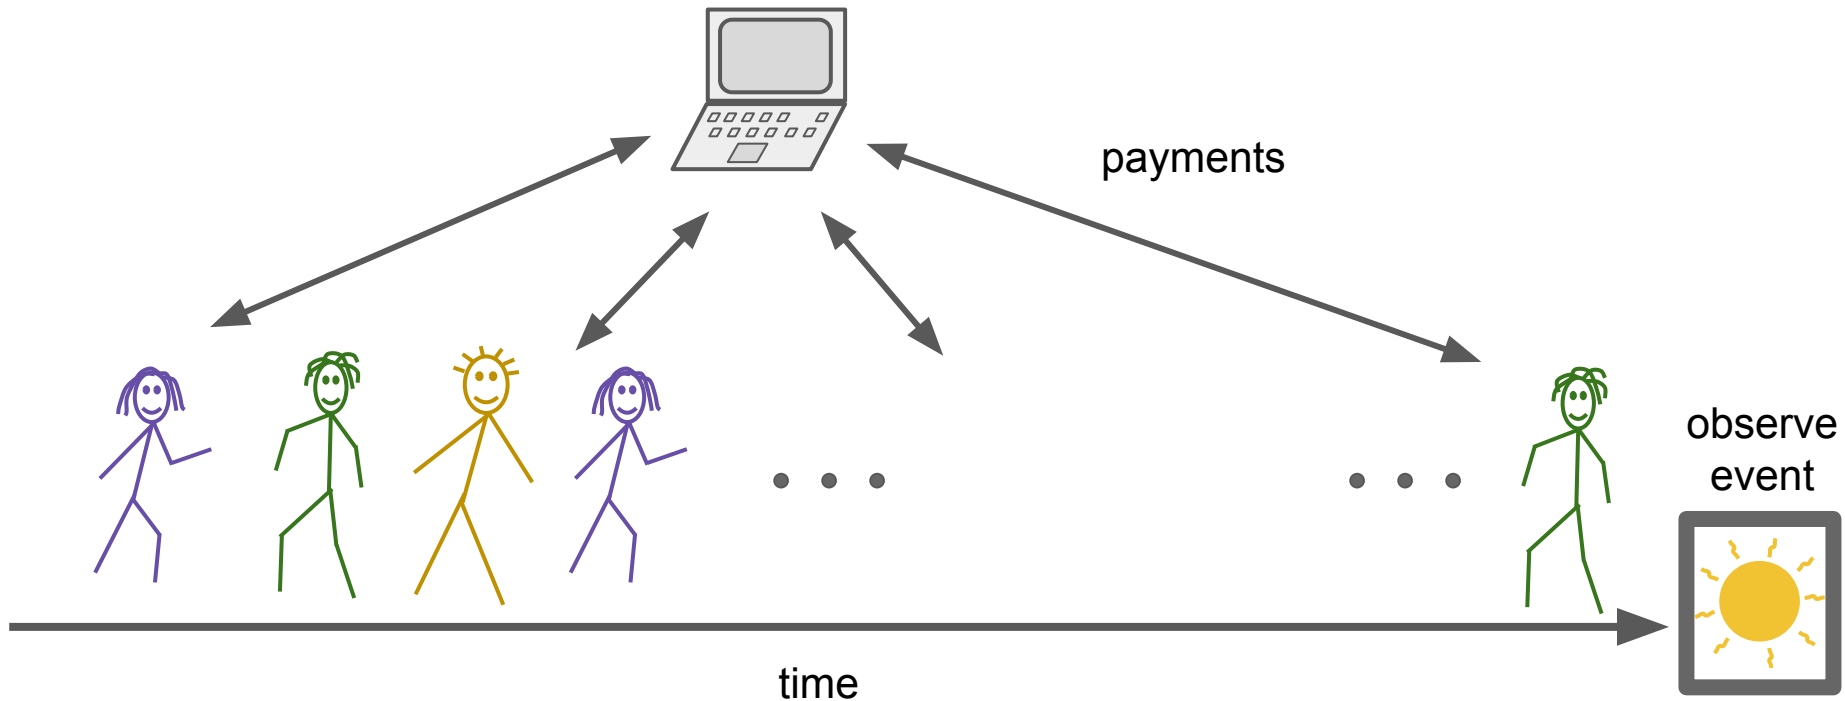
Prediction markets



Prediction markets




Prediction markets



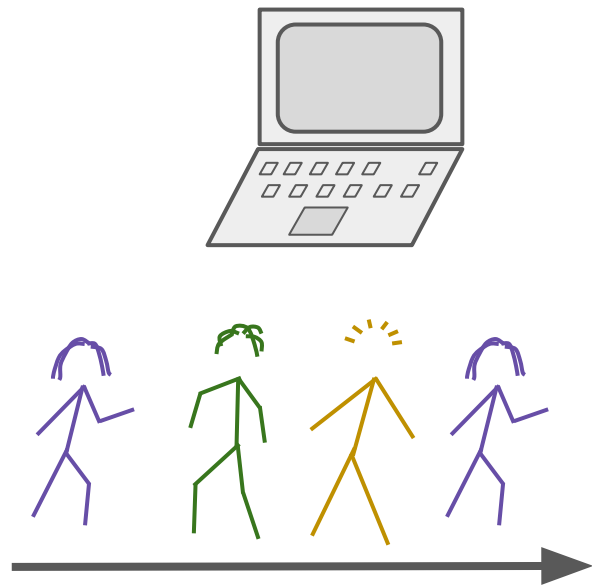
Prior work

- Abernethy-F.-W. 2015 (*Neural Information Processing Systems*)




Prior work

- Abernethy-F.-W. 2015 (*Neural Information Processing Systems*)
– *differentially private prediction markets*




Prior work

- Abernethy-F.-W. 2015 (*Neural Information Processing Systems*)
 - *differentially private* prediction markets
 - financial loss may not be bounded!



Prior work


- Abernethy-F.-W. 2015 (*Neural Information Processing Systems*)
 - *differentially private* prediction markets
 - financial loss may not be bounded!
- Cummings, Pennock, Wortman Vaughan 2016 (*EC*)
 - Impossibility: all private market makers have unbounded financial loss!



This paper

Main result: a market construction with:


- differential privacy
- incentives to participate
- accuracy/fidelity of predictions
- **Bounded worst-case financial loss**




This paper

Main result: a market construction with:

- differential privacy (ϵ)
- incentives to participate (α)
- accuracy/fidelity of predictions (α , dimension d)
- **Bounded worst-case financial loss $\tilde{O}(d / \epsilon \alpha)$**




Extensions (cf AFW'15):
purchasing data, kernel methods....



This paper

Main result: a market construction with:

- differential privacy (ϵ)
- incentives to participate (α)
- accuracy/fidelity of predictions (α , dimension d)
- **Bounded worst-case financial loss $\tilde{O}(d / \epsilon \alpha)$**



Extensions (cf AFW'15):
purchasing data, kernel methods....

Thanks!

