Differentially Private, Bounded-Loss Prediction Markets

Bo Waggoner UPenn⟶Microsoft
with Rafael Frongillo Colorado

WADE, June 2018

Outline

→ A. Cost function based prediction markets

B. Summary of results and prior work

C. Construction

Prediction markets

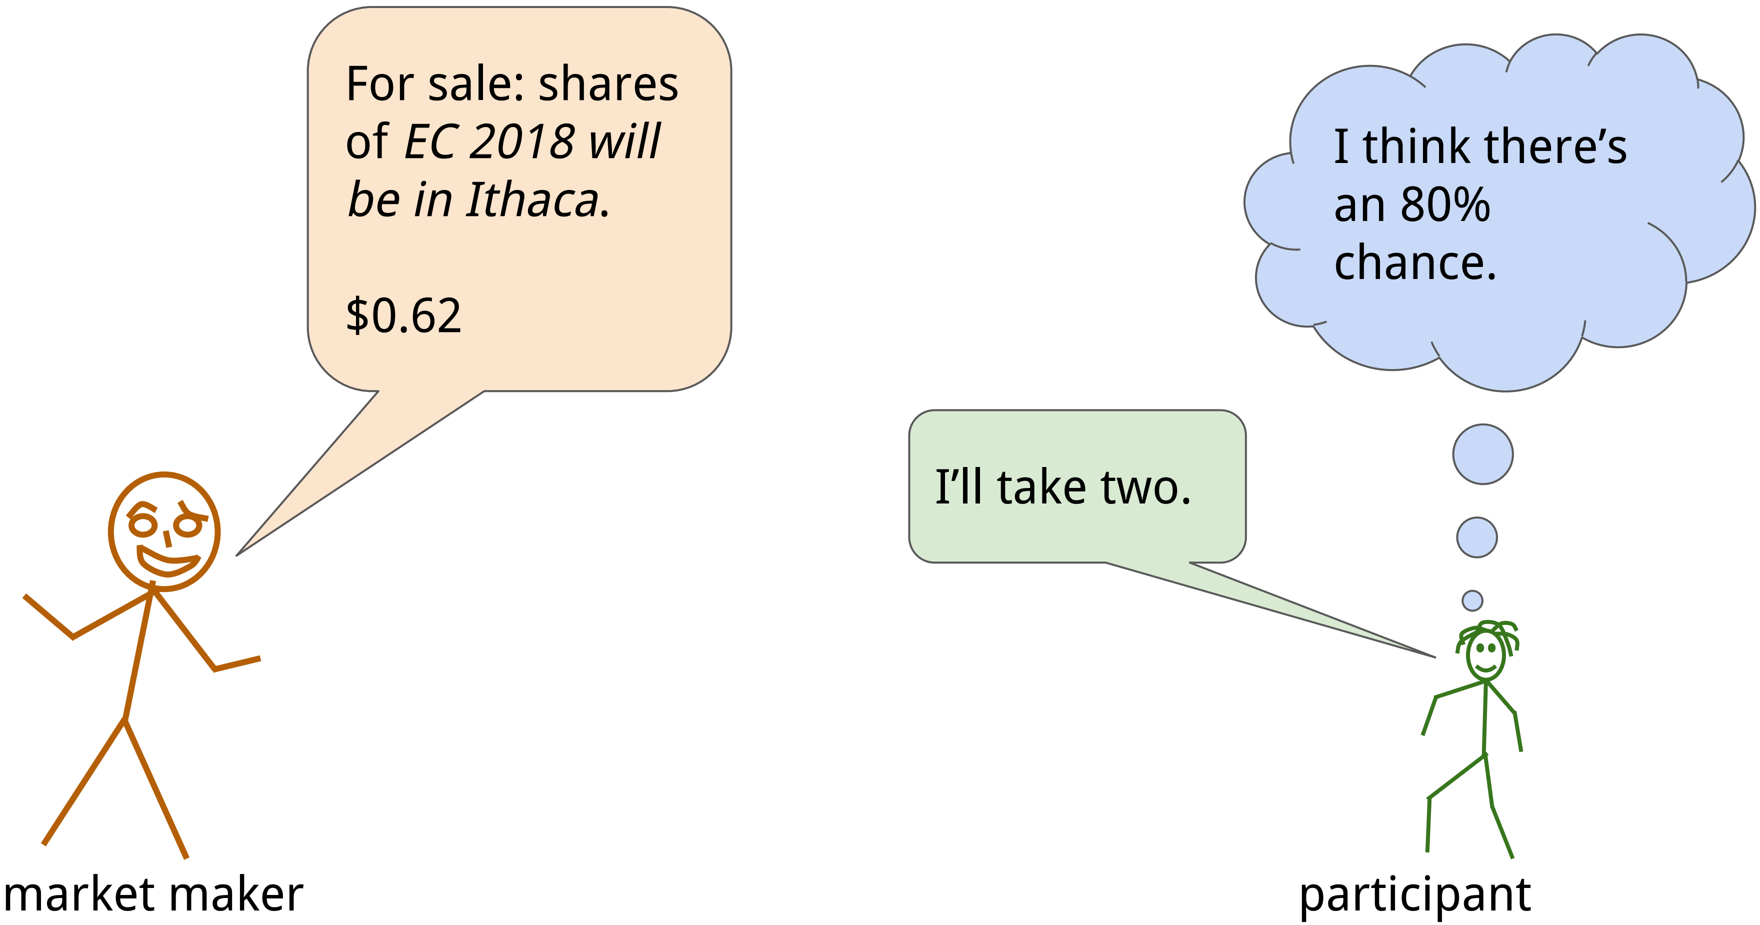

Prediction markets

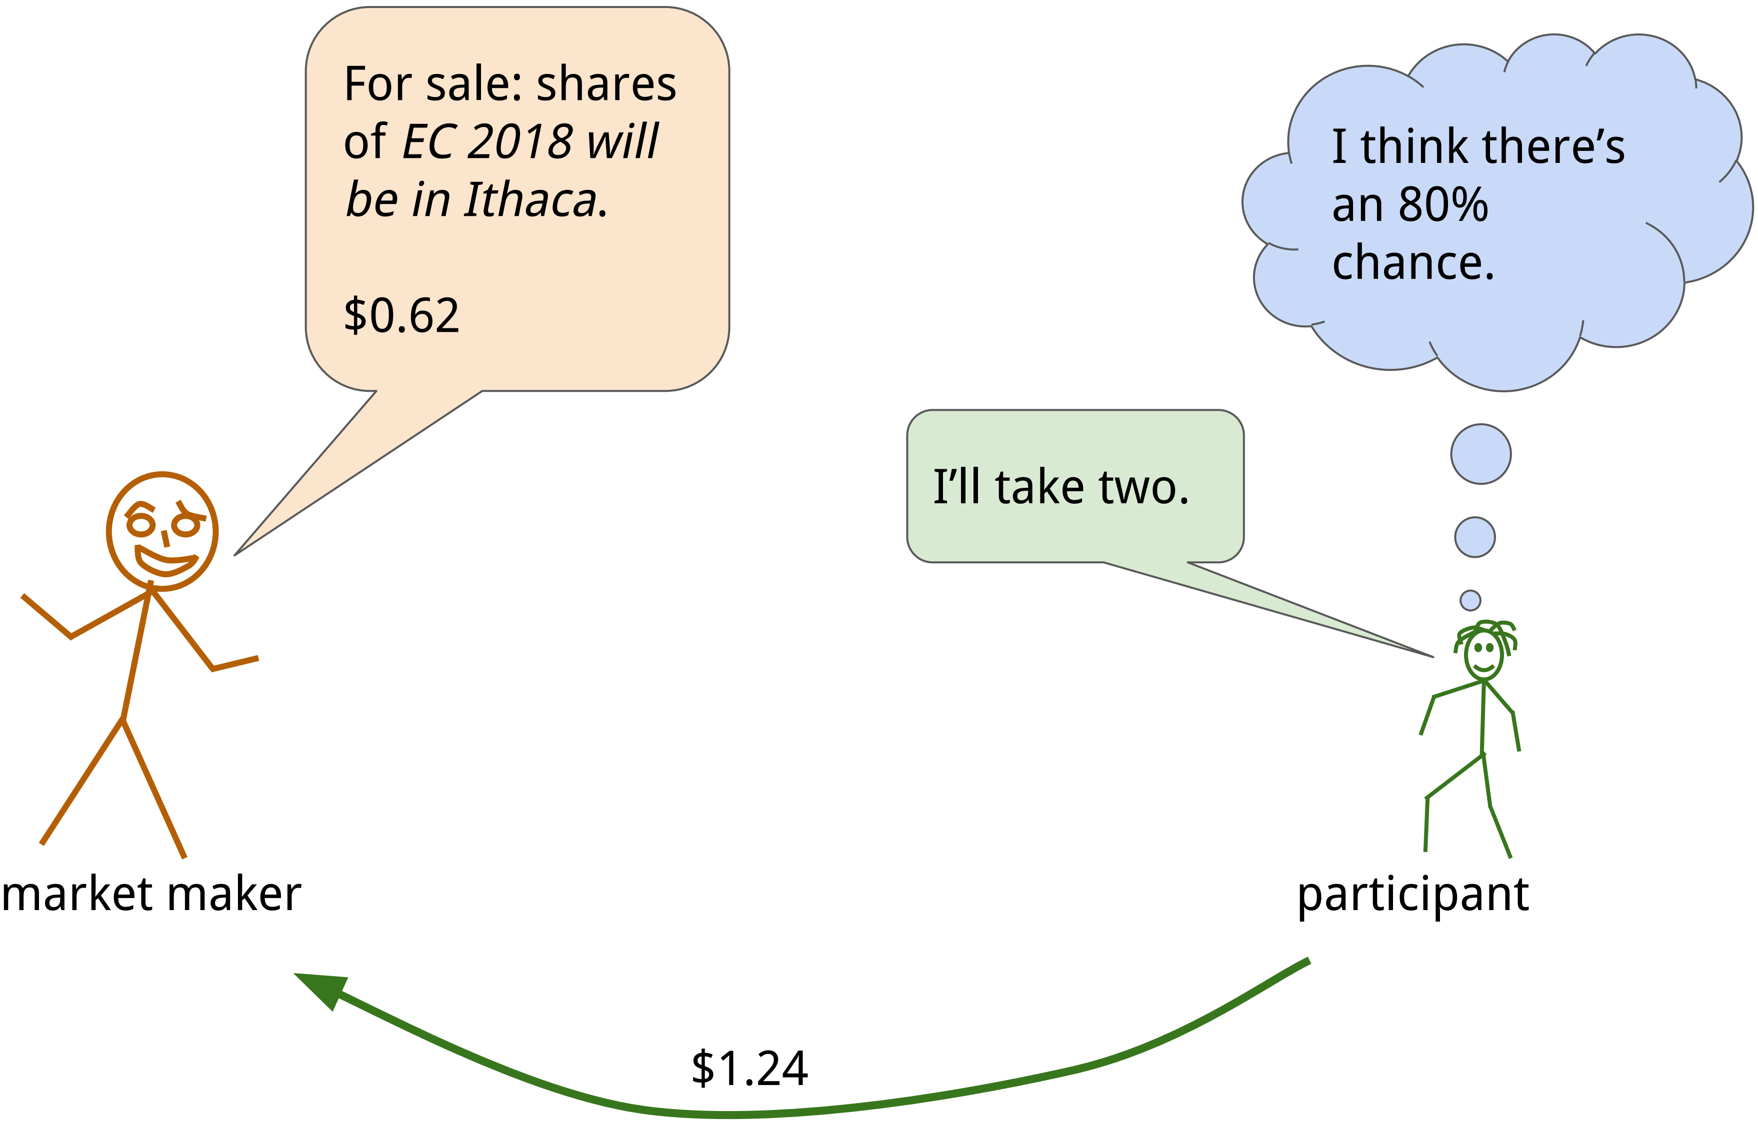

Later

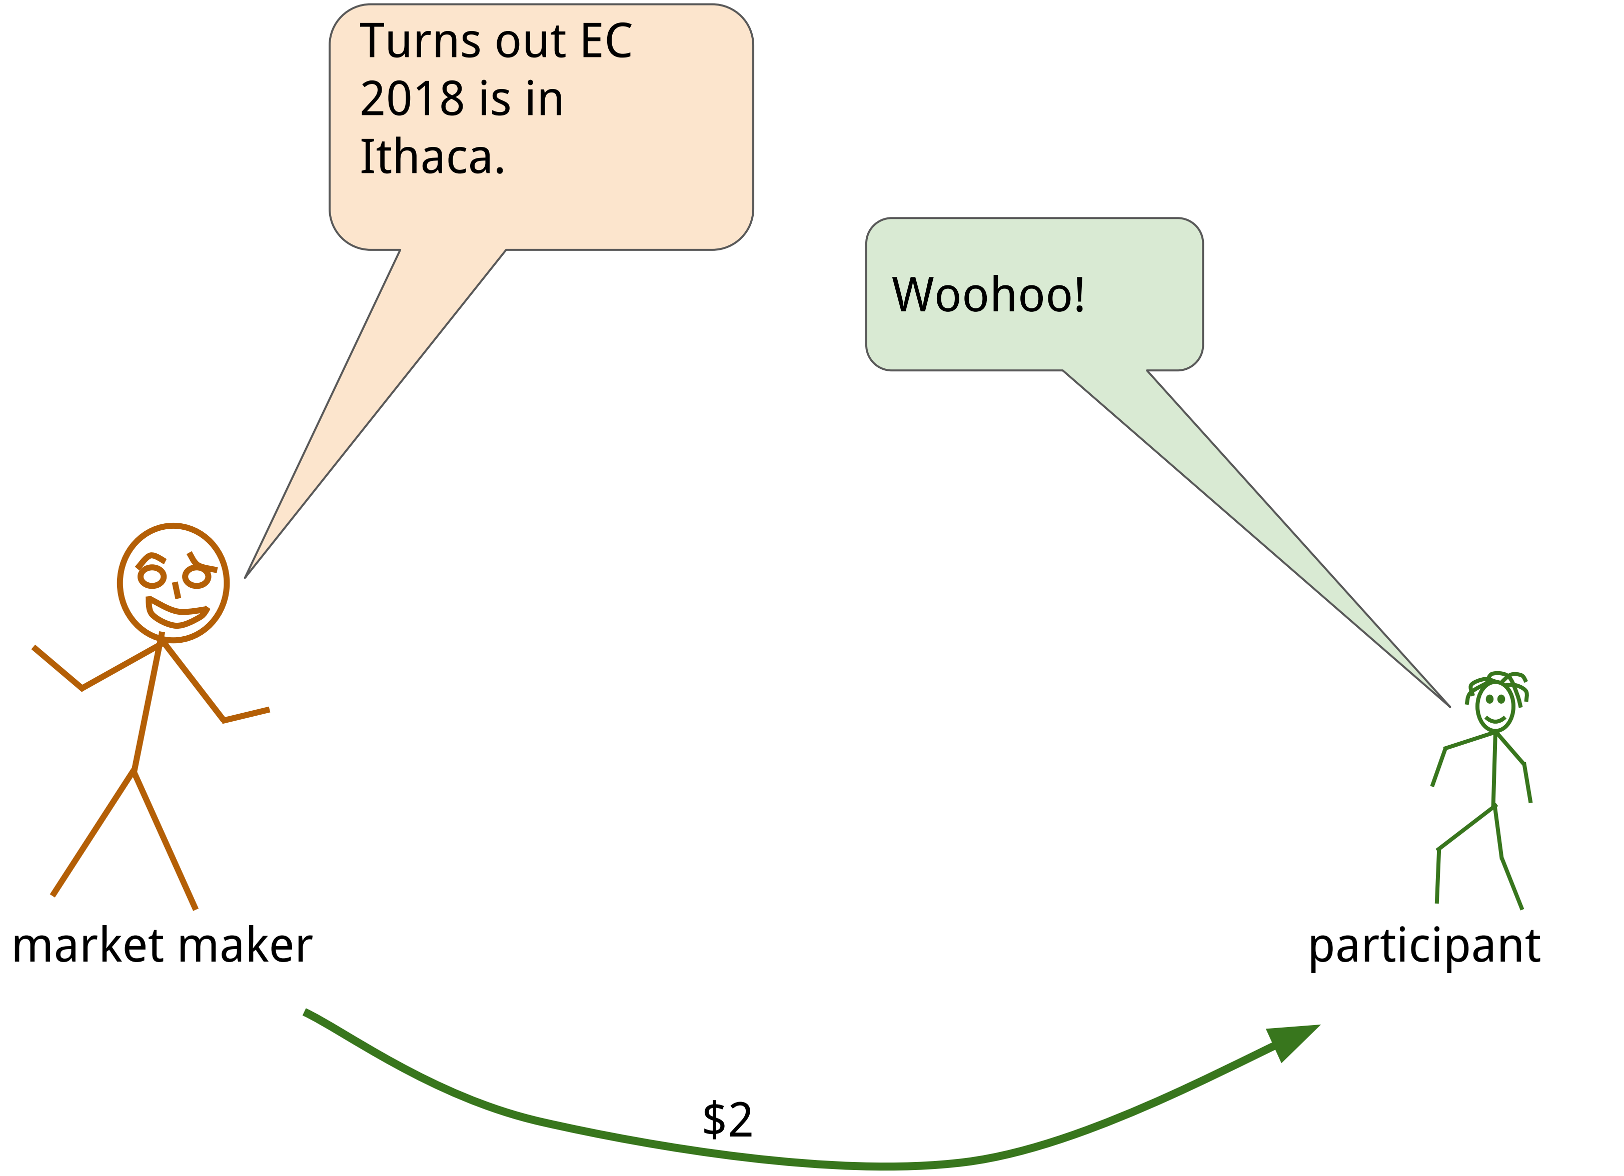
Turns out EC 2018 is in Ithaca.
Woohoo!
market maker
participant
$2

(In an alternate universe)

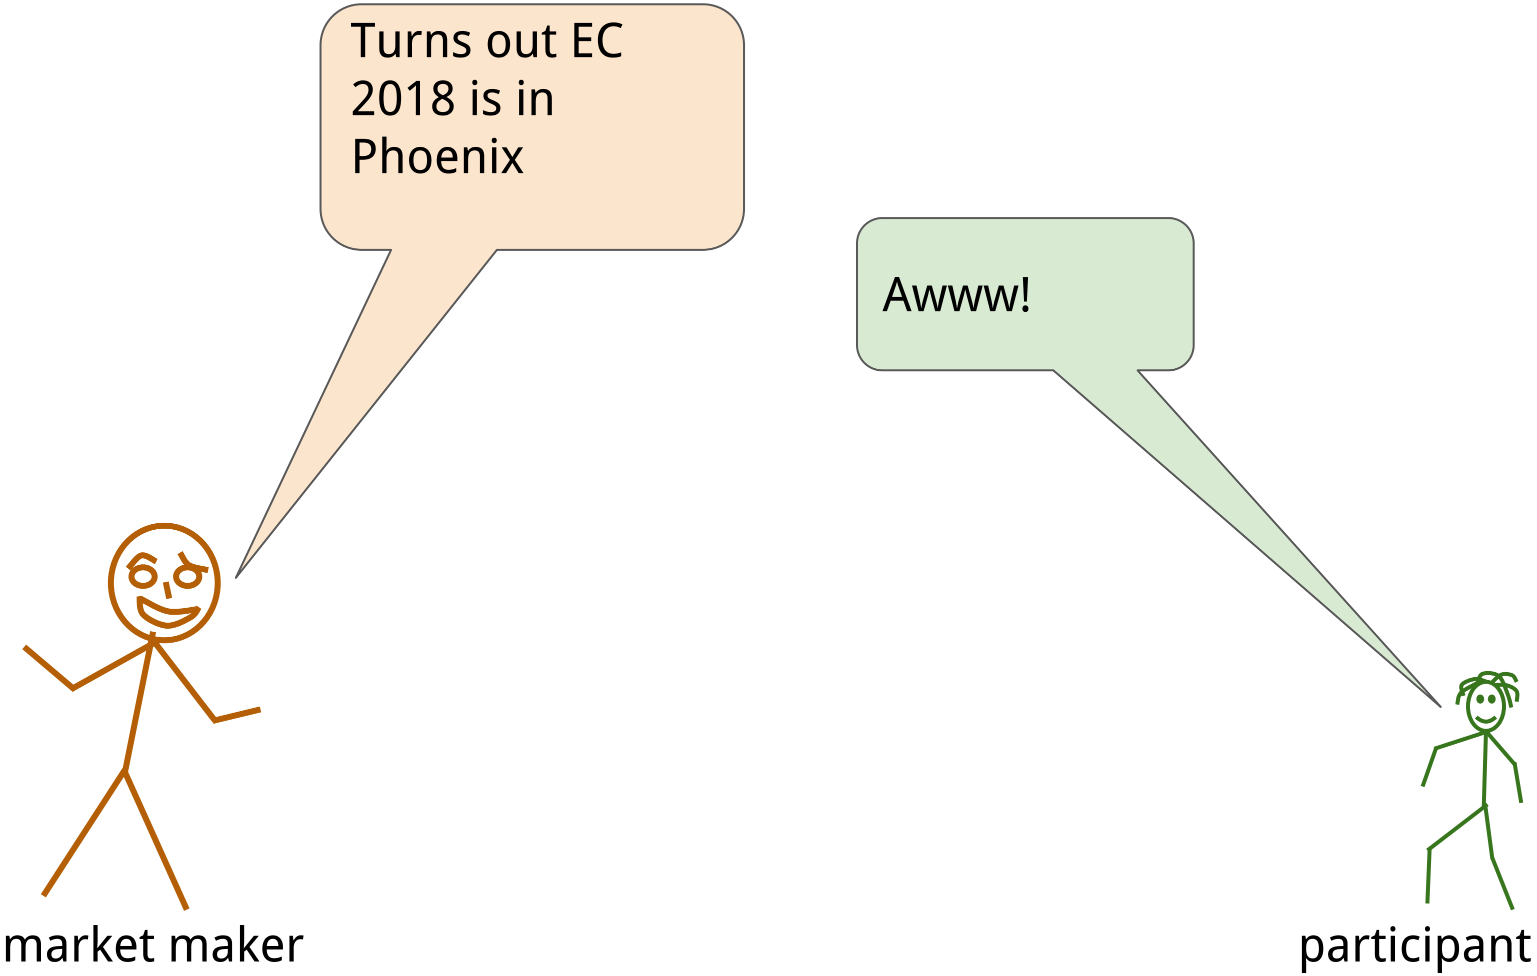

(no payoff)

Short selling

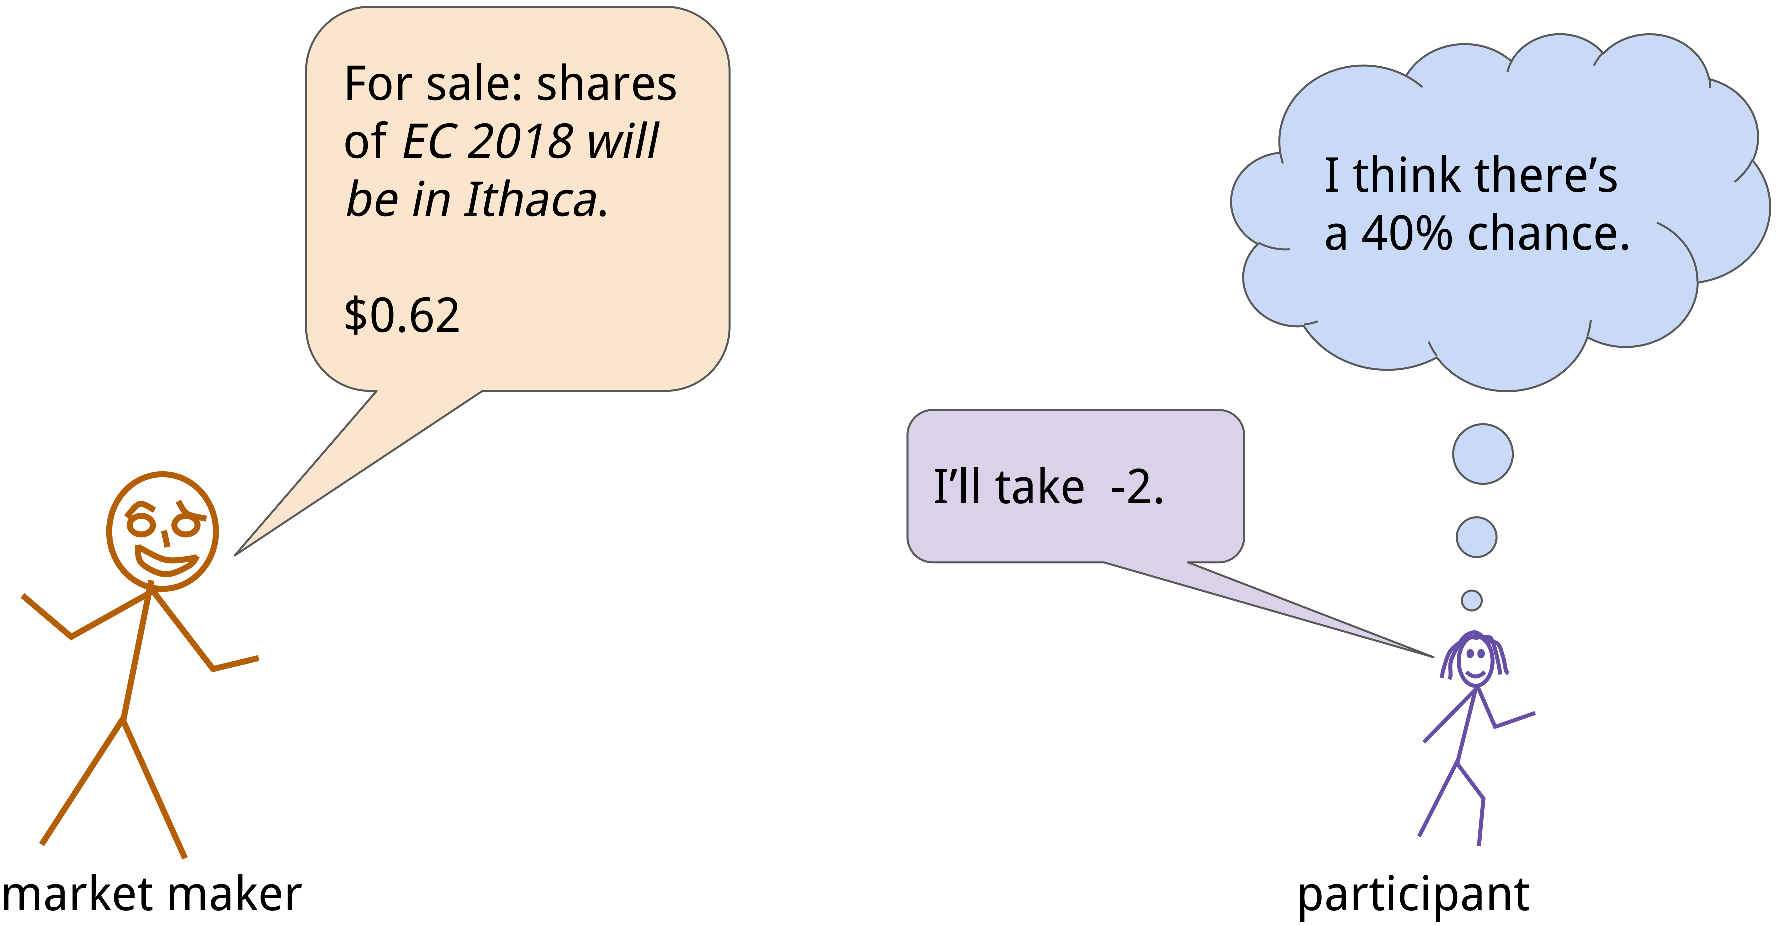

Short selling

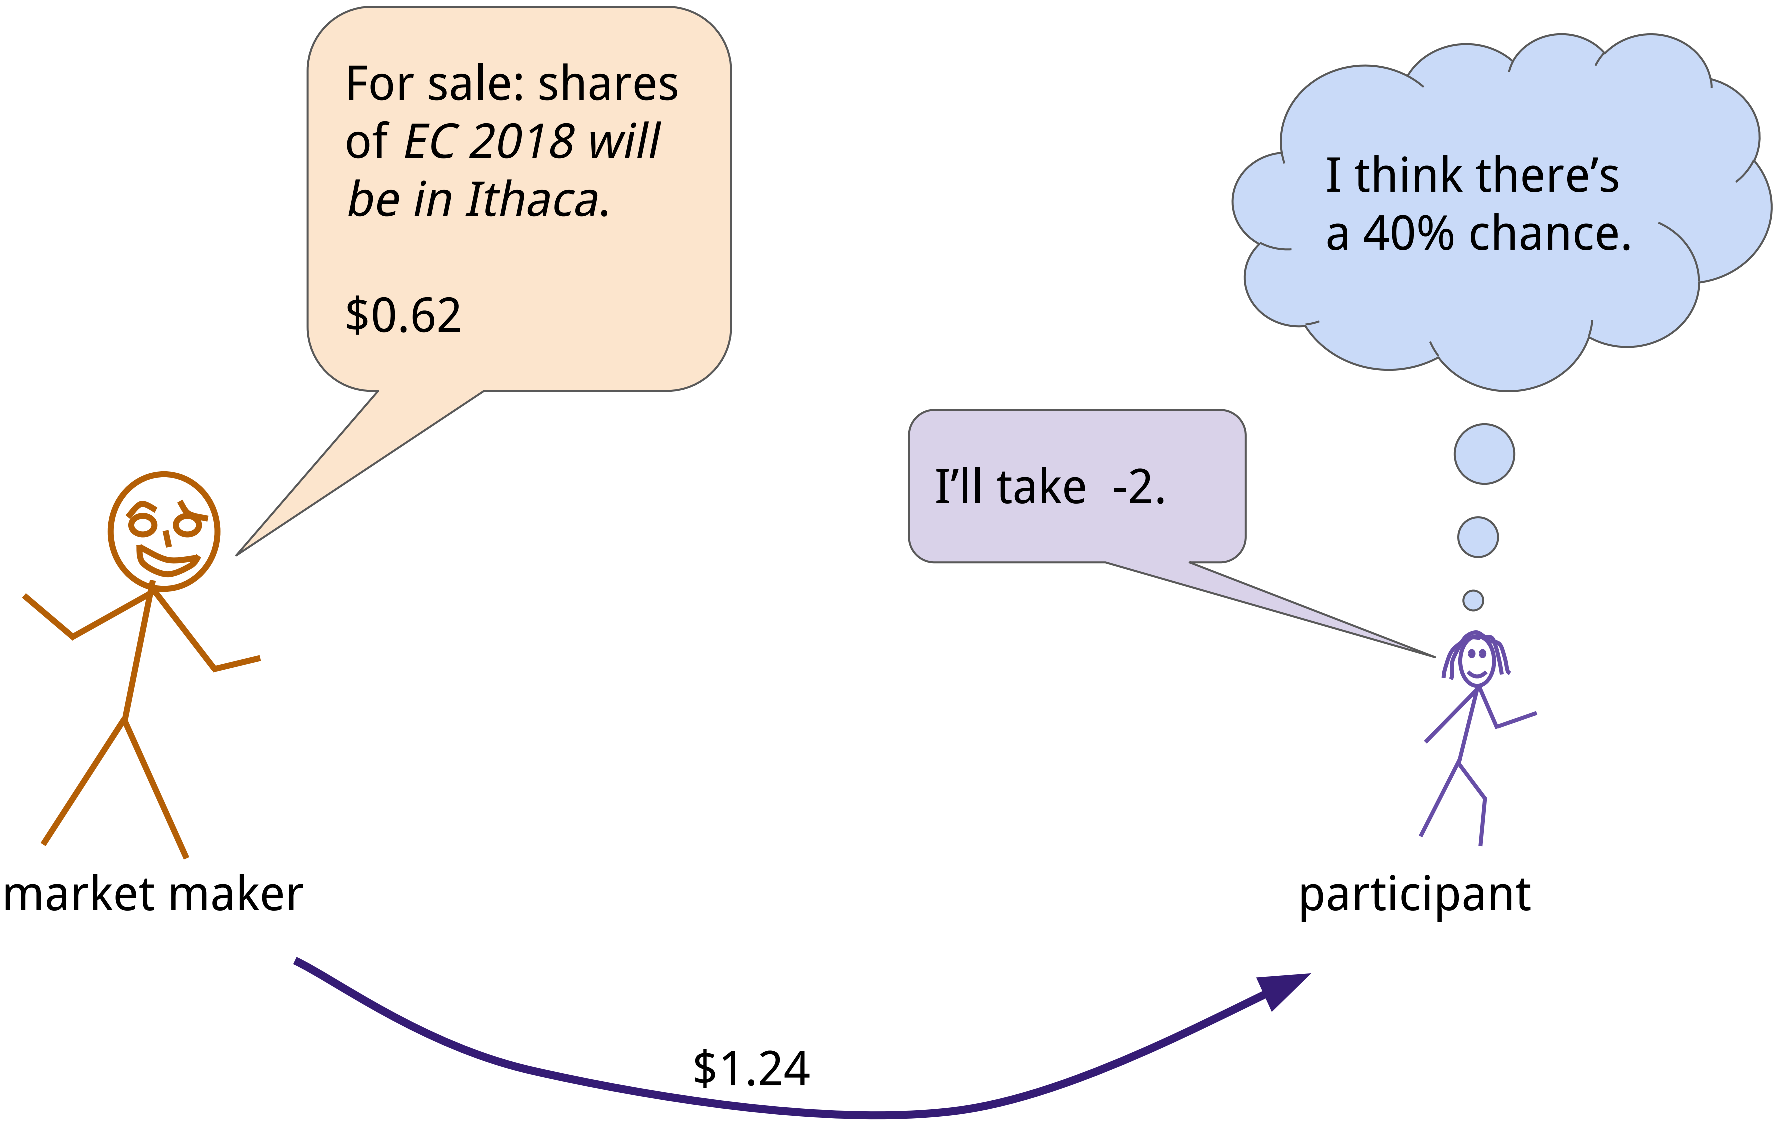

Later

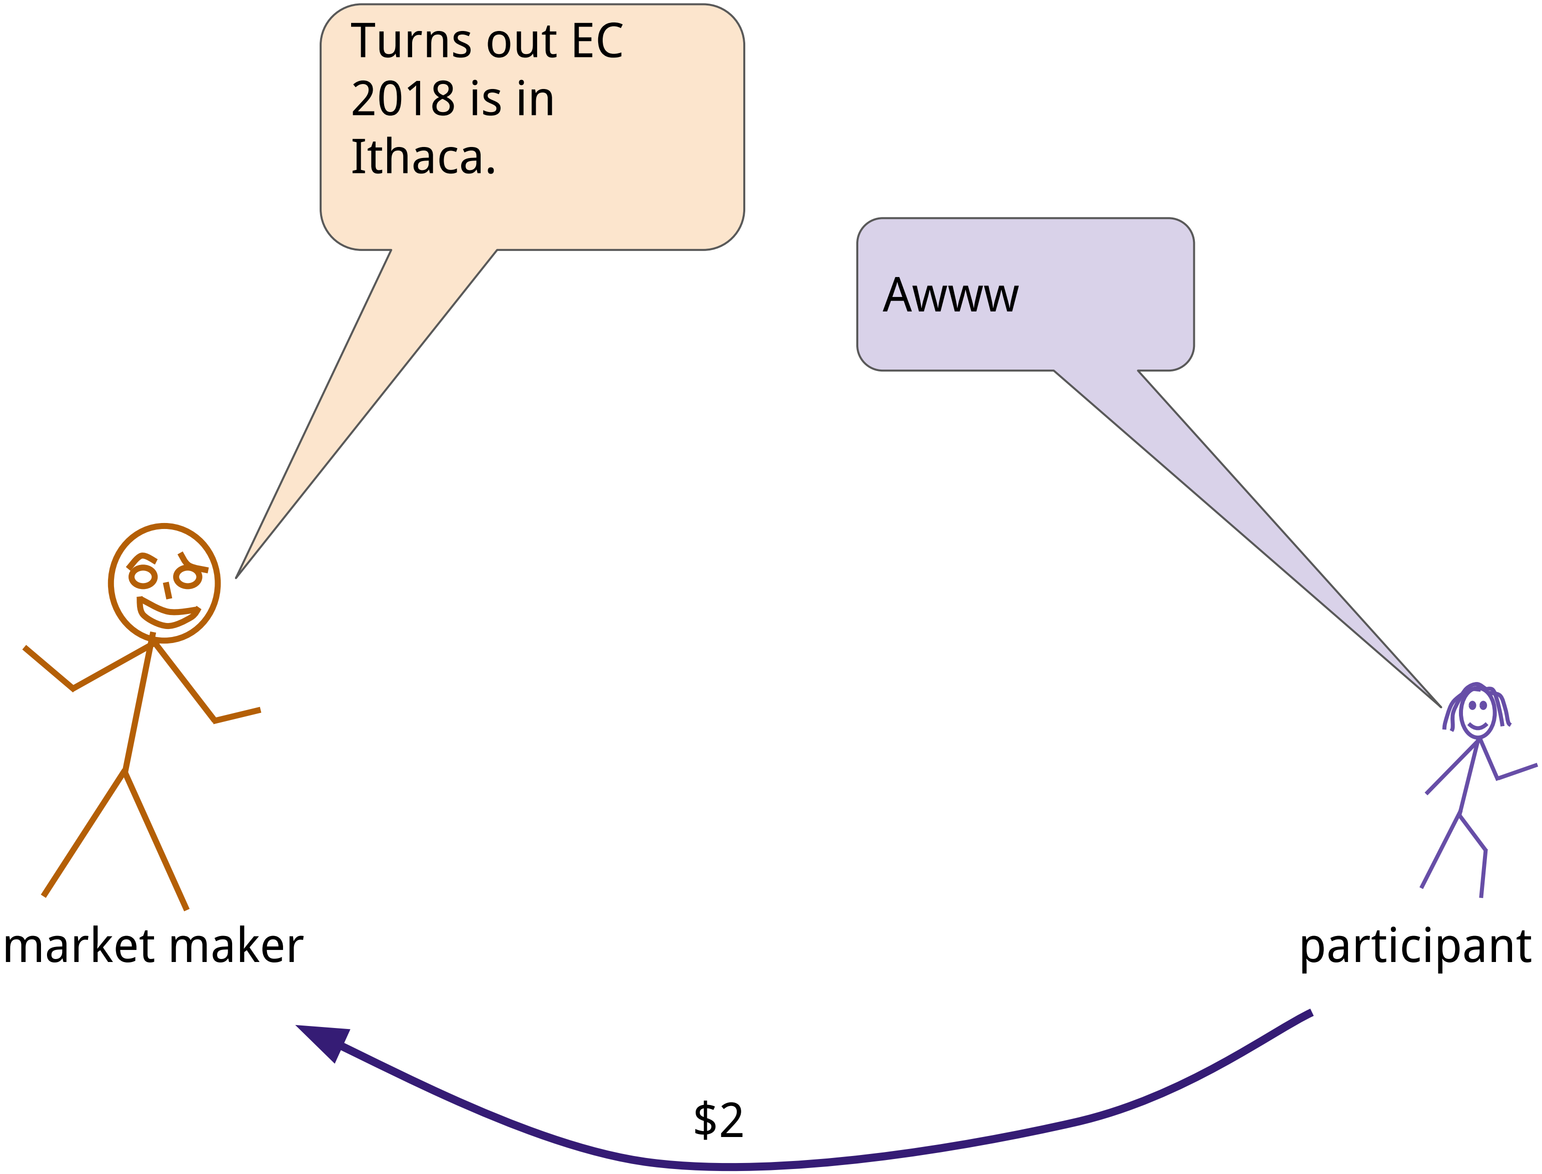
Turns out EC 2018 is in Ithaca.
Awww
market maker
participant
$2

(In an alternate universe)

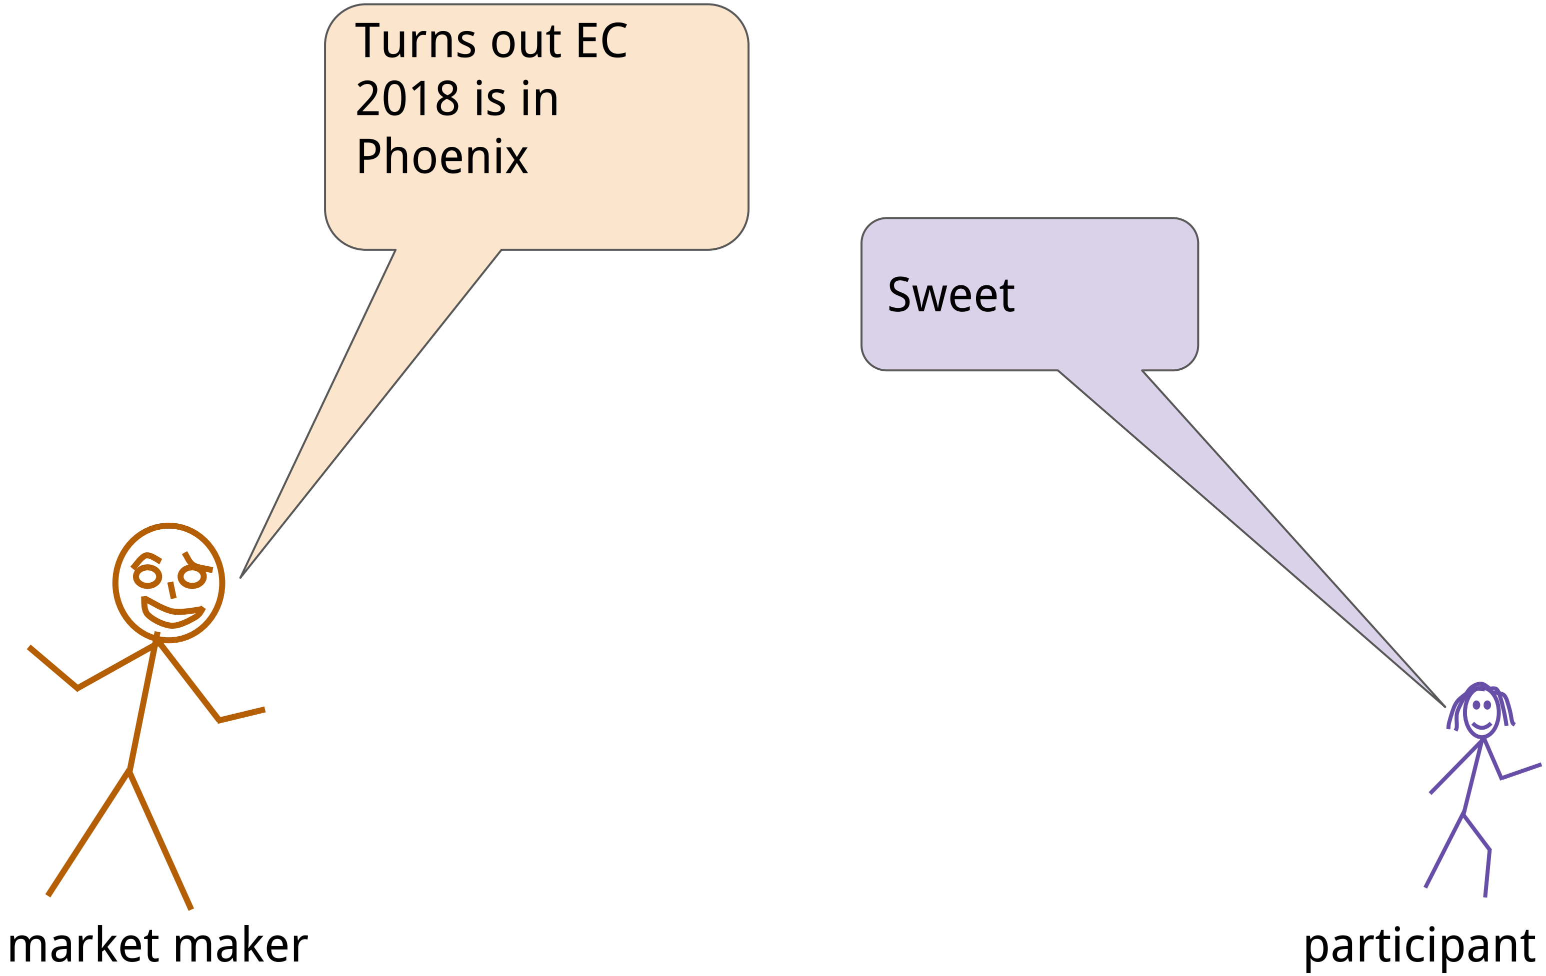

(no payoff)

Prediction markets - dynamics

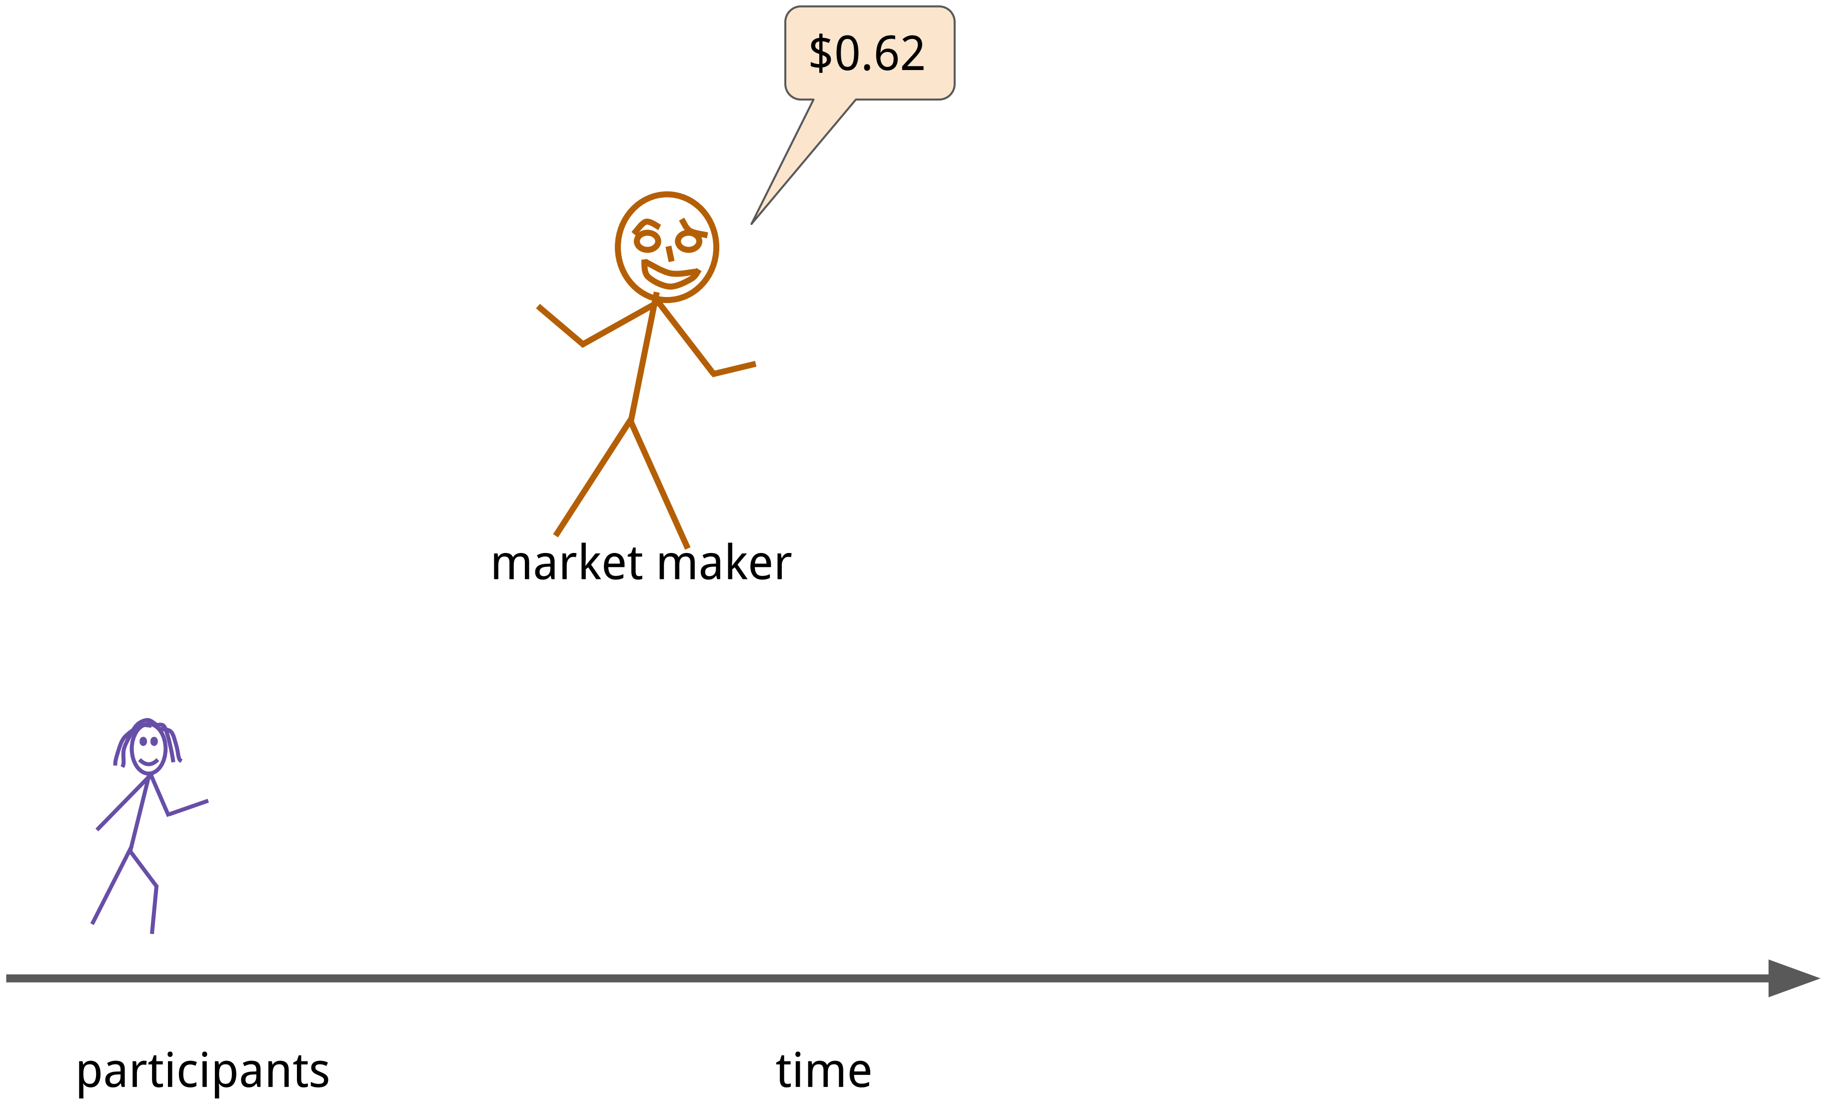

Prediction markets - dynamics

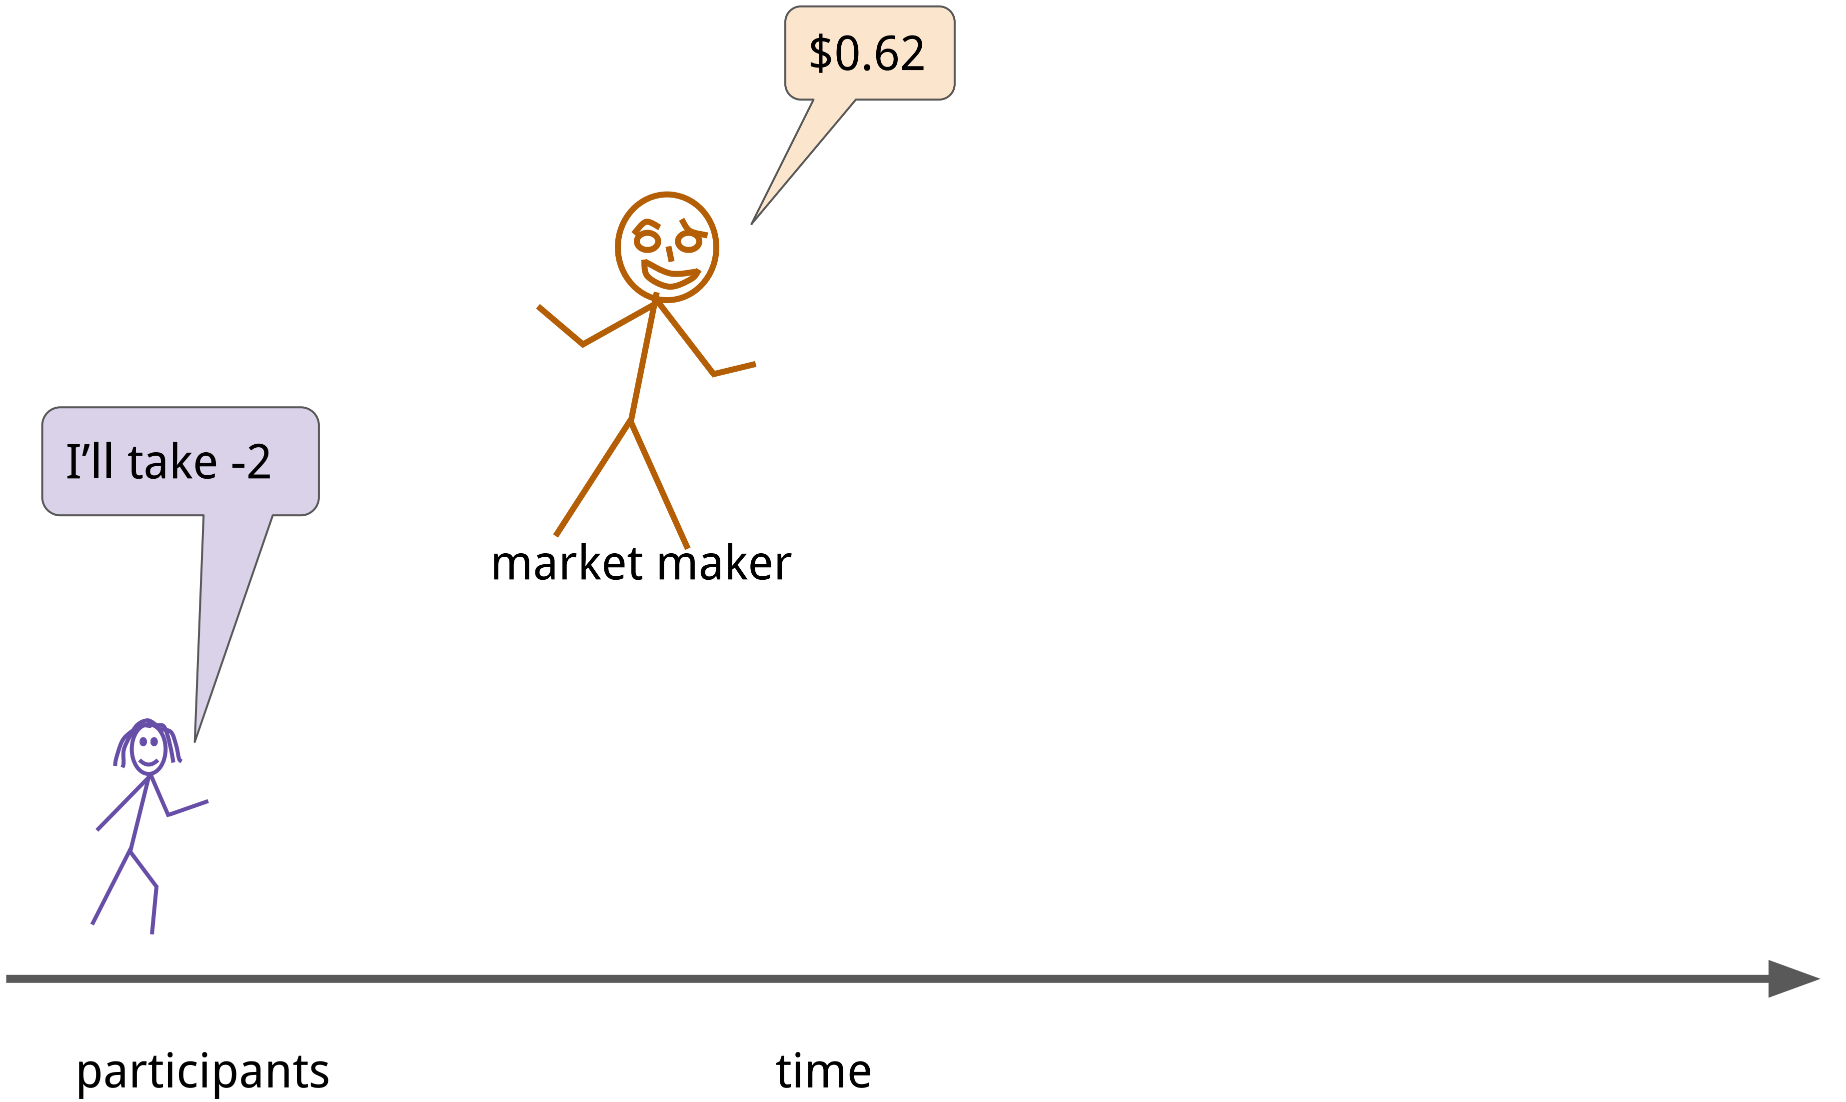

Prediction markets - dynamics

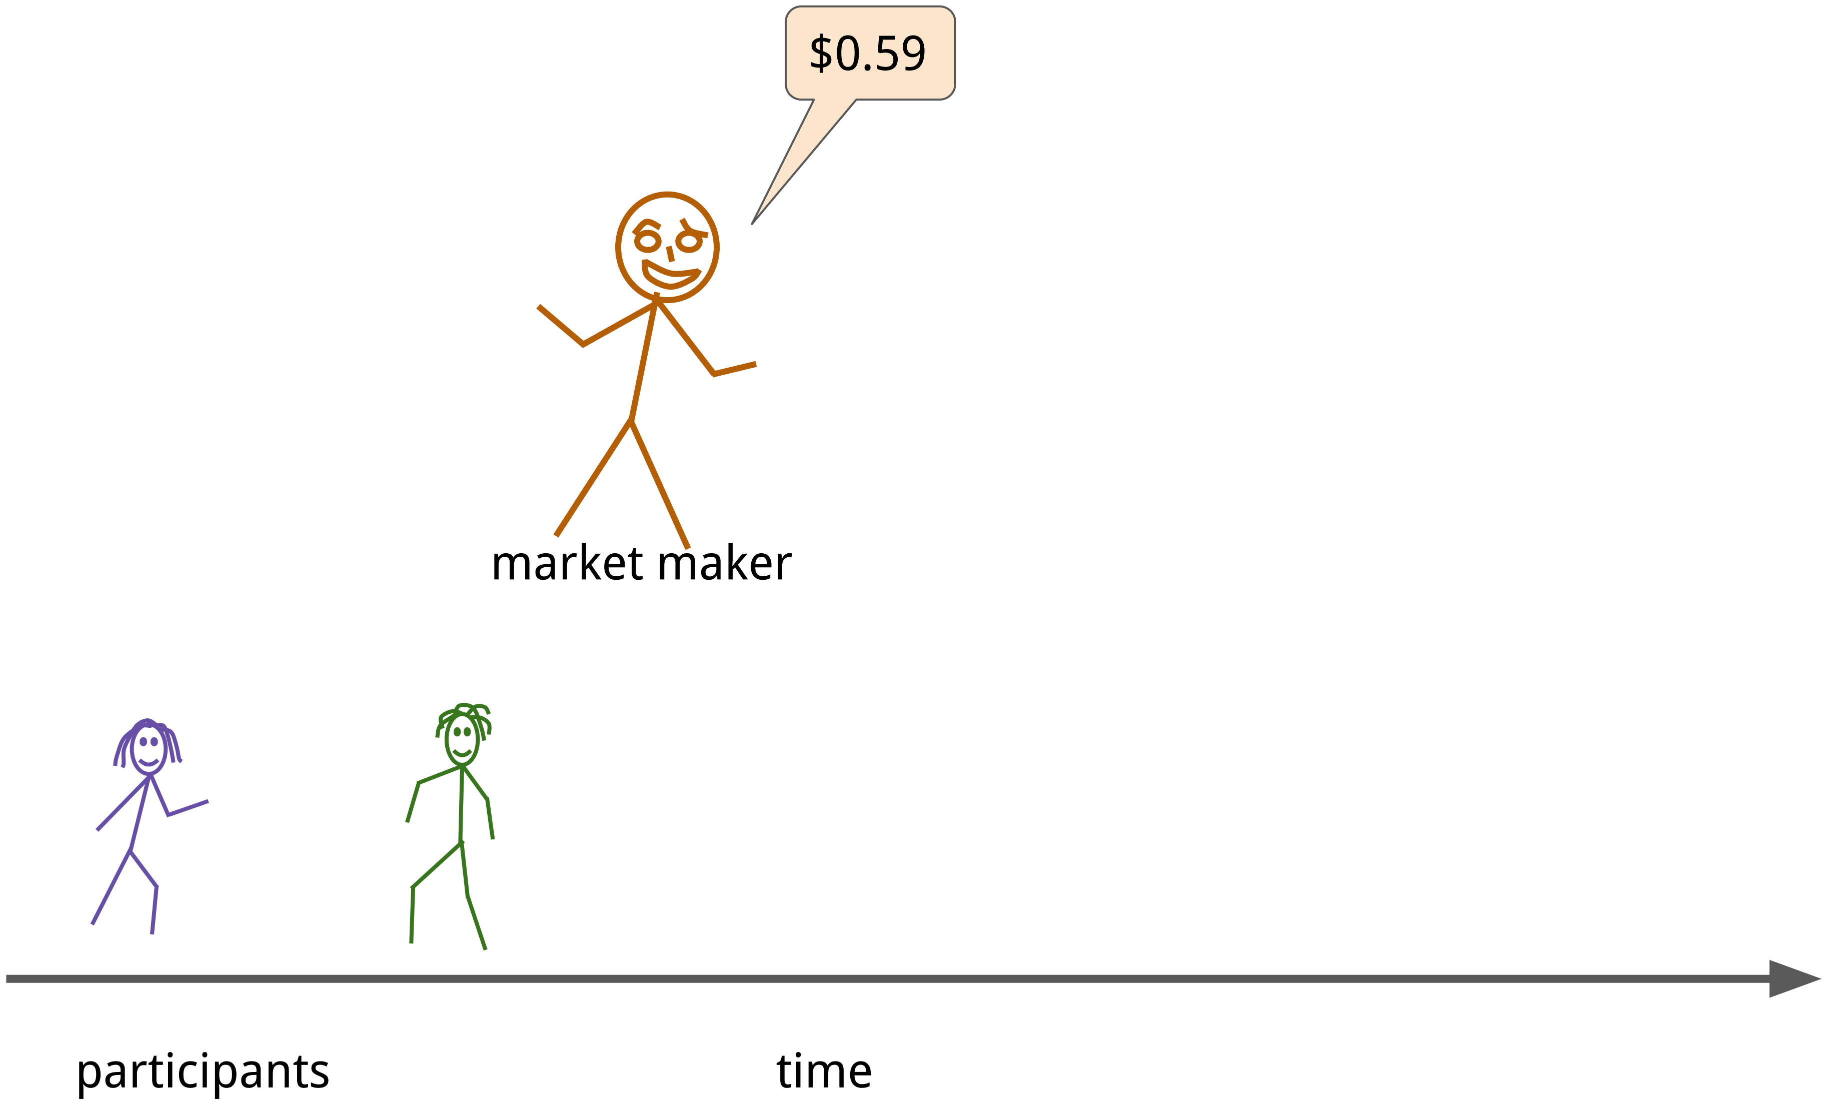

Prediction markets - dynamics

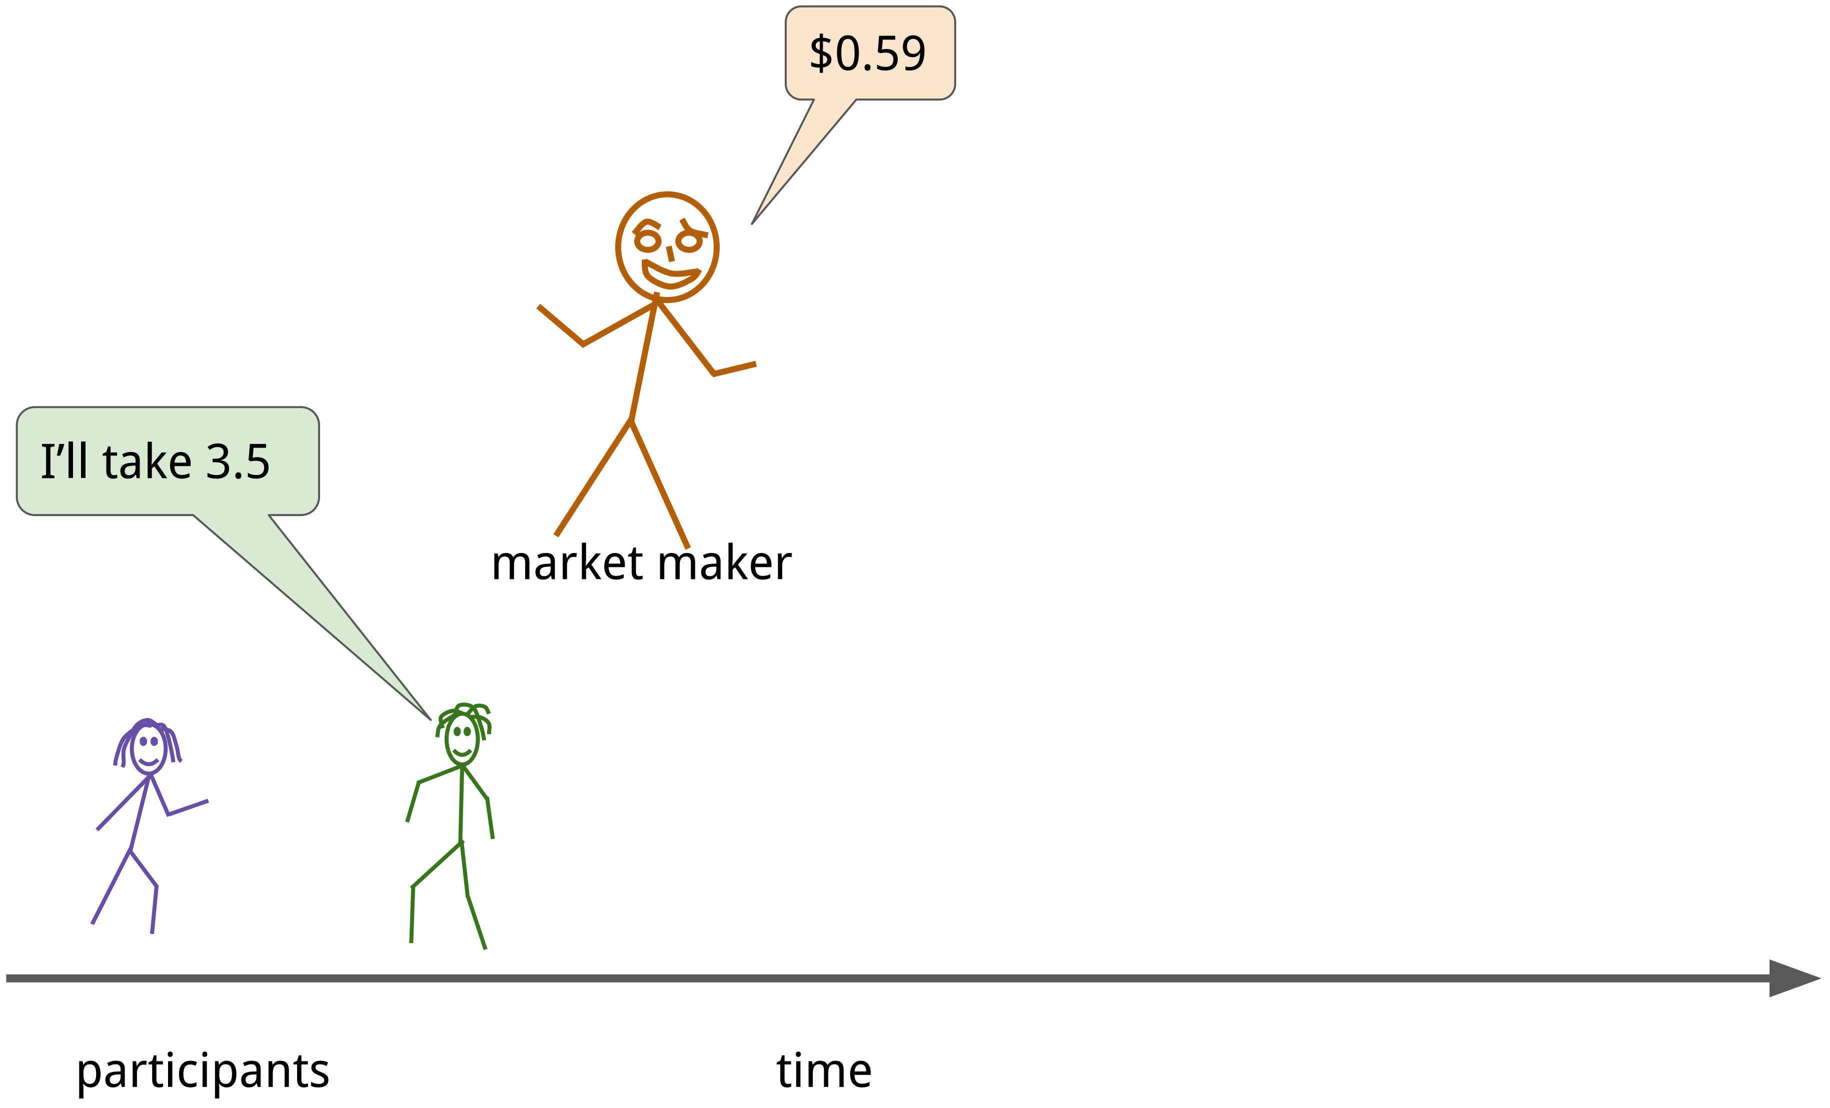

Prediction markets - dynamics

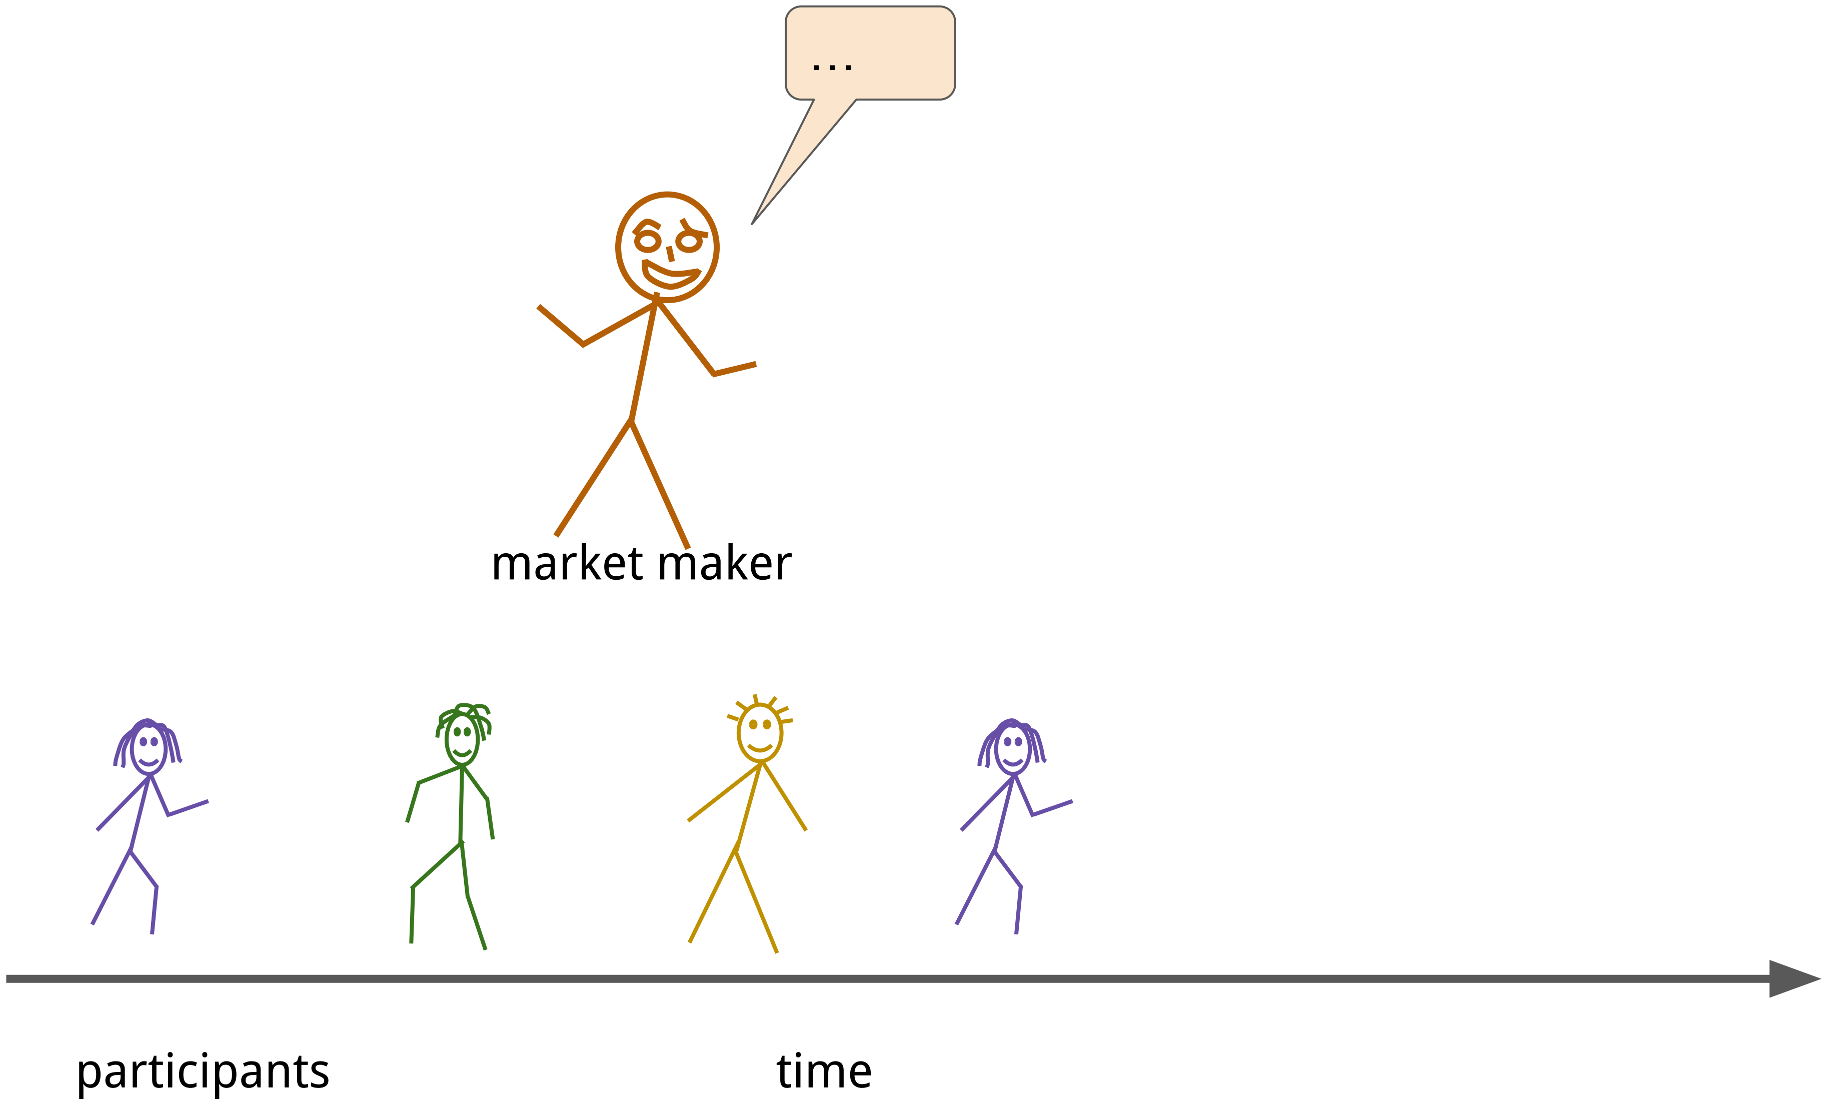

Prediction markets - dynamics

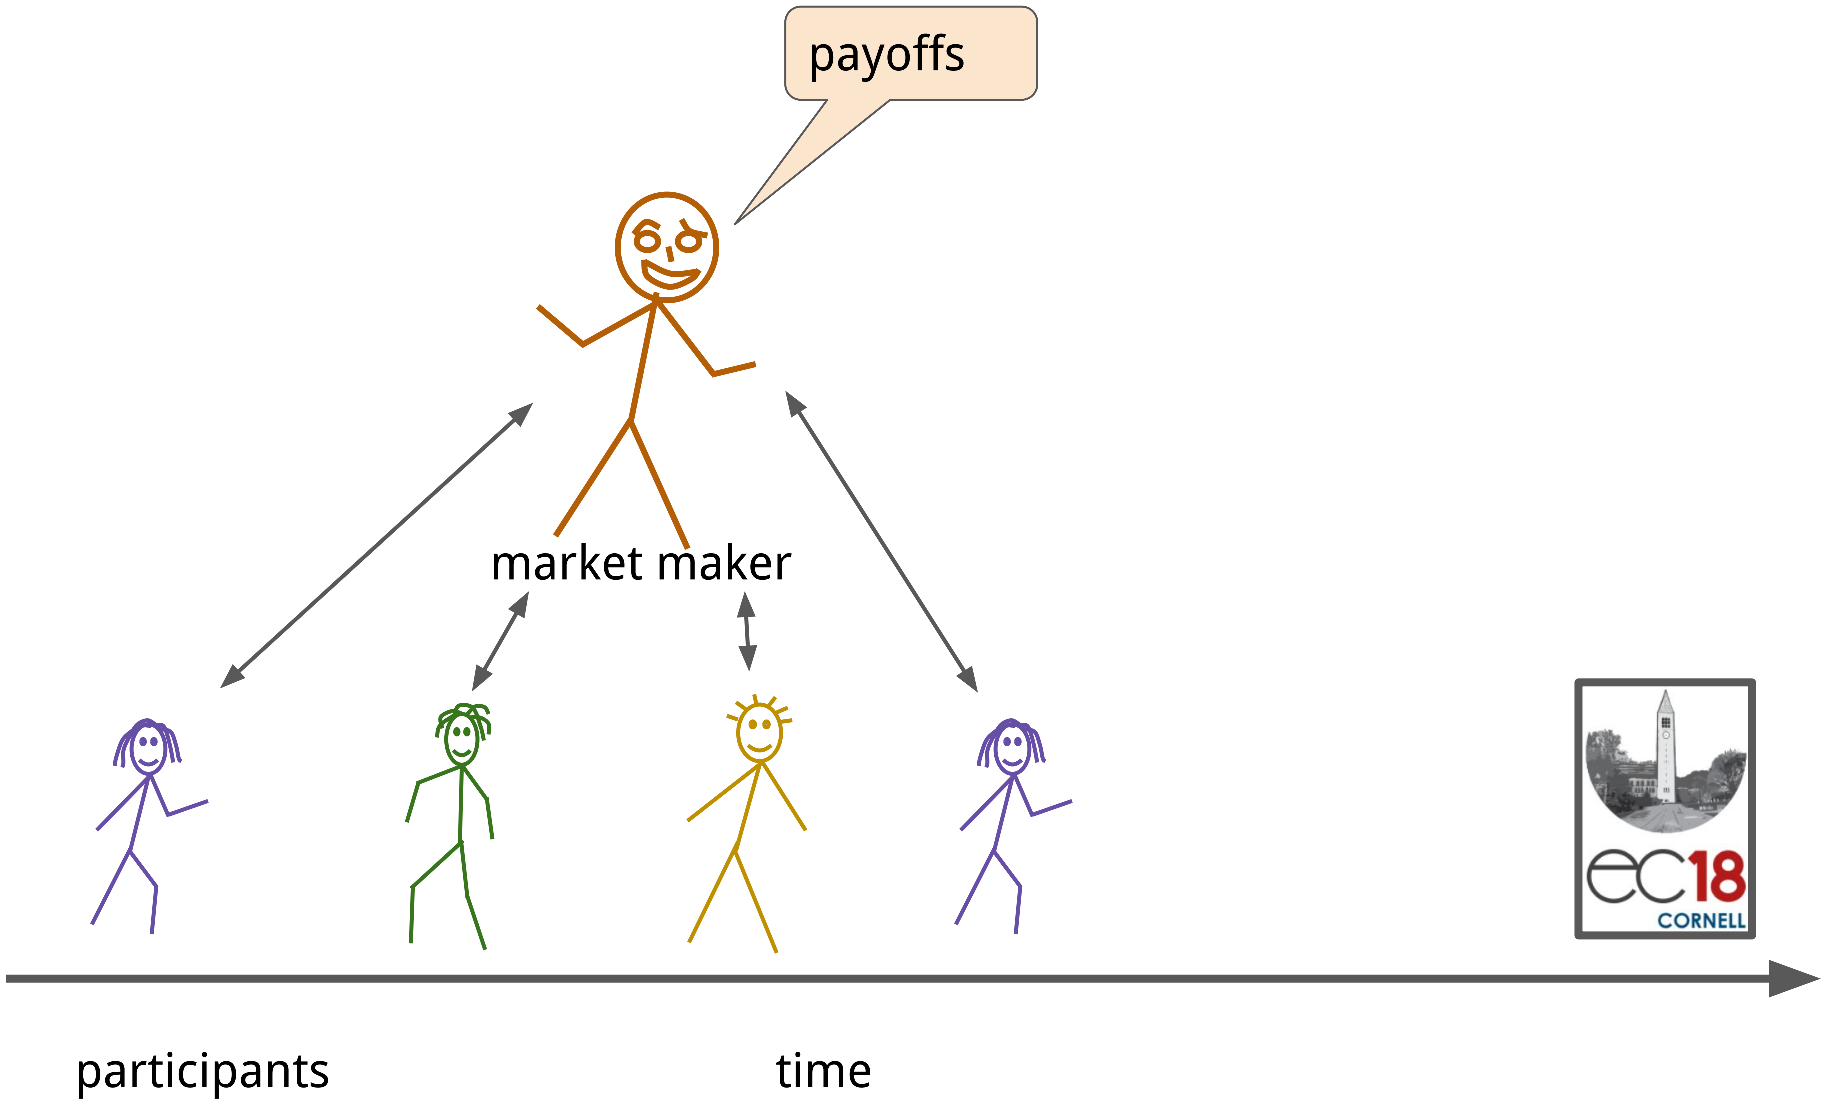

Design question

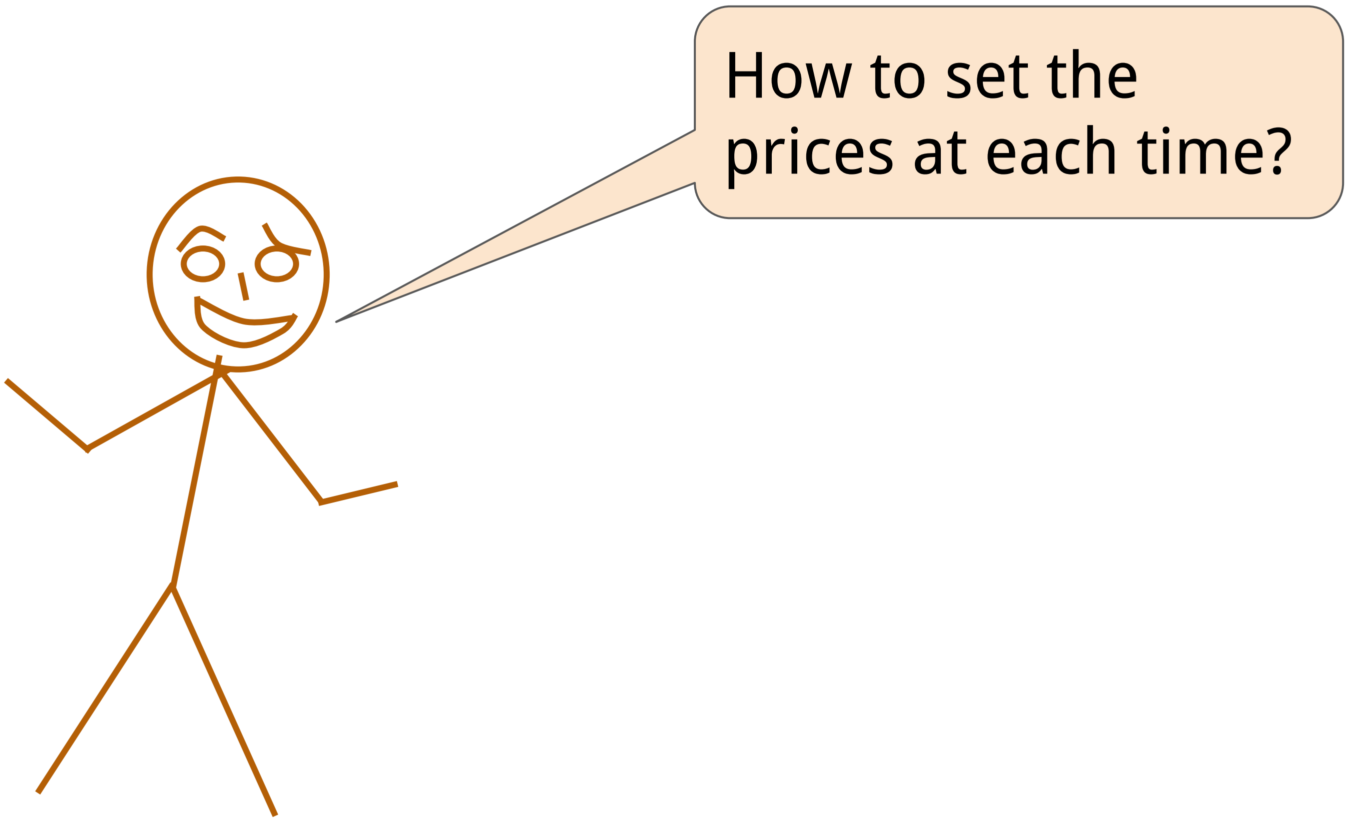

Design question

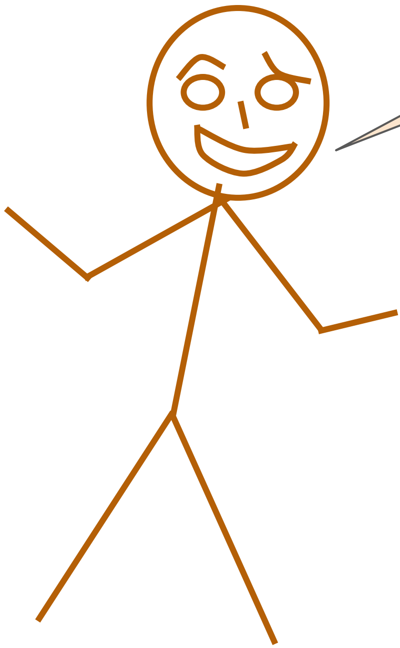

How to set the prices at each time?

Convex function

C: (total shares sold) $\longrightarrow$ (total price paid)

Design question

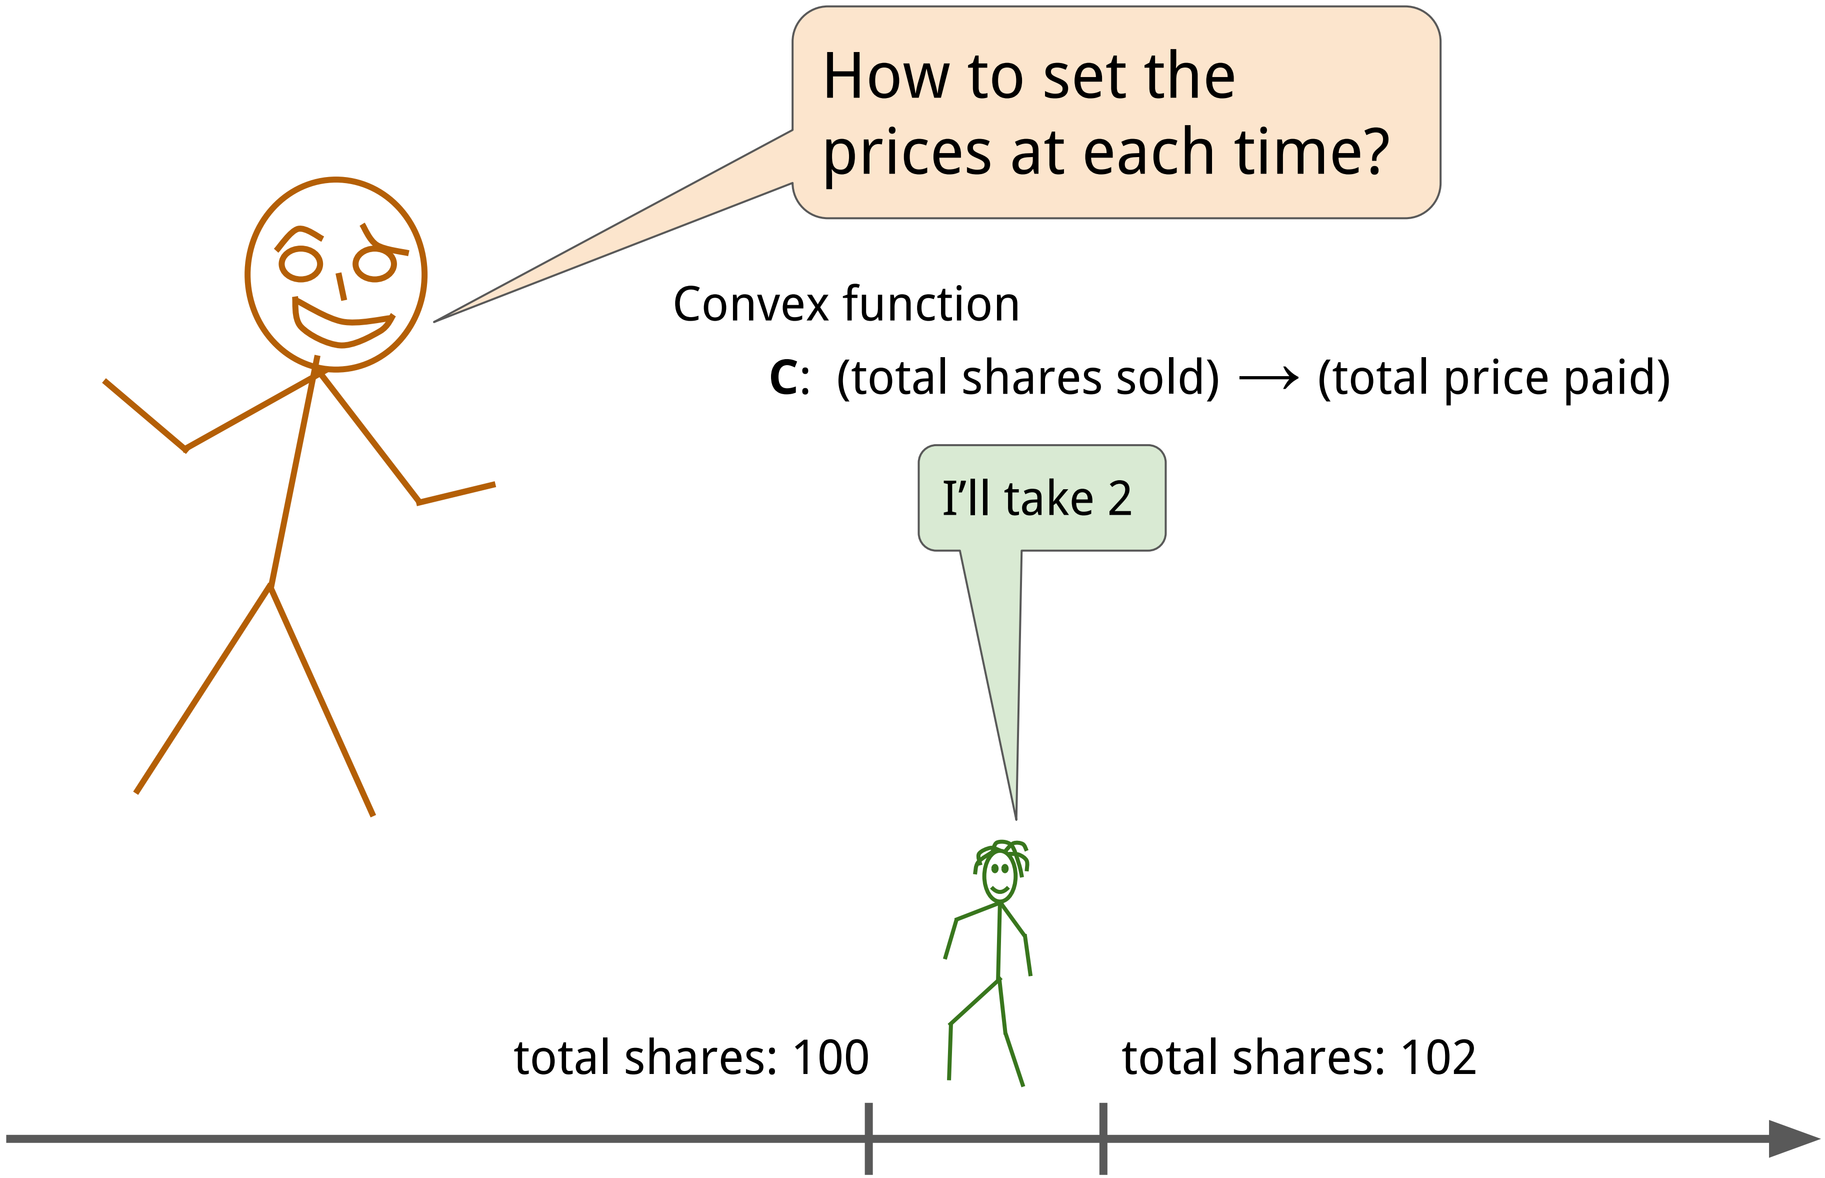

Design question

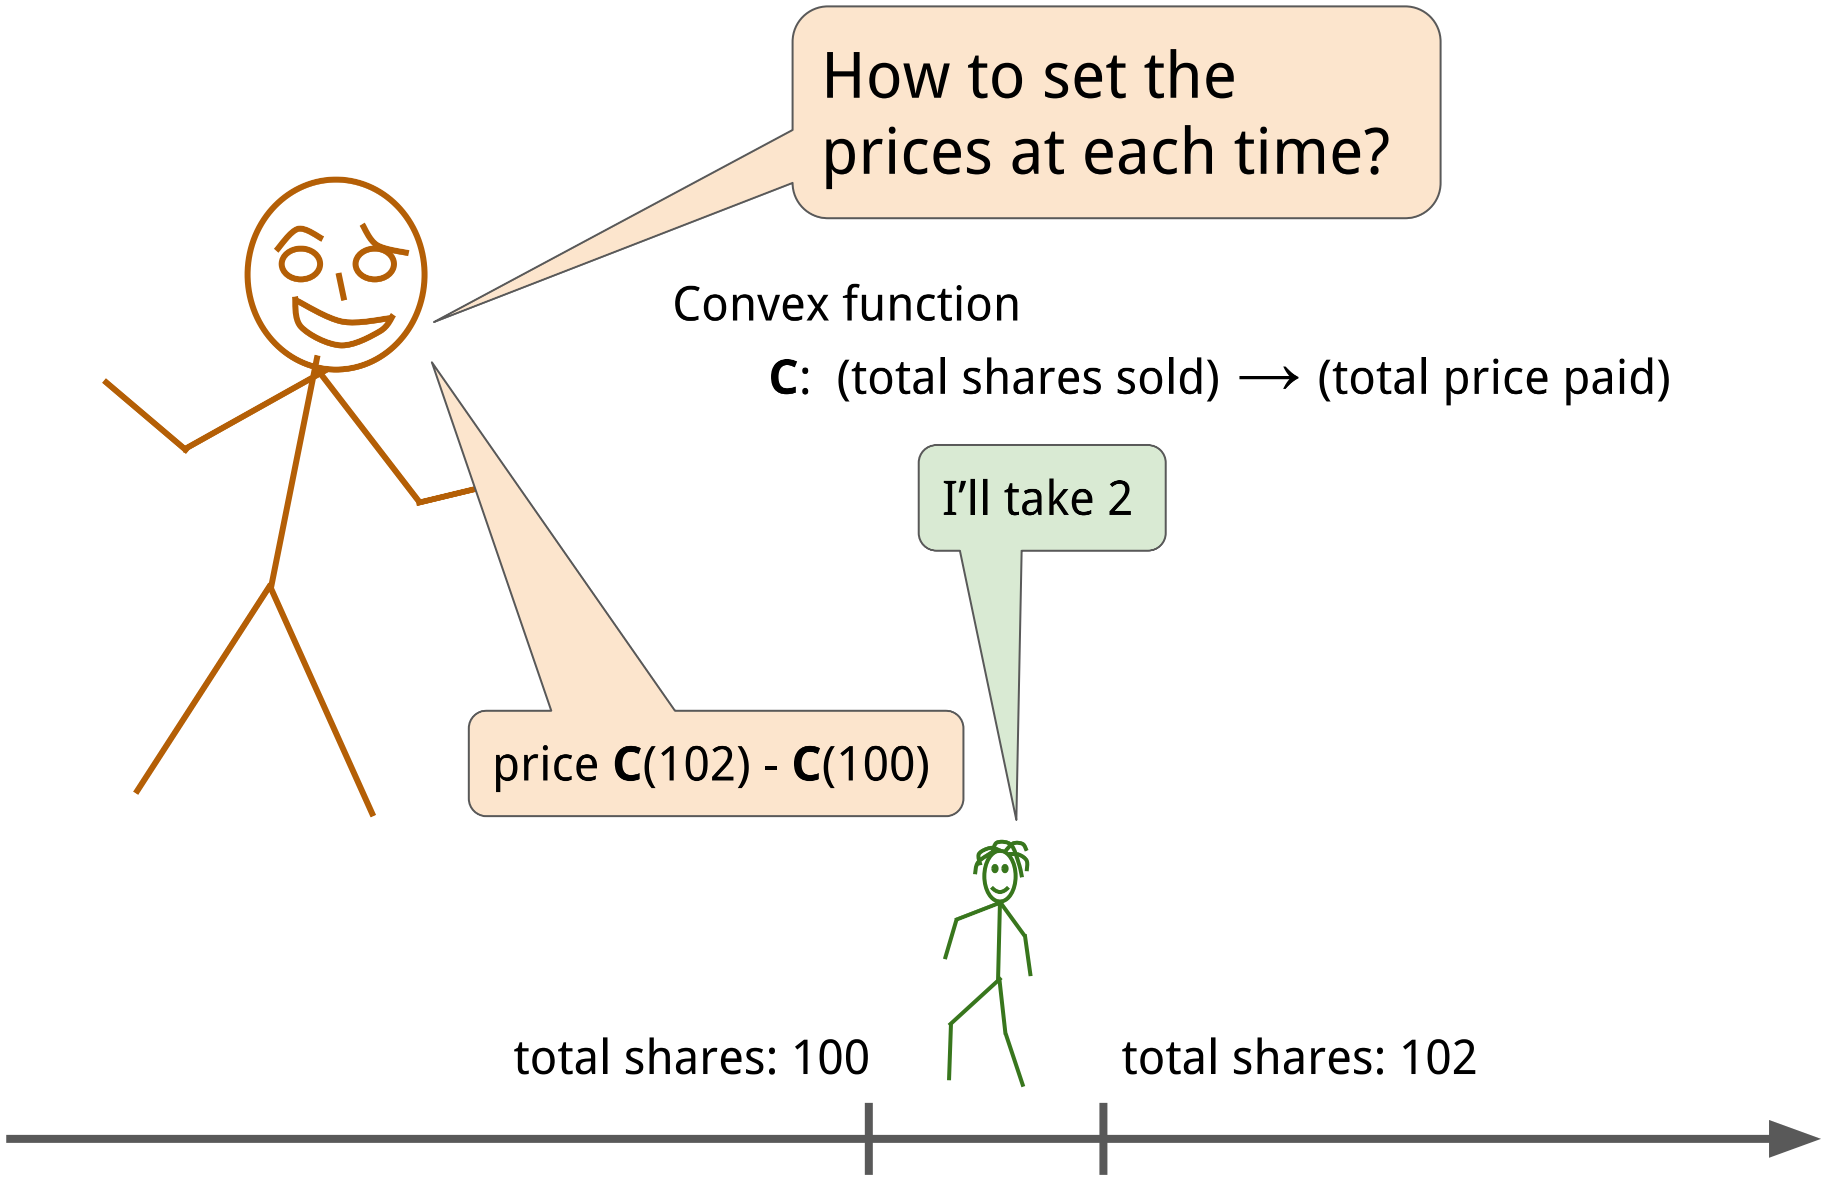

The cost function

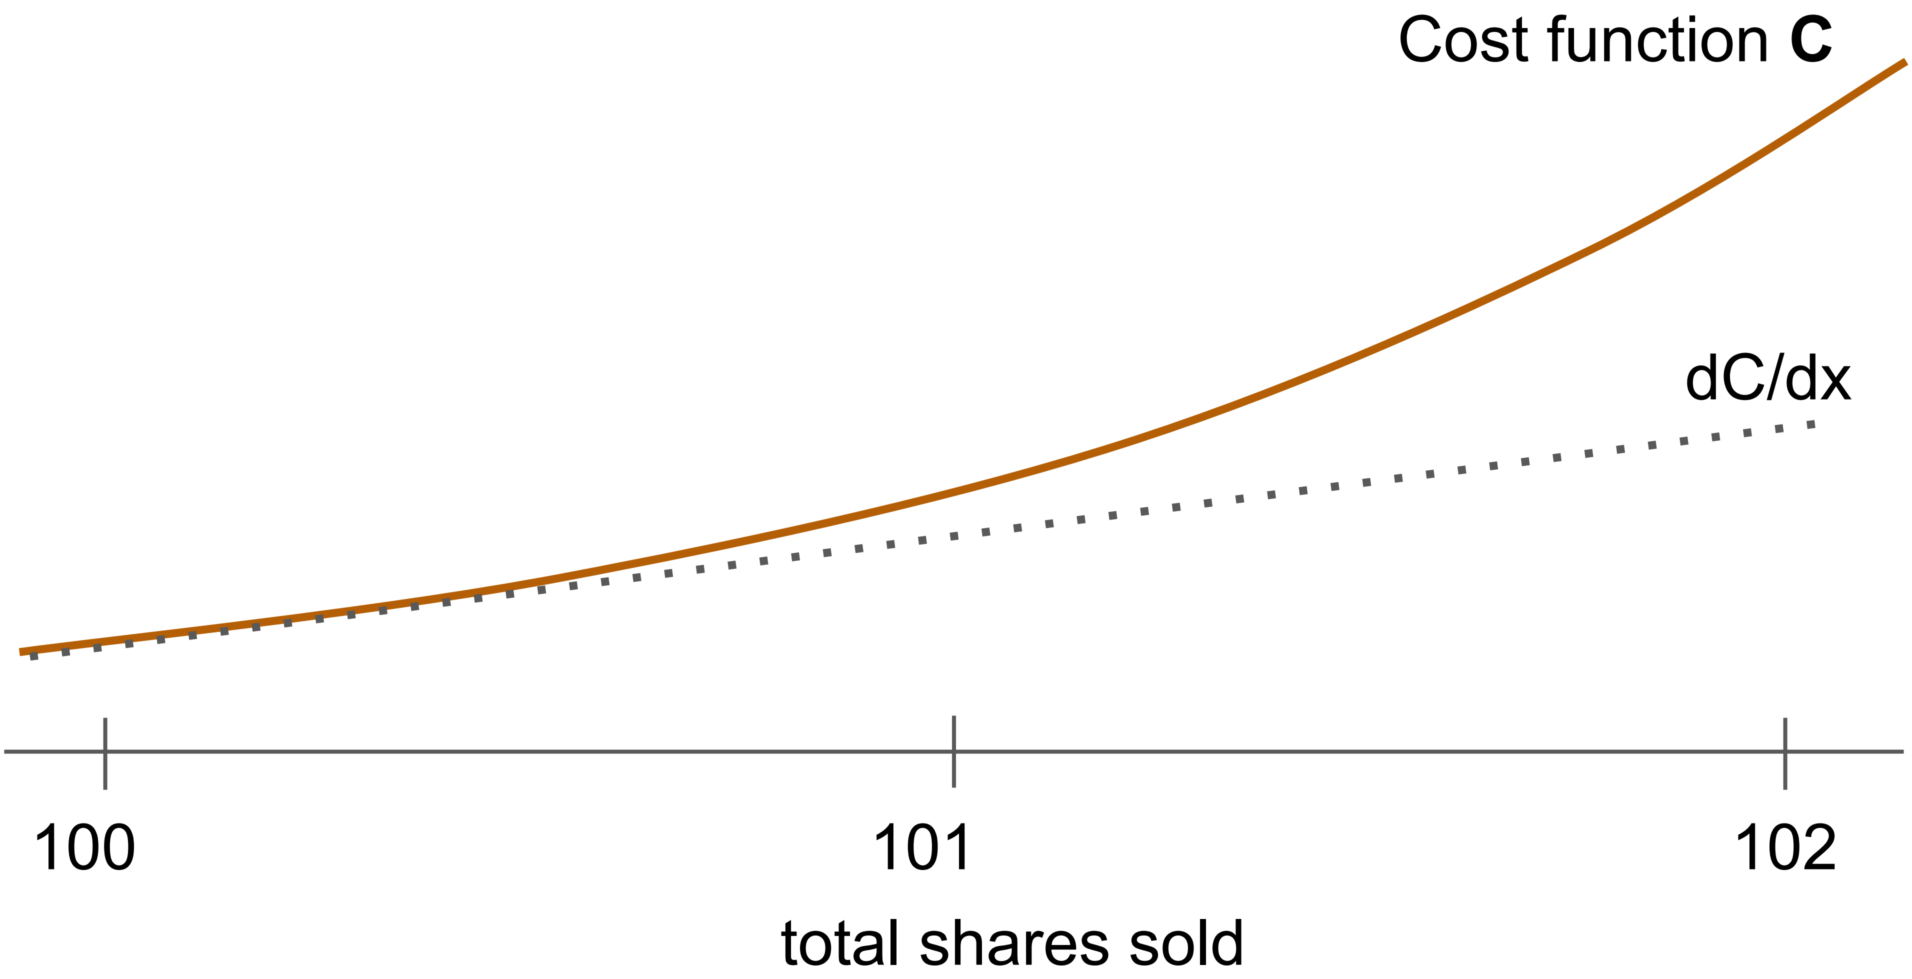

The cost function

instantaneous price = dC/dx
= Pr[event].

convexity ⇐⇒ price ↑ when you buy

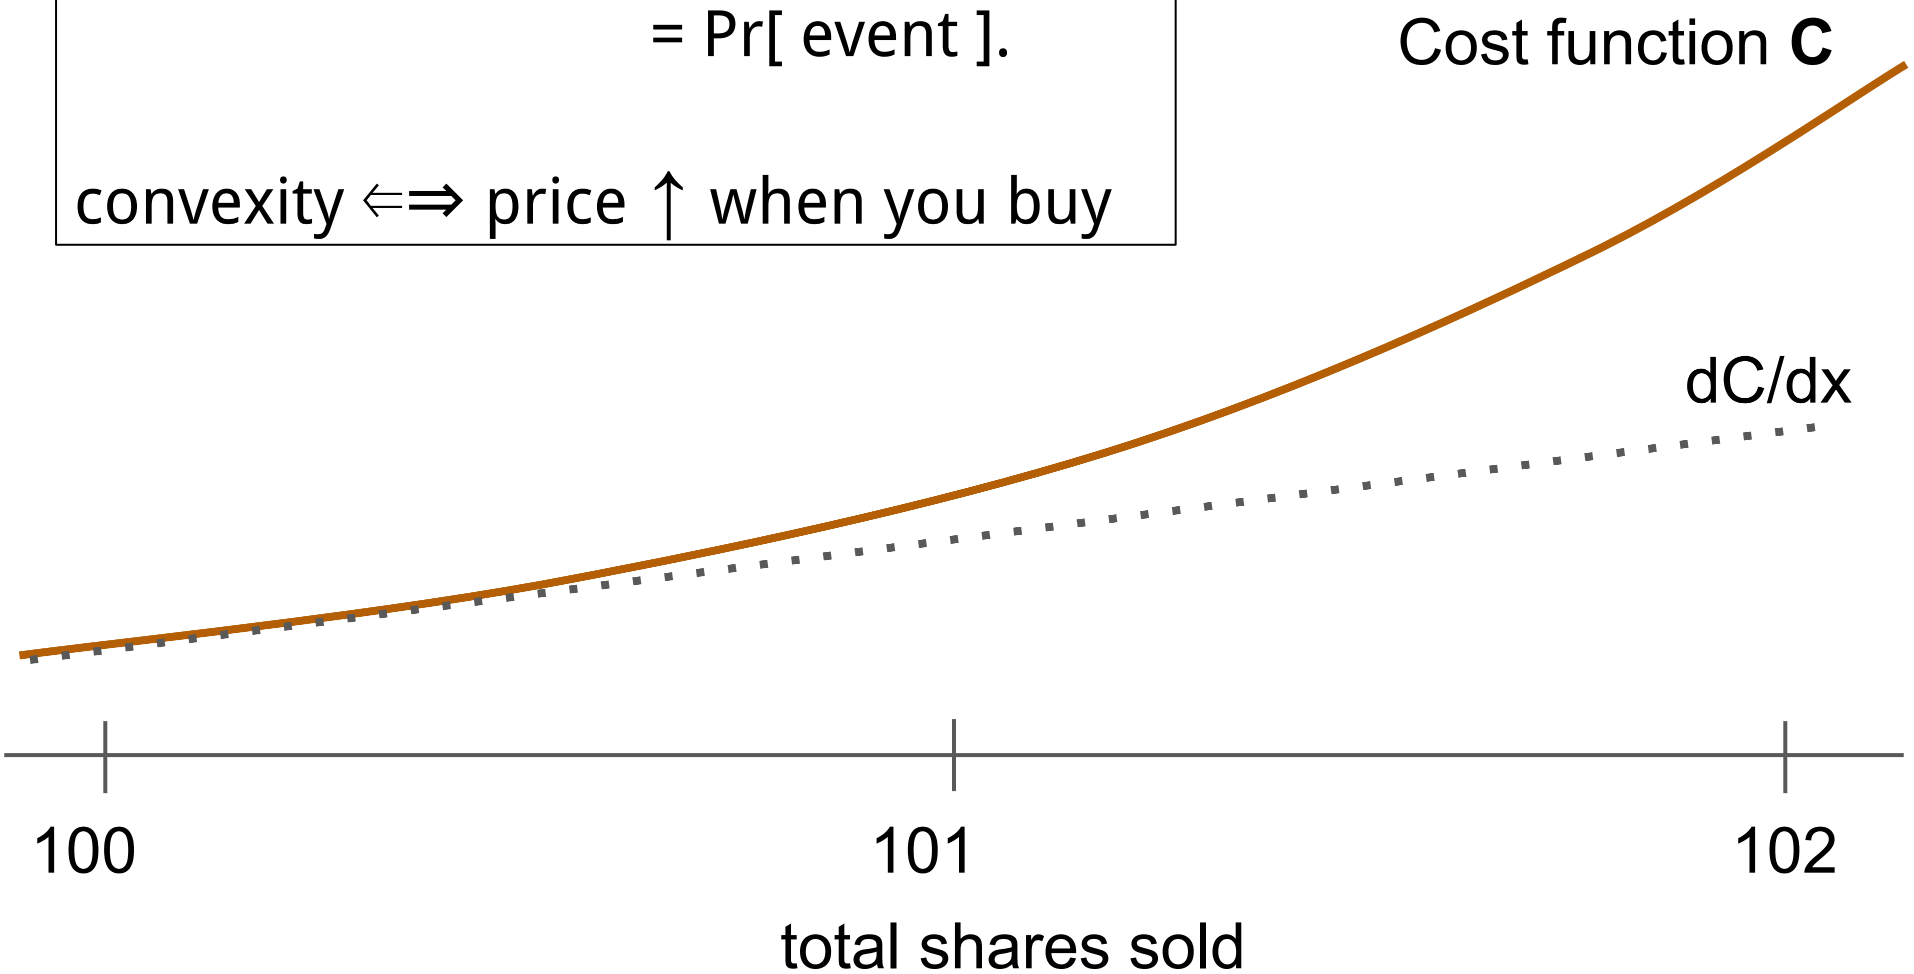

Key idea: price sensitivity λ

How quickly do prices respond to trades?

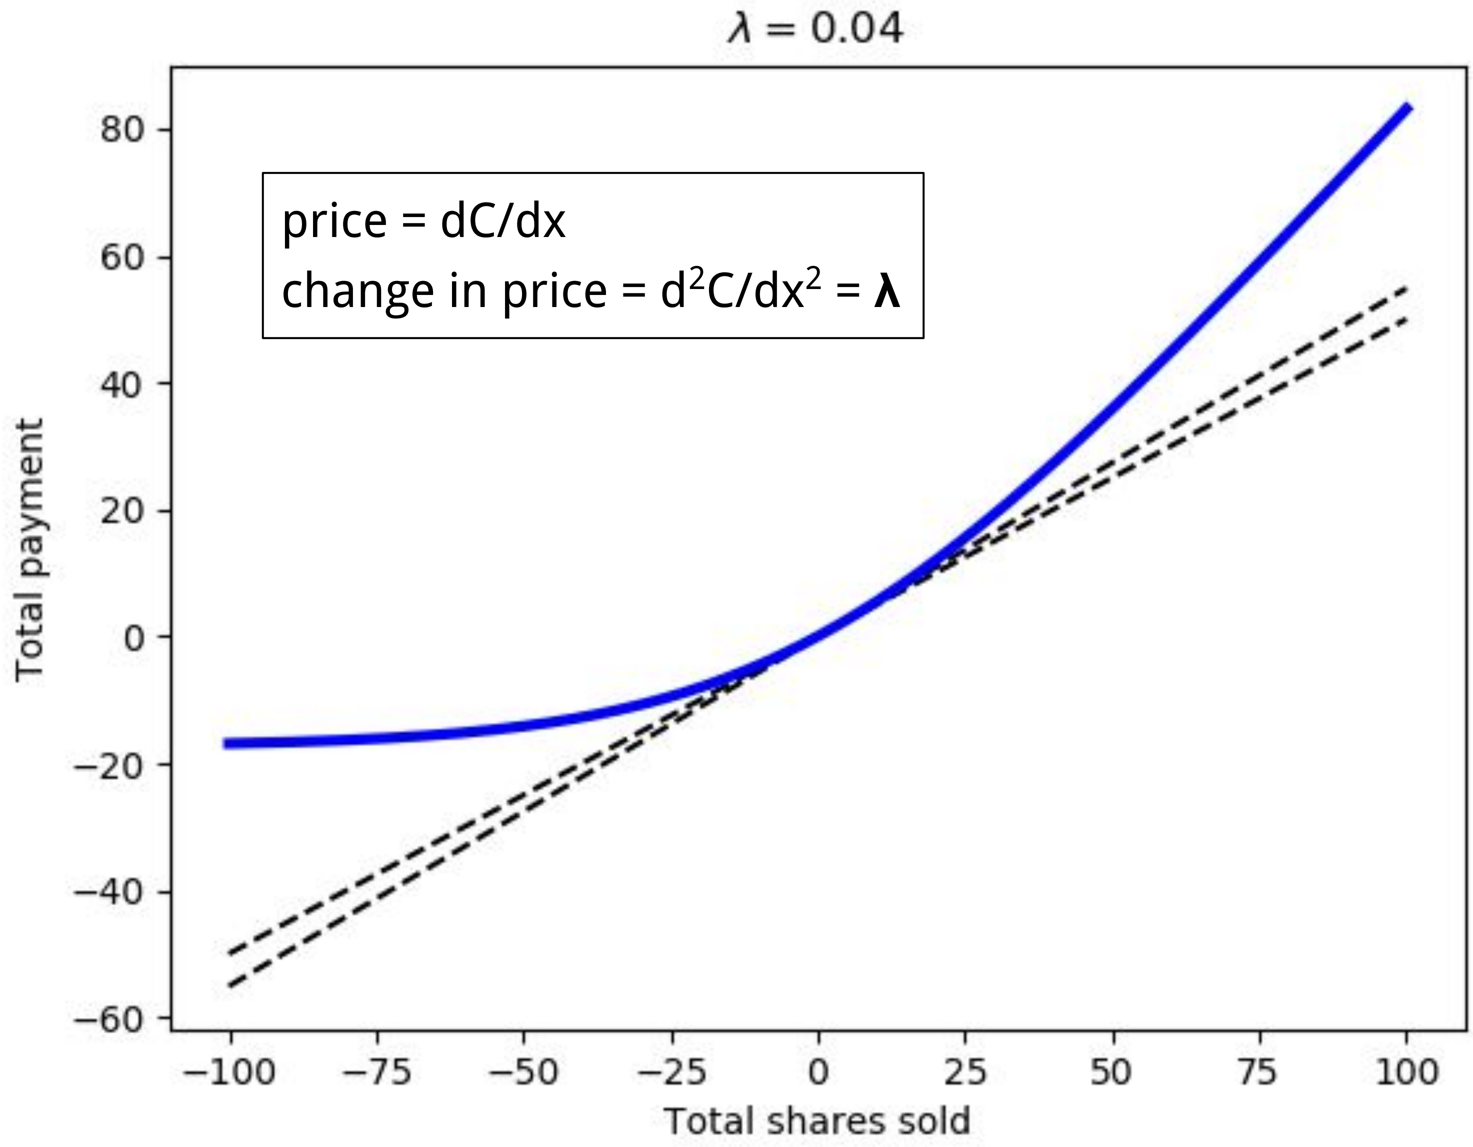

Key idea: price sensitivity λ

How quickly do prices respond to trades?

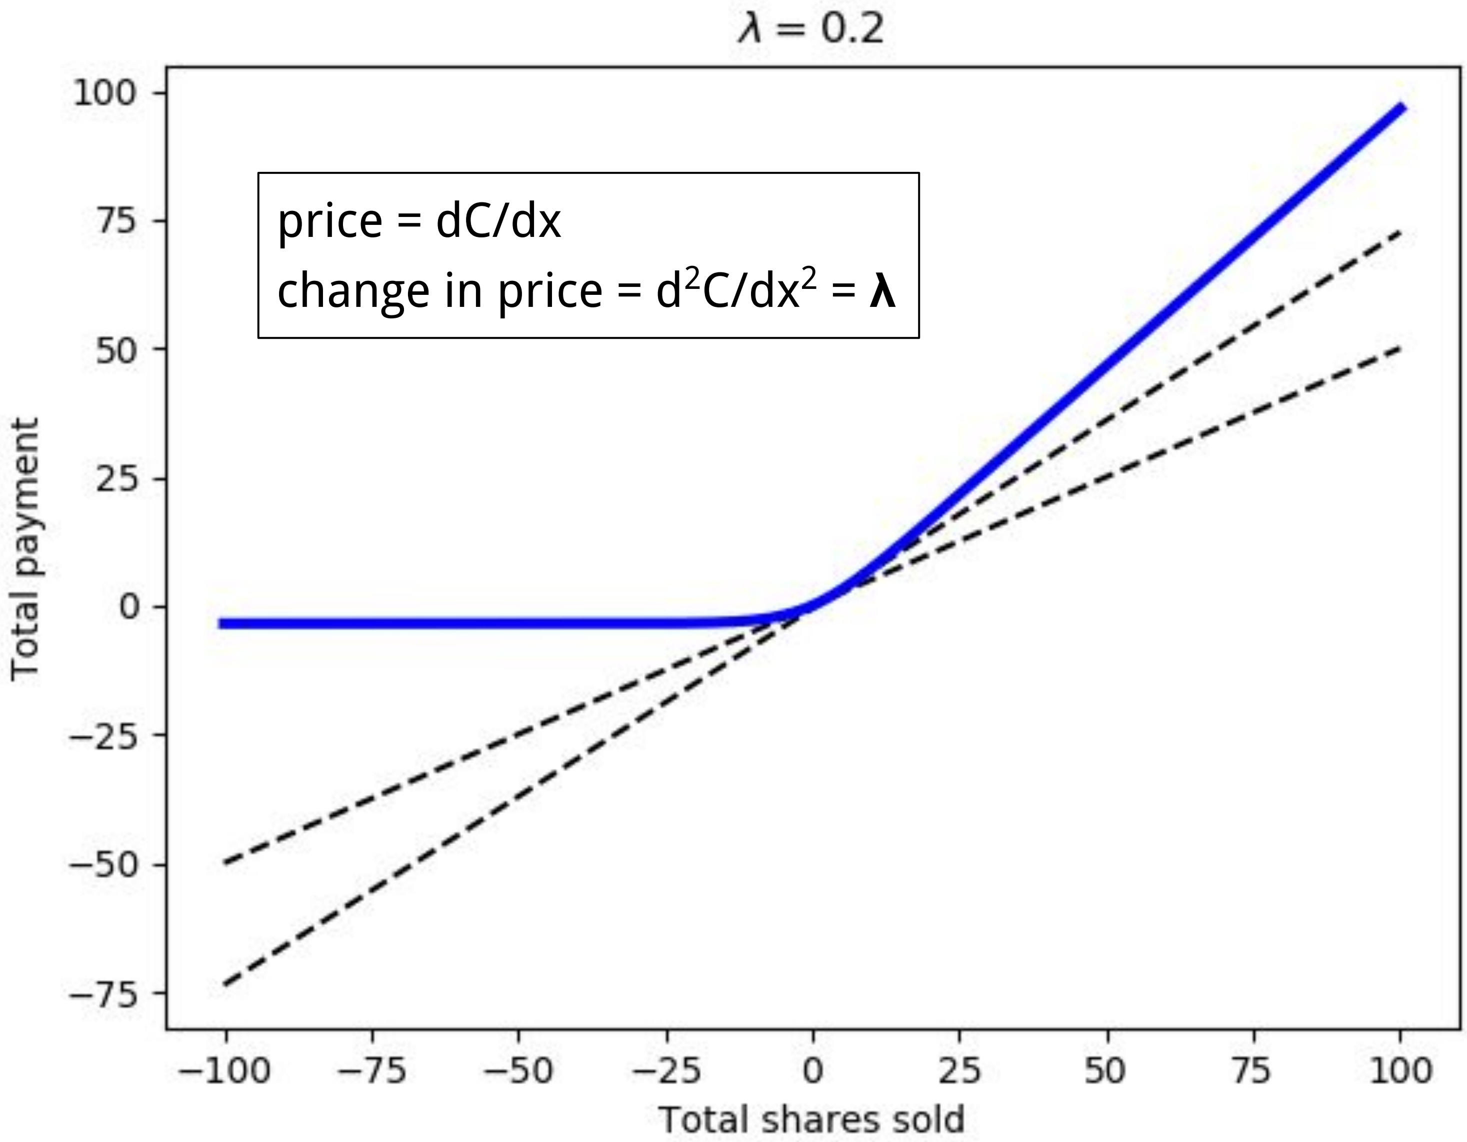

Key idea: price sensitivity λ

How quickly do prices respond to trades?

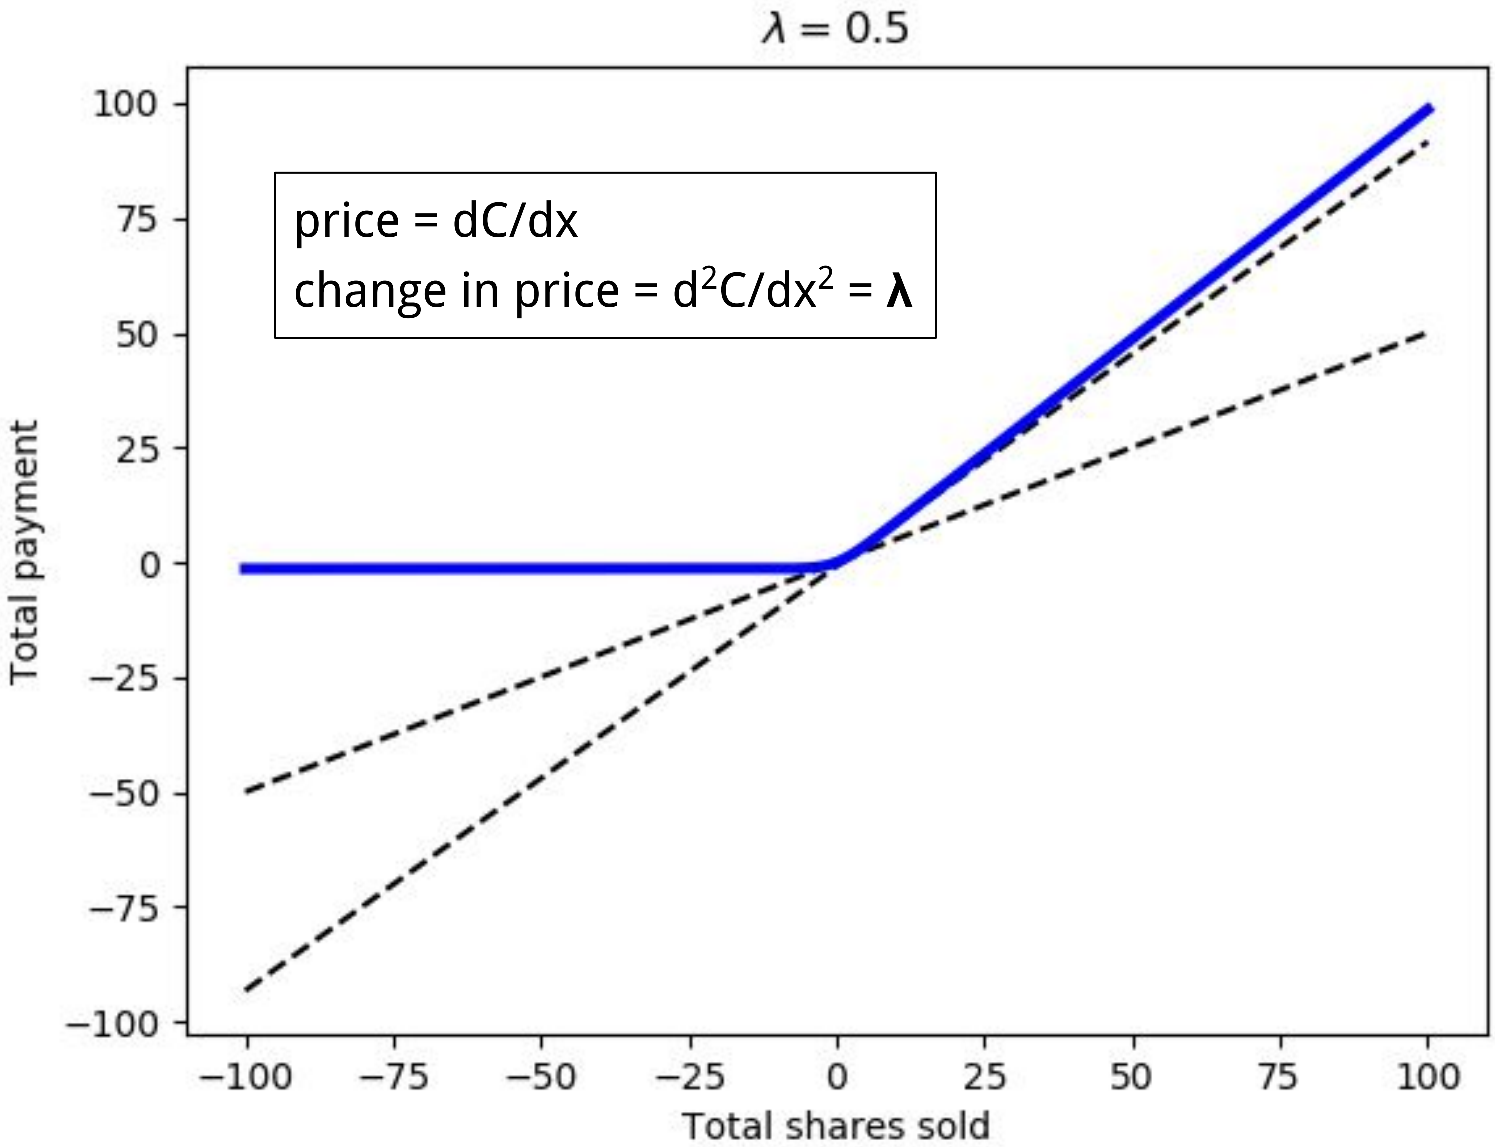

Worst Case Loss ≈ 1 / λ

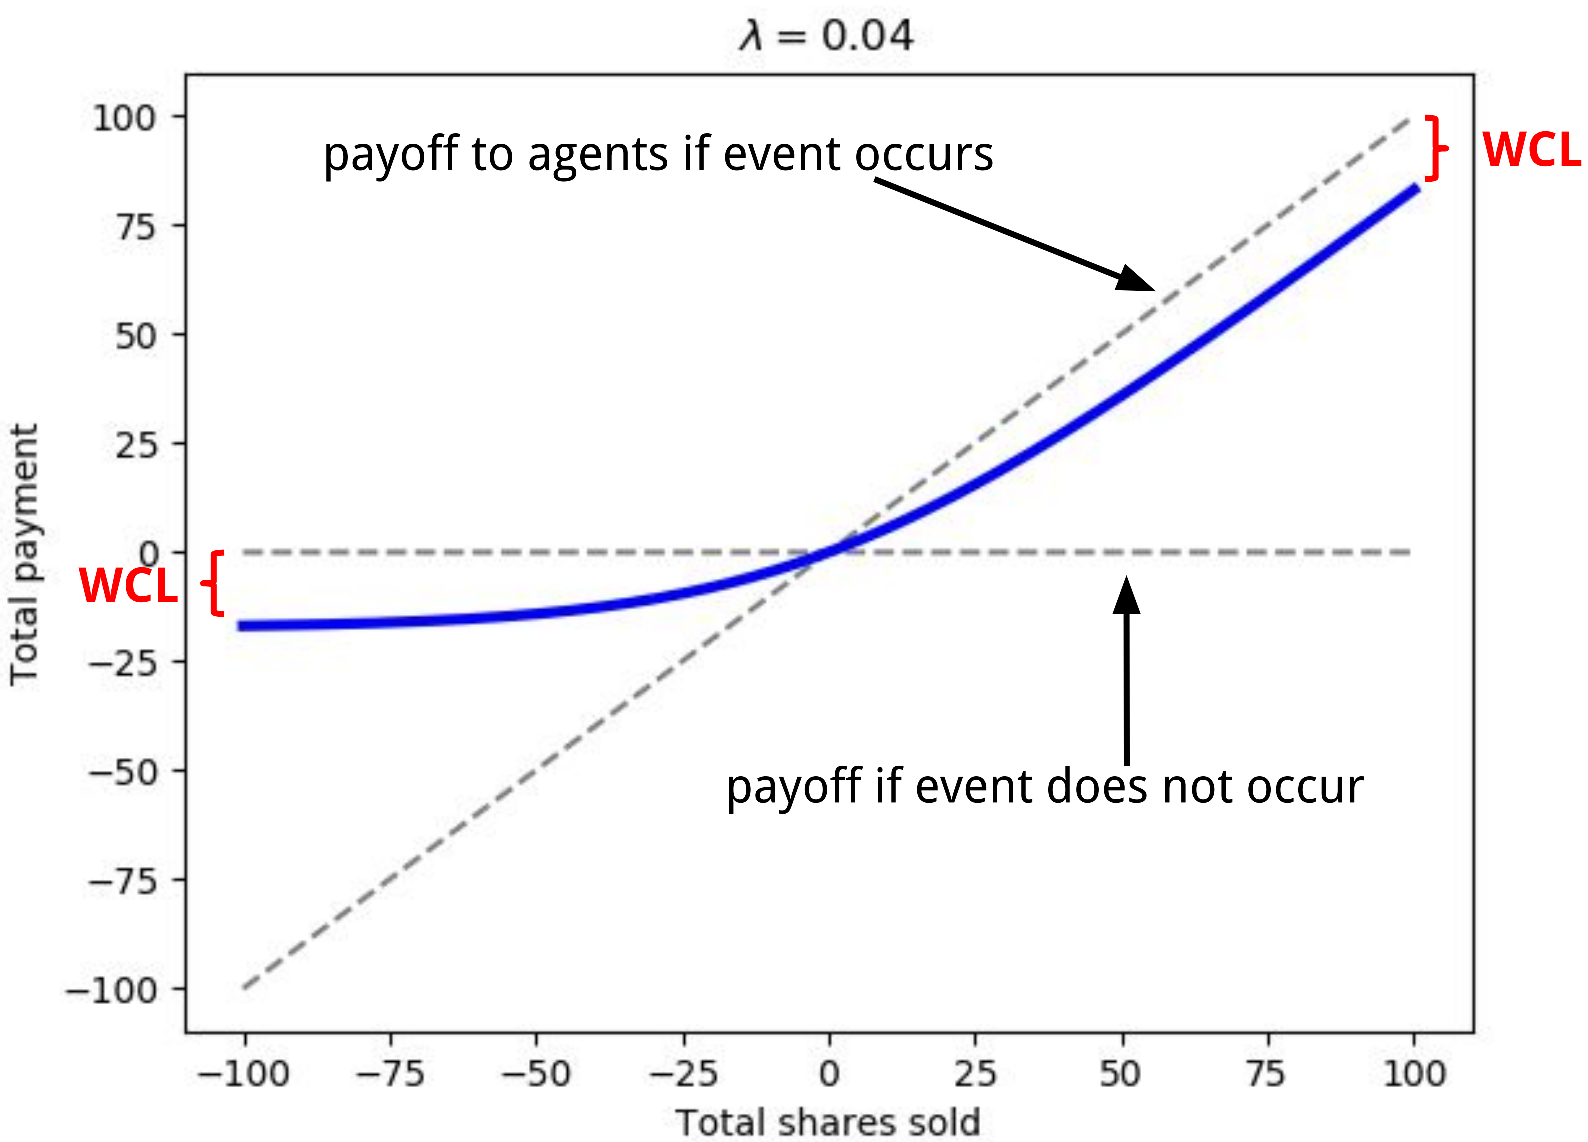

Worst Case Loss ≈ 1 / λ

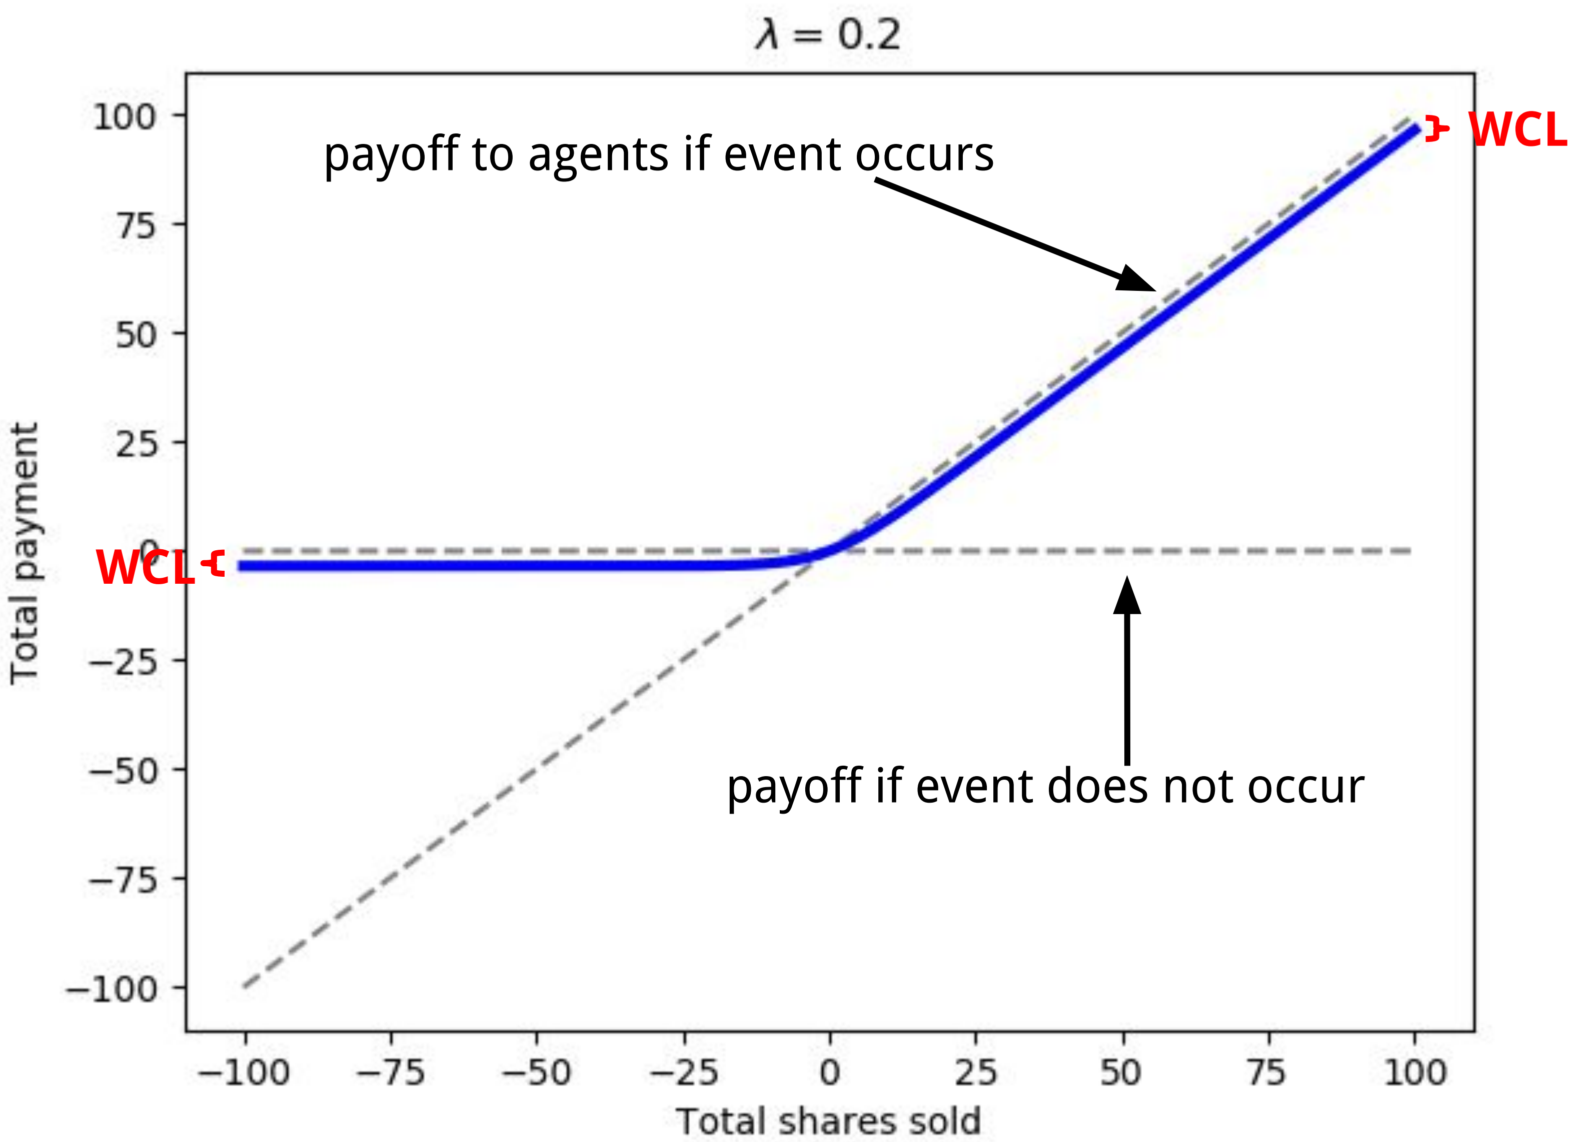

Worst Case Loss ≈ 1 / λ

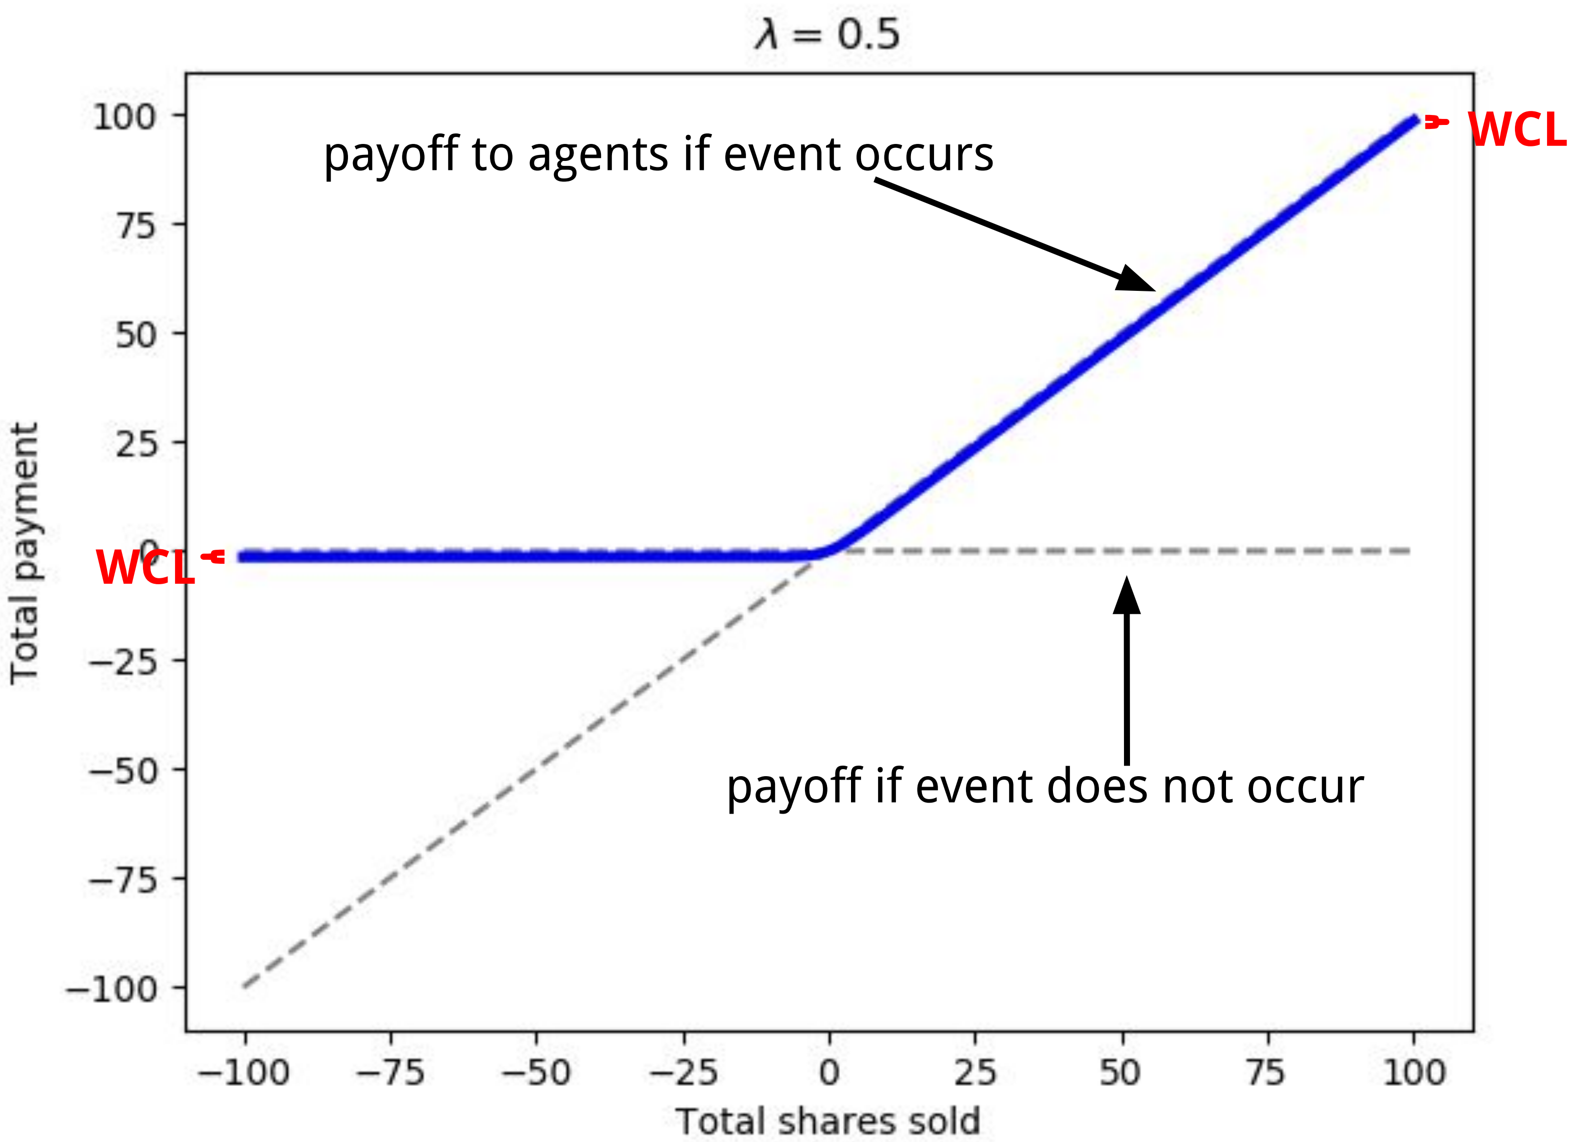

Outline

A. Cost function based prediction markets

→ B. Summary of results and prior work

C. Construction

Privacy in markets: history

- *Waggoner, Frongillo, Abernethy. NIPS 2015*
  - includes a proposal for private prediction markets
  - focused on ML extensions; private markets not well explained

- *Cummings, Pennock, Wortman Vaughan. EC 2016*
  - every private prediction market has **unbounded financial loss**

- *Frongillo, Waggoner. 2018 (manuscript)*
  - modified market achieving **bounded** loss (with unbounded participants)
  - idea 1: transaction fee
  - idea 2: adaptive **price sensitivity** (liquidity)

Private prediction markets (with unbounded loss)

Participant arrives, makes a trade, then we add noise.

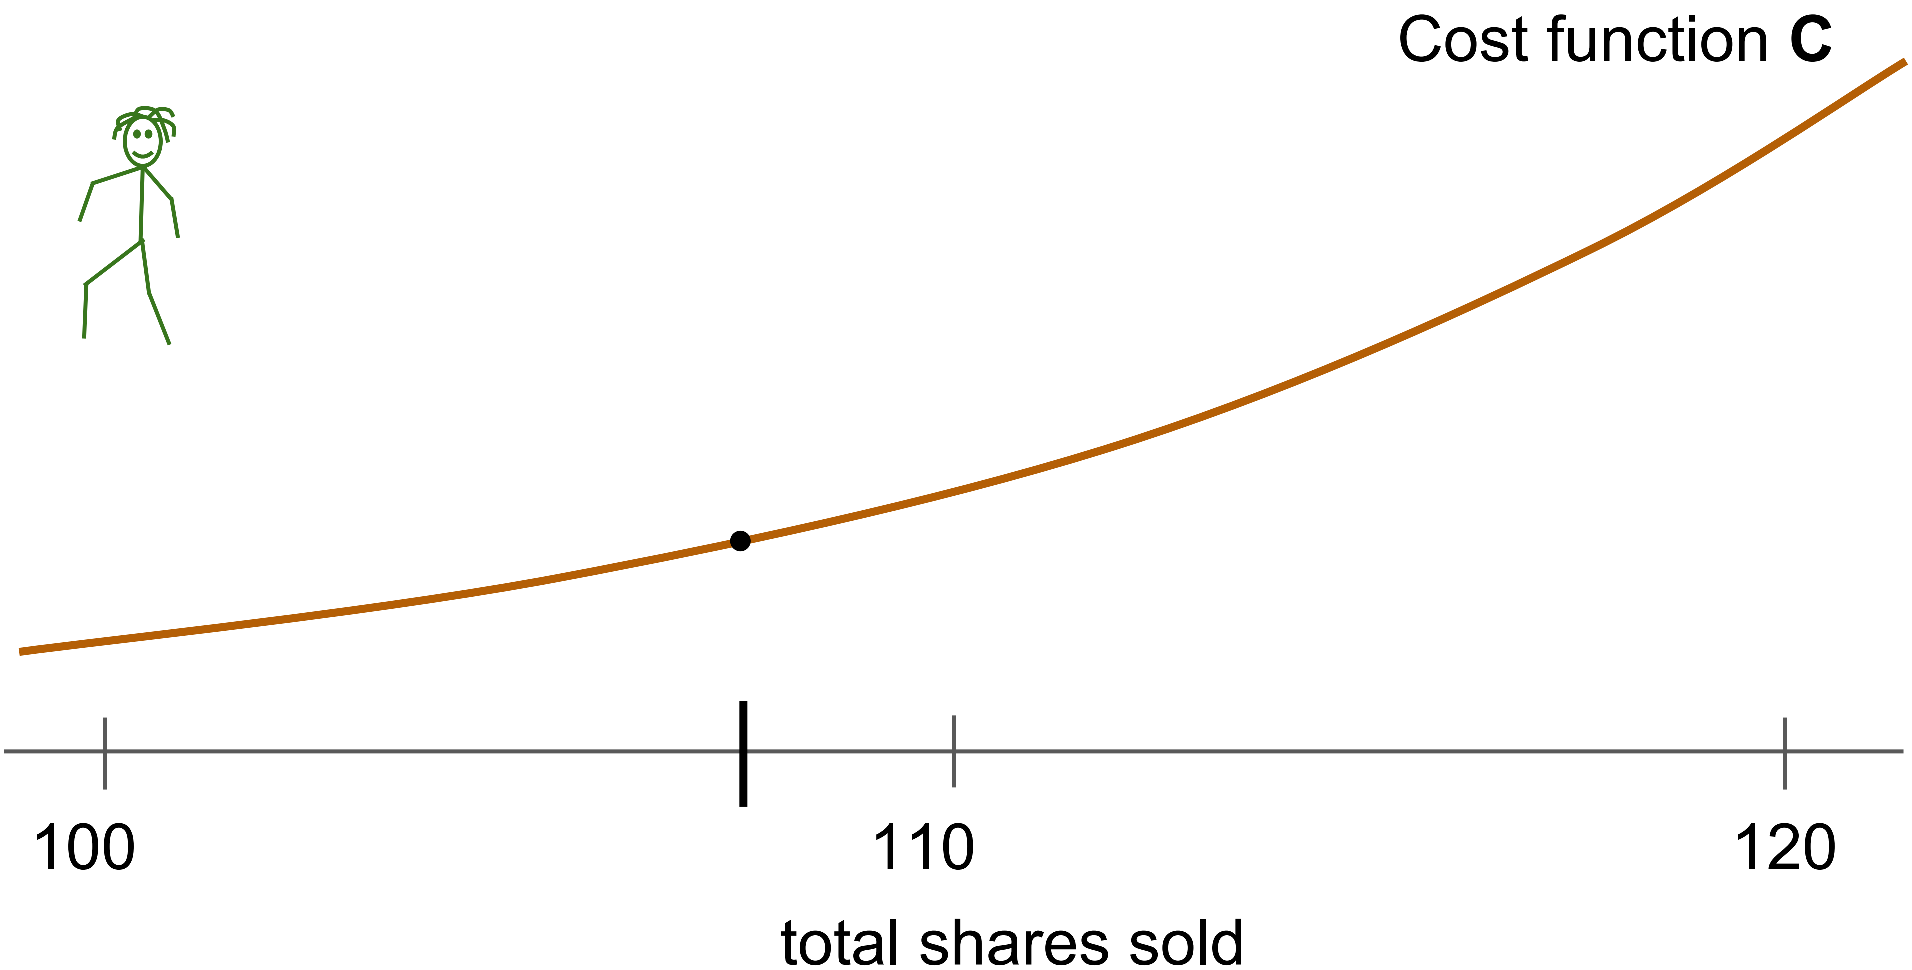

Private prediction markets (with unbounded loss)

Participant arrives, makes a trade, then we add noise.

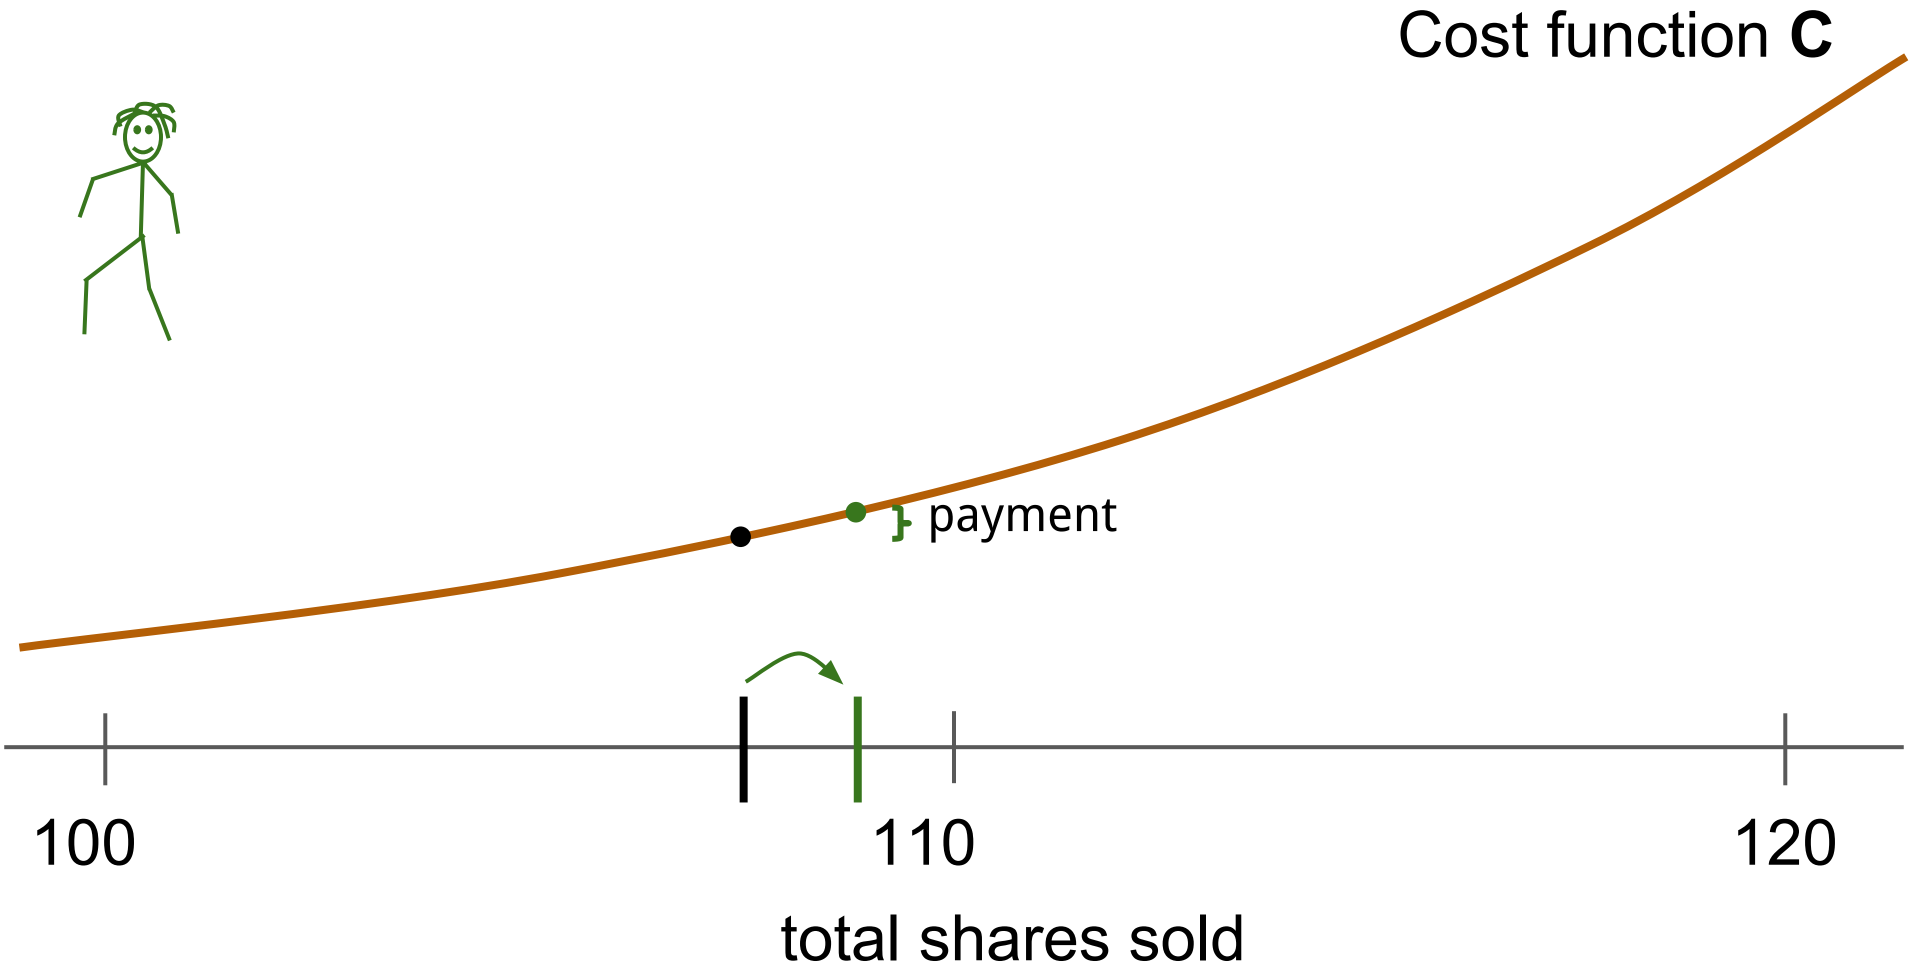

Private prediction markets (with unbounded loss)

Participant arrives, makes a trade, then we add noise.

Everyone else sees only the new market state.

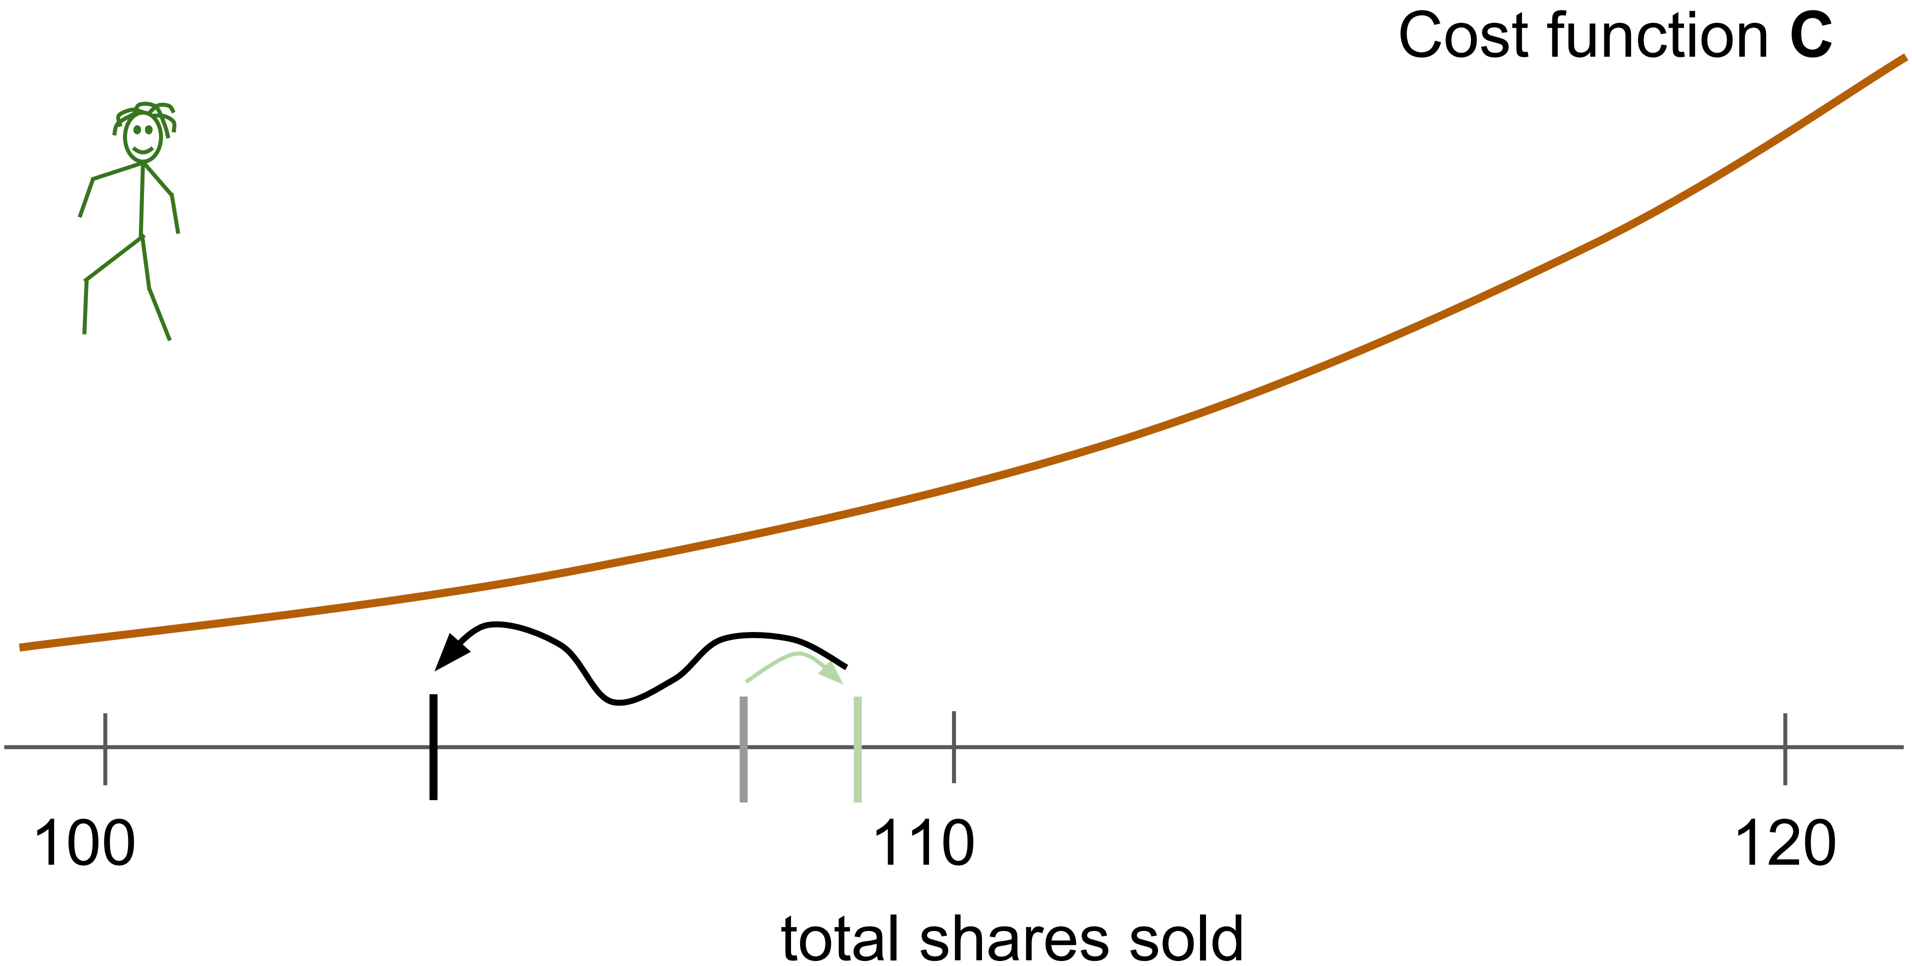

Private prediction markets (with unbounded loss)

Given privacy level ε, set amount of noise.

Then, given accuracy level α, set price sensitivity λ s.t. noise doesn't hurt accuracy.

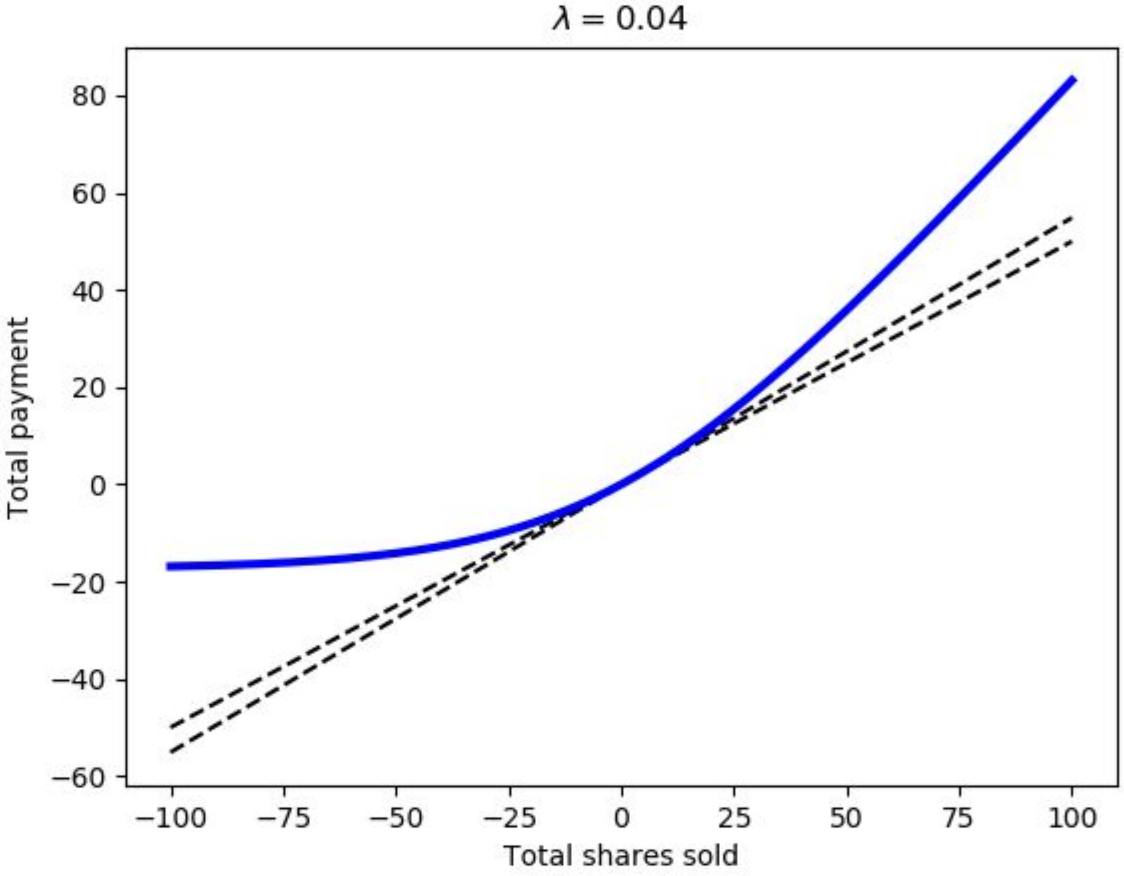

Better privacy-accuracy tradeoffs

Independent noise each step, **T total participants** ⇒ error O(sqrt(T)).

Best privacy technique ("continual observation"): add O(log T) noise each step…
… coordinated across time steps s.t. total noise is always $O(\log^2 T)$.

⇒ $\lambda = \Theta(1 / \log^2 T)$.

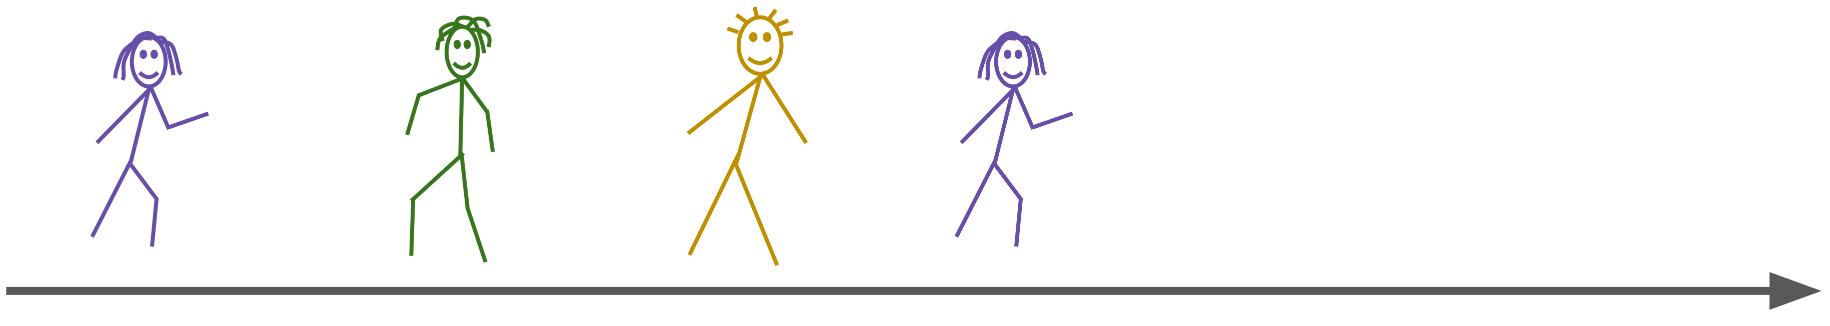

Better privacy-accuracy tradeoffs

Independent noise each step, **T total participants** ⇒ error O(sqrt(T)).

Best privacy technique (“continual observation”): add O(log T) noise each step…
… coordinated across time steps s.t. total noise is always $O(\log^2 T)$.

⇒ $\lambda = \Theta(1 / \log^2 T)$.

Interpretation: “noise trader” makes random purchases after each arrival;
total loss = loss of market maker + loss of noise trader.

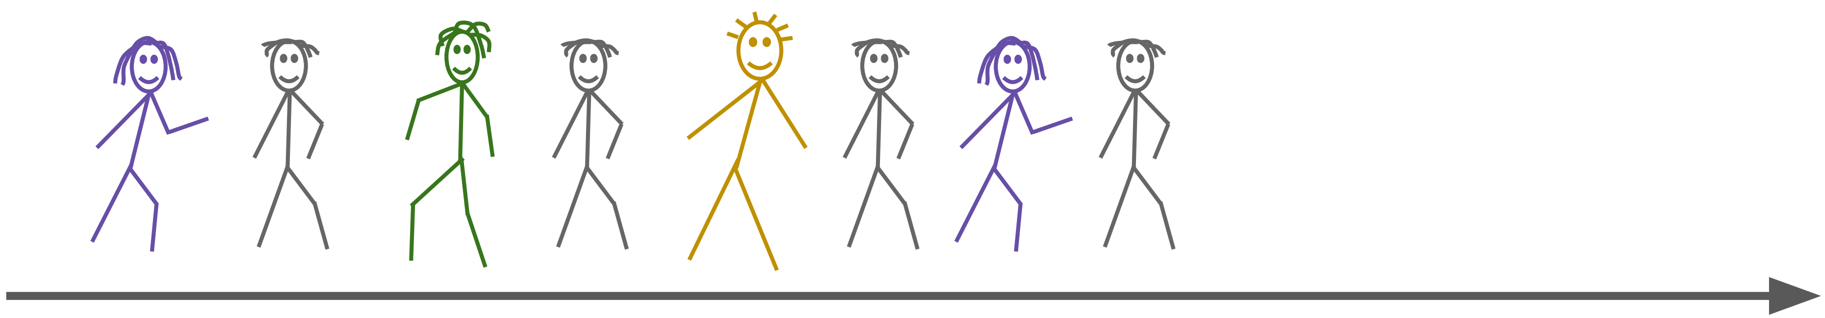

Better privacy-accuracy tradeoffs

Independent noise each step, **T total participants** ⇒ error O(sqrt(T)).

Best privacy technique (“continual observation”): add O(log T) noise each step…
… coordinated across time steps s.t. total noise is always O($\log^2$ T).

⇒ $\lambda = \Theta(1 / \log^2 T)$.

Interpretation: “noise trader” makes random purchases after each arrival;
total loss = loss of market maker + loss of noise trader.

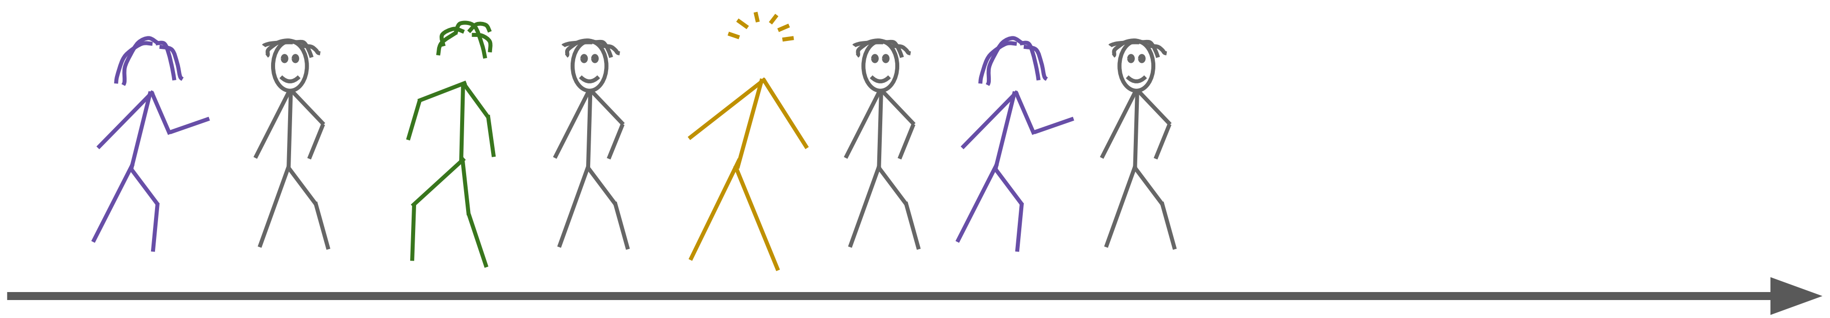

Private prediction markets (with unbounded loss)

Theorem (based on Waggoner, Frongillo, Abernethy 2015)

The private market achieves:

- ε-differential privacy
- α-precision with high probability (noise affects prices by at most α)
- incentive to participate (if prices are wrong, an agent can profit by changing them)

all with

$\lambda = \Theta(1 / \log^2 T)$.

(So about $\log^2 T$ participants coordinate a useful prediction.)

Private prediction markets (with unbounded loss)

Theorem (based on Waggoner, Frongillo, Abernethy 2015)

The private market achieves:

- ε-differential privacy
- α-precision with high probability (noise affects prices by at most α)
- incentive to participate (if prices are wrong, an agent can profit by changing them)

all with

$$\lambda = \Theta(1 / \log^2 T).$$

(So about $\log^2 T$ participants coordinate a useful prediction.)

Problem: worst case loss is at least $O(\log^2 T)$ …

Private prediction markets (with unbounded loss)

Theorem (based on Waggoner, Frongillo, Abernethy 2015)

The private market achieves:

- ε-differential privacy
- α-precision with high probability (noise affects prices by at most α)
- incentive to participate (if prices are wrong, an agent can profit by changing them)

all with

$$\lambda = \Theta(1 / \log^2 T).$$

(So about $\log^2 T$ participants coordinate a useful prediction.)

Theorem (Cummings et al. 2016)

Every private cost-function based market has financial loss **unbounded in T**.

Outline

A. Cost function based prediction markets

B. Summary of results and prior work

→ C. Construction

Initial approach: add a transaction fee

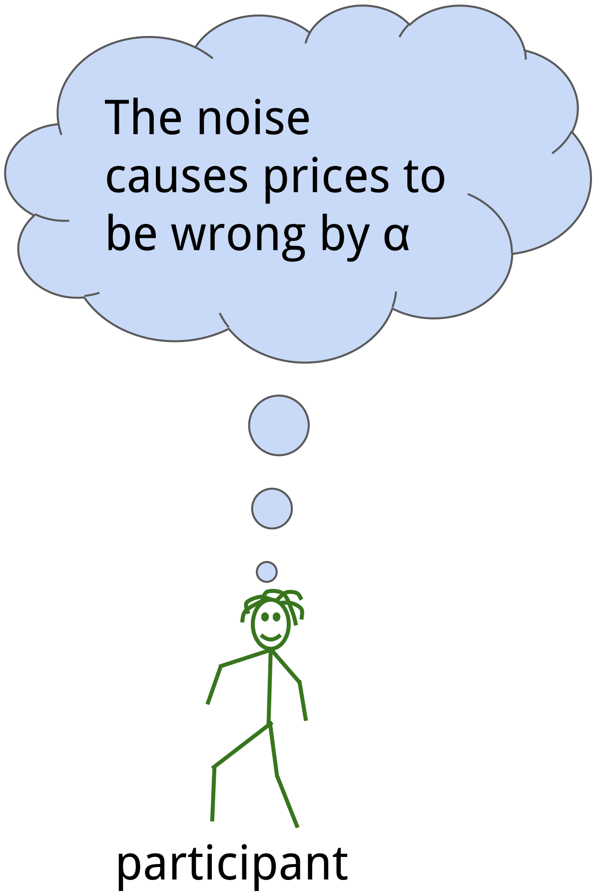

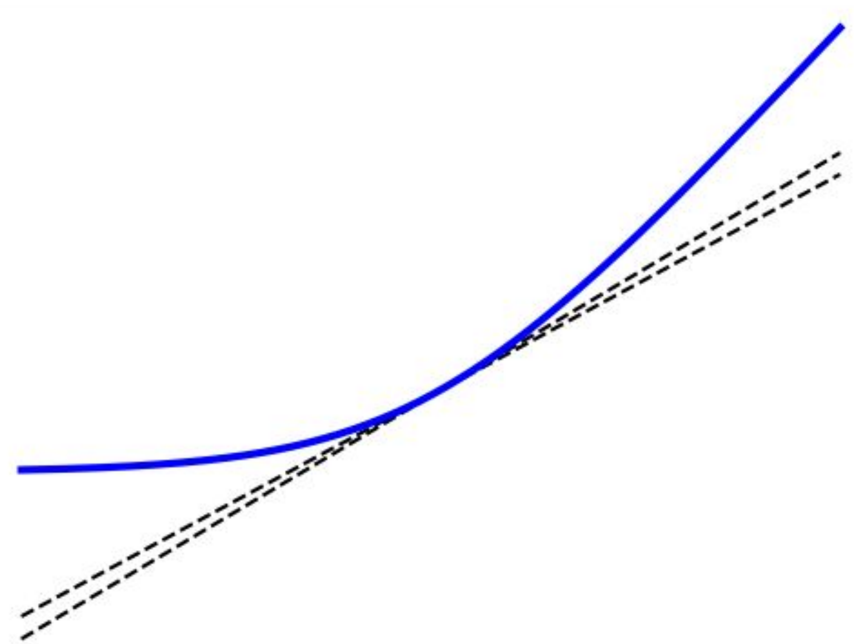

Initial approach: add a transaction fee

The noise causes prices to be wrong by α

Arbitrage opp of α

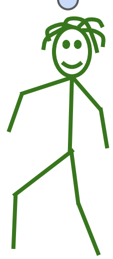

participant

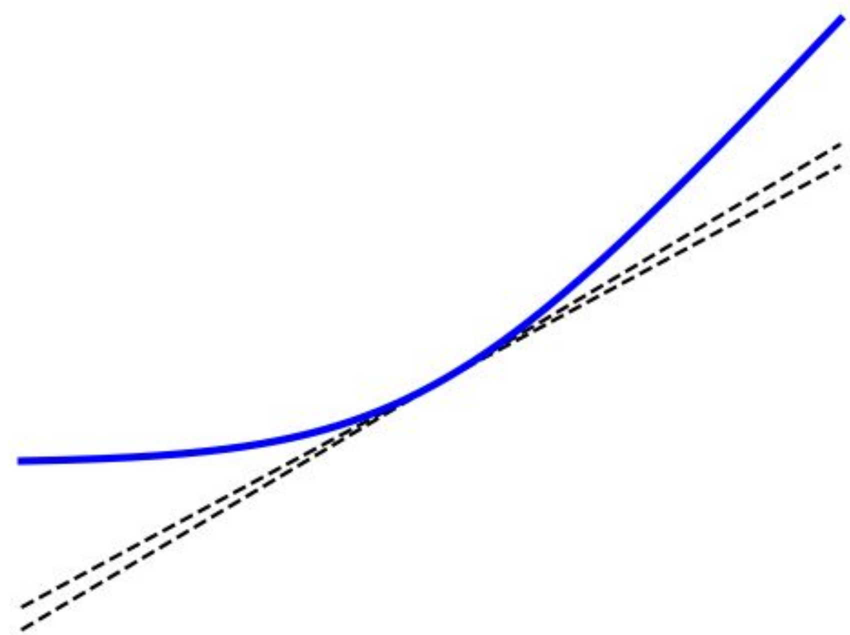

Initial approach: add a transaction fee

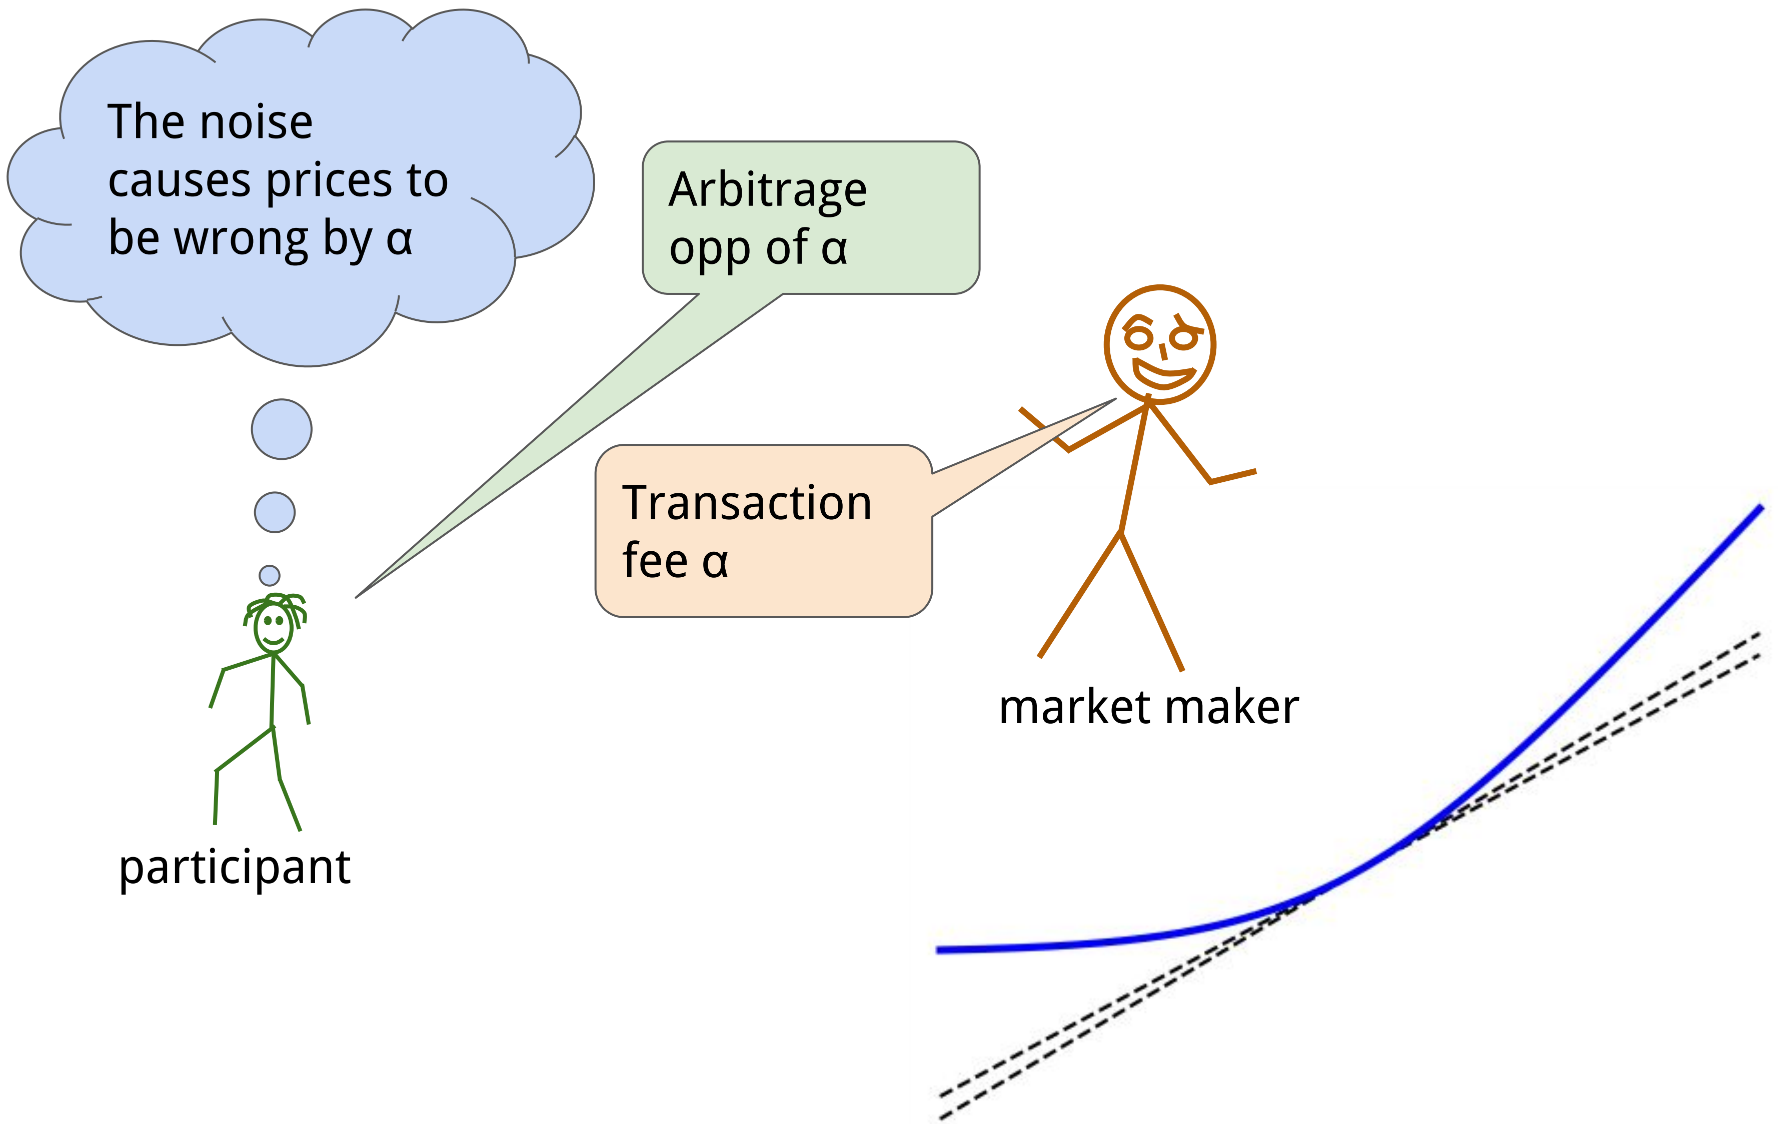

Transaction fee result (stepping stone)

Theorem

The same private market, but with transaction fee α, achieves:

- ε-differential privacy
- α-precision with high probability
- α-incentive to participate (prices are wrong by α ⇒ profit opportunity)
- worst-case loss $O(1/\lambda) = O(\log^2 T)$.

Transaction fee result (stepping stone)

Theorem

The same private market, but with transaction fee α, achieves:

- ε-differential privacy
- α-precision with high probability
- α-incentive to participate (prices are wrong by $\alpha \Rightarrow$ profit opportunity)
- worst-case loss $O(1/\lambda) = O(\log^2 T)$.

Proof idea:

$$\text{Loss} = \underbrace{(\text{Market maker loss})}_{O(1/\lambda)} + \underbrace{(\text{noise trader loss})}_{???} - \underbrace{(\text{transaction fees})}_{\alpha T}$$

Noise trader loss $\leq \alpha T$

Slightly intricate, depends on the details of the privacy scheme!

α is a convenient transaction fee that works, but not fundamental in the analysis.

Bounding noise trader loss

Each step, sell some number of previous bundle and buy a new bundle.

Bundle held for t steps ⇒ price changes at most λt ⇒ loss at most λt (size).

Sum over all bundles.

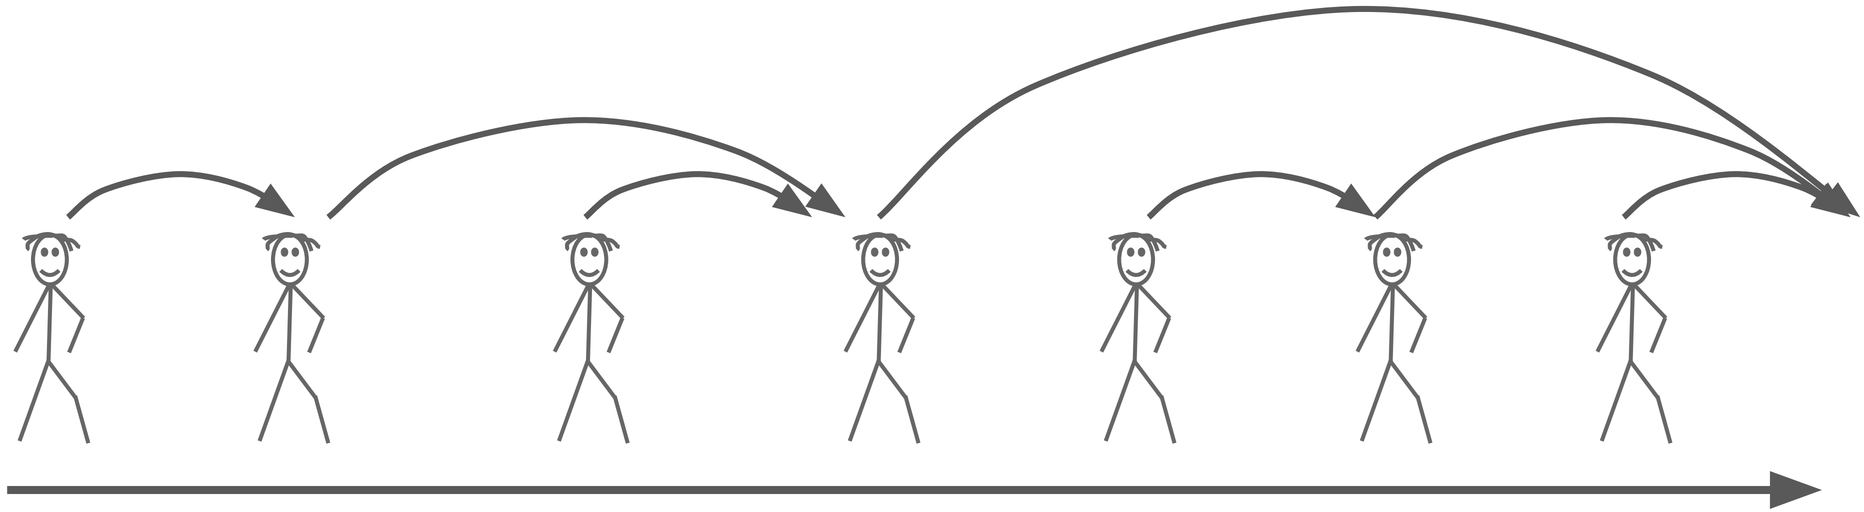

Wait a minute!

Let's try transaction fee 2α.

$$\text{Loss} = \underbrace{(\text{Market maker loss})}_{\log^2 T} + \underbrace{(\text{noise trader loss})}_{\alpha T} - \underbrace{(\text{transaction fees})}_{2\alpha T}$$

= Profit $\Omega(T)$!

Is this market guaranteed to make a profit??

Wait a minute!

Let's try transaction fee 2α.

$$\text{Loss} = \underbrace{(\text{Market maker loss})}_{\log^2 T} + \underbrace{(\text{noise trader loss})}_{\alpha T} - \underbrace{(\text{transaction fees})}_{2\alpha T}$$

= Profit $\Omega(T)$!

Is this market guaranteed to make a profit??

No … not if only $\log^2 T$ participants show up.

Wait a minute!

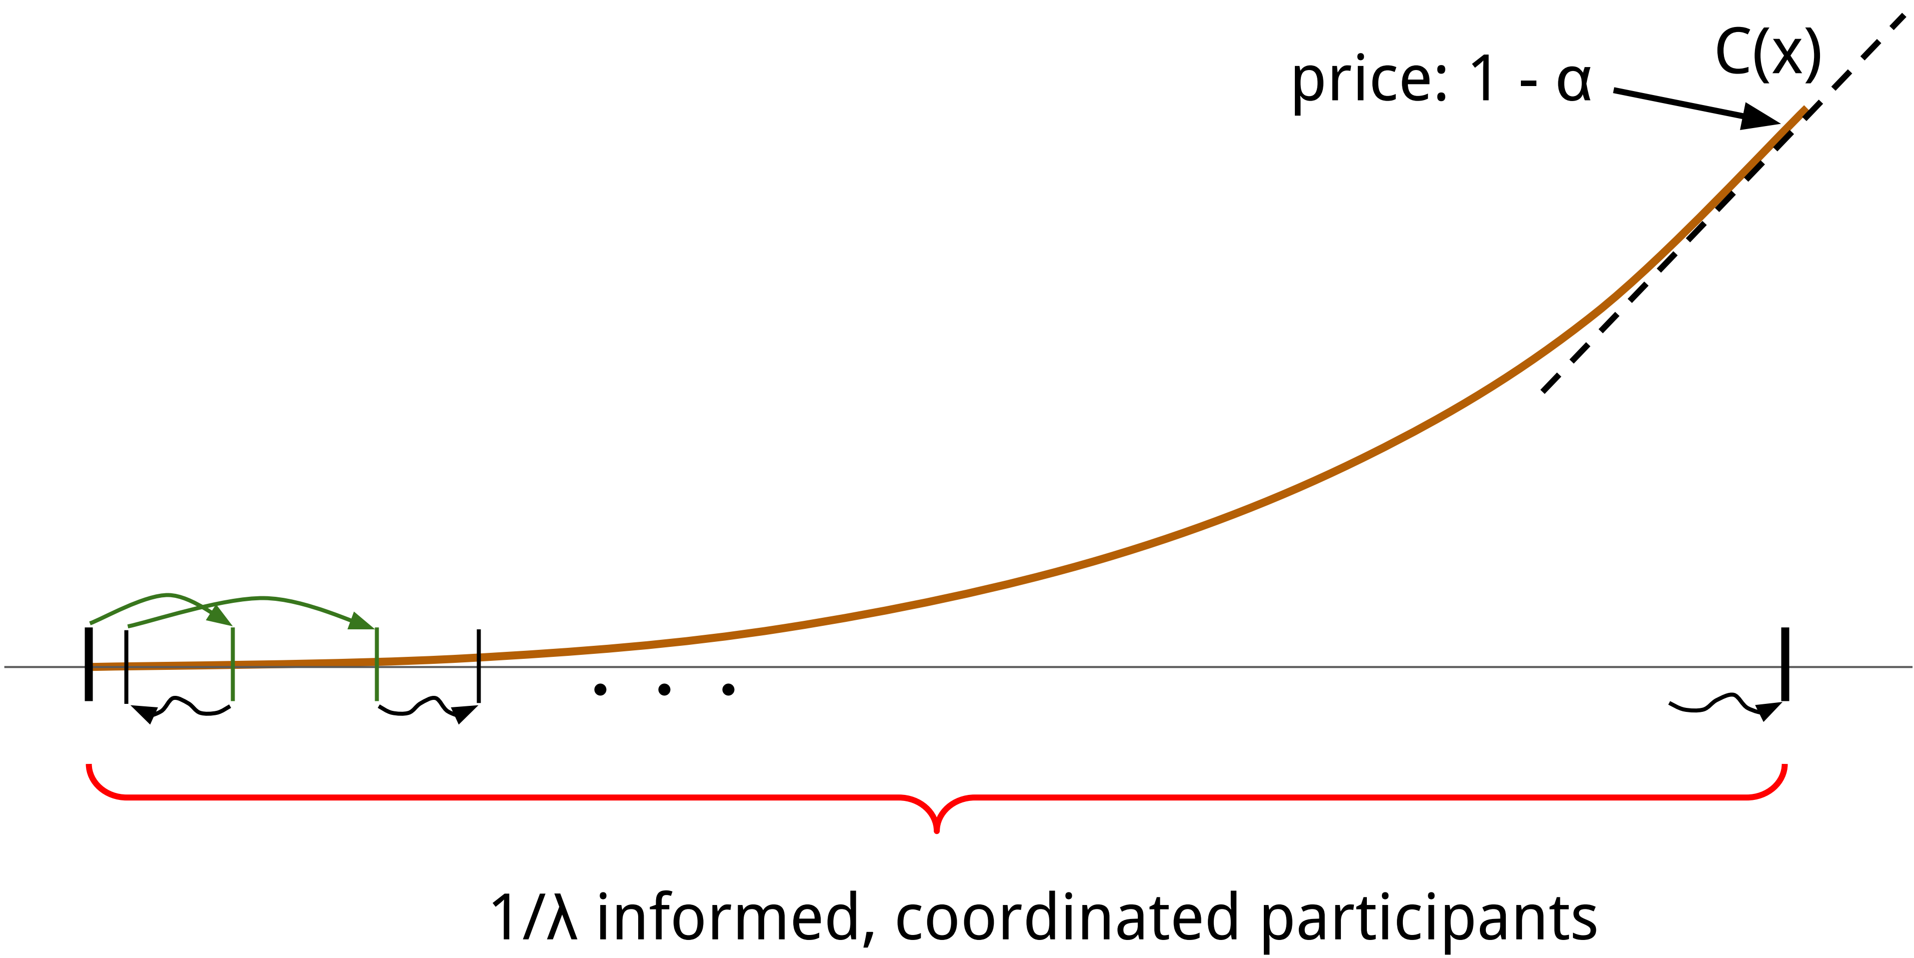

Wait a minute!

Let's try transaction fee 2α.

Loss = (Market maker loss) + (noise trader loss) - (transaction fees)

$\log^2 T \quad + \quad \alpha T \quad - \quad 2\alpha T$

= Profit Ω(T)!

Is this market guaranteed to make a profit??

No … not if only $\log^2 T$ participants show up.

So worst-cast loss is still $\log^2 T$.

Wait a minute!

Let's try transaction fee 2α.

Loss = (Market maker loss) + (noise trader loss) - (transaction fees)

$$\log^2 T \quad + \quad \alpha T \quad - \quad 2\alpha T$$

= Profit Ω(T)!

Is this market guaranteed to make a profit??

No … not if only $\log^2 T$ participants show up.

So worst-cast loss is still $\log^2 T$.

But if all T participants arrive … then yes!

Why?

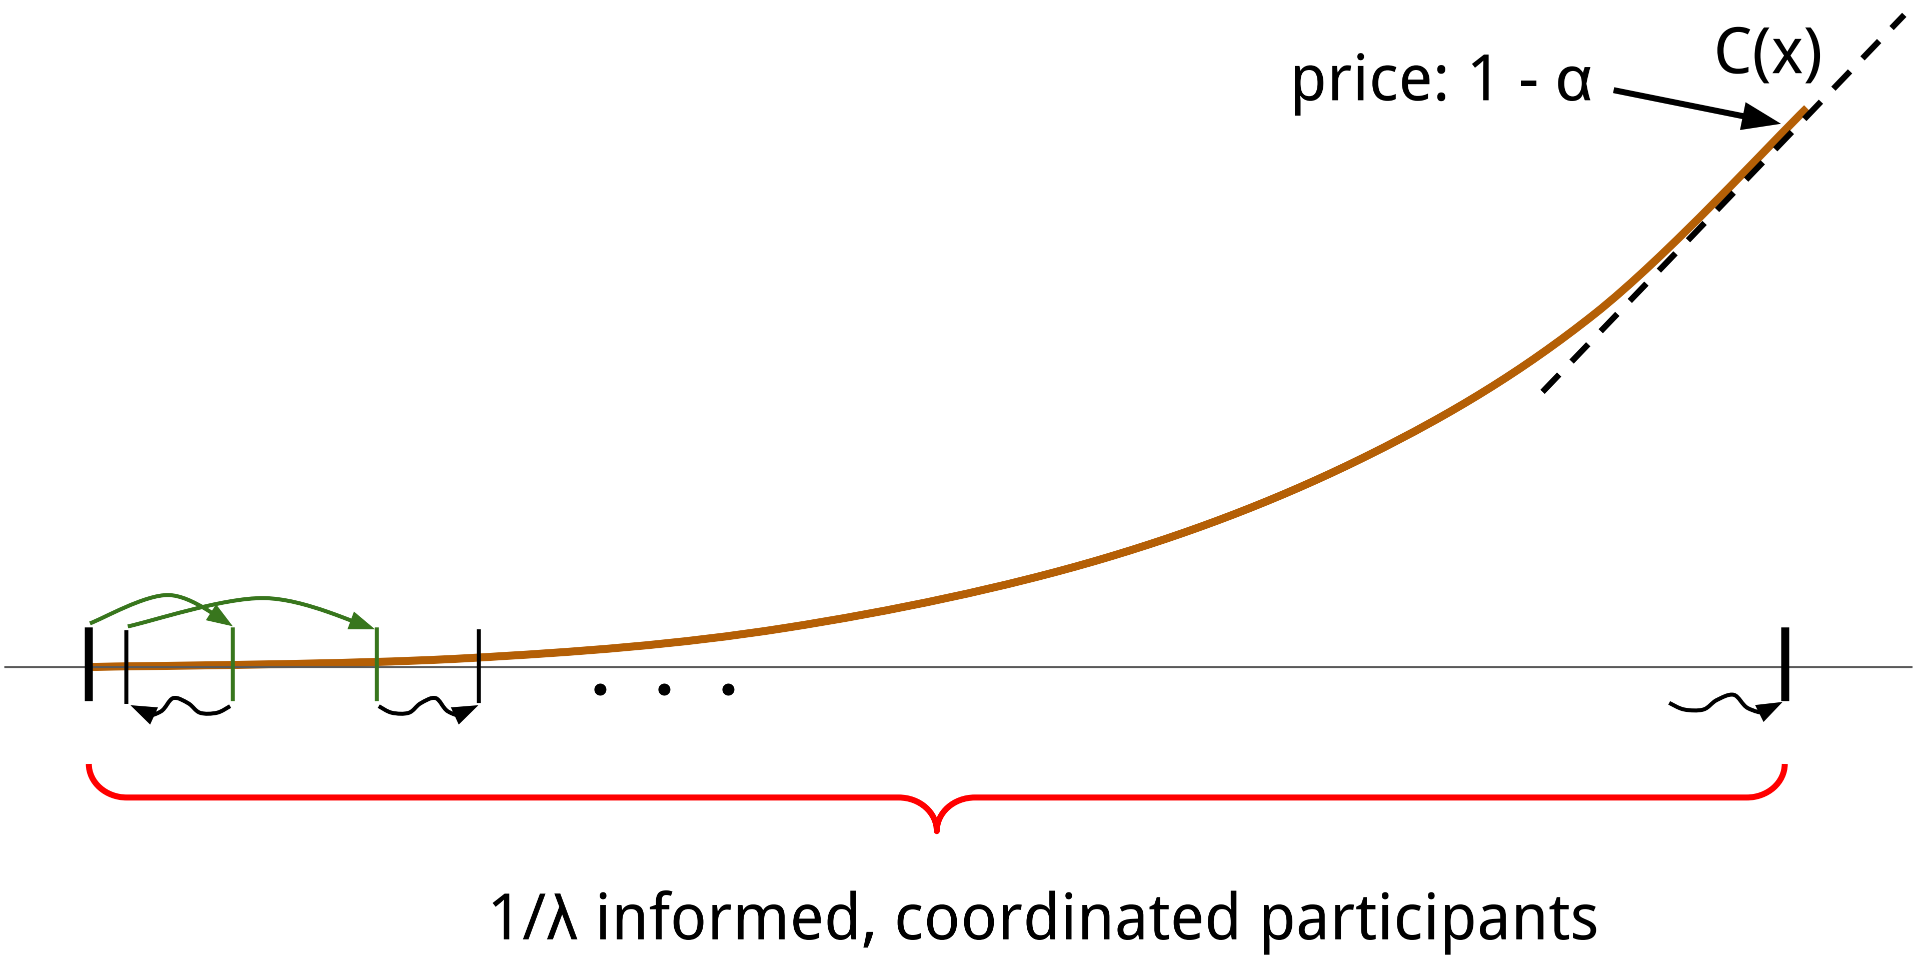

Why?

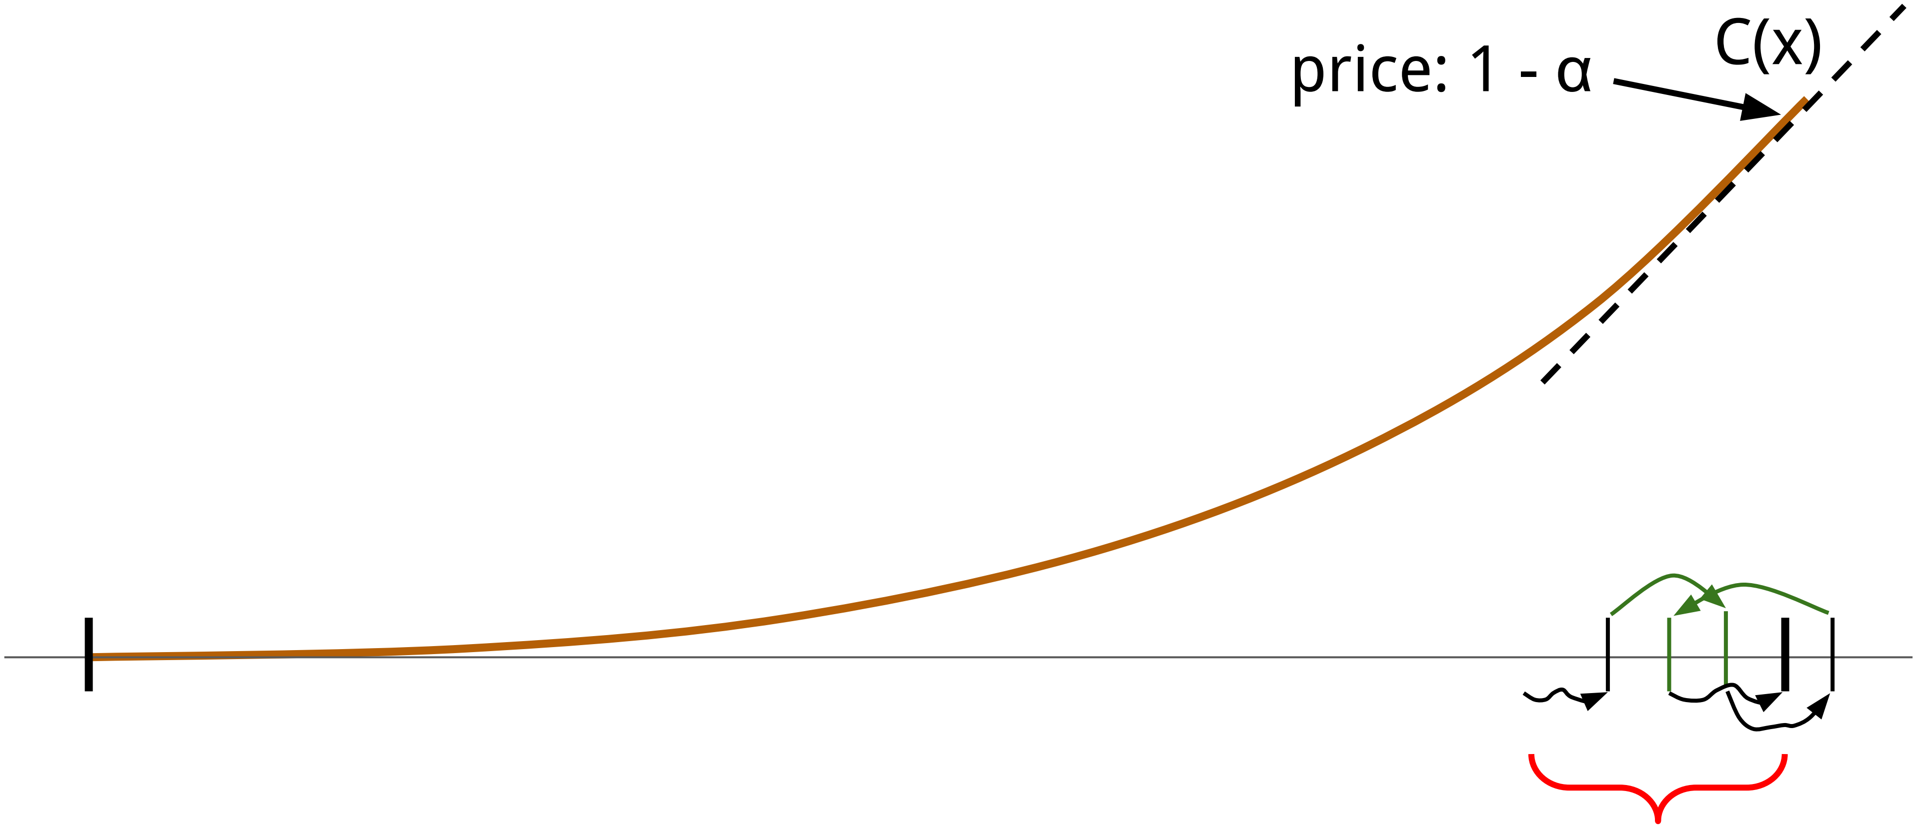

… mixed with lots of **disagreement!**

Why?

Disagreement is pure profit (transaction fees) for the market maker.

At most 1 / λ arrivals can agree!

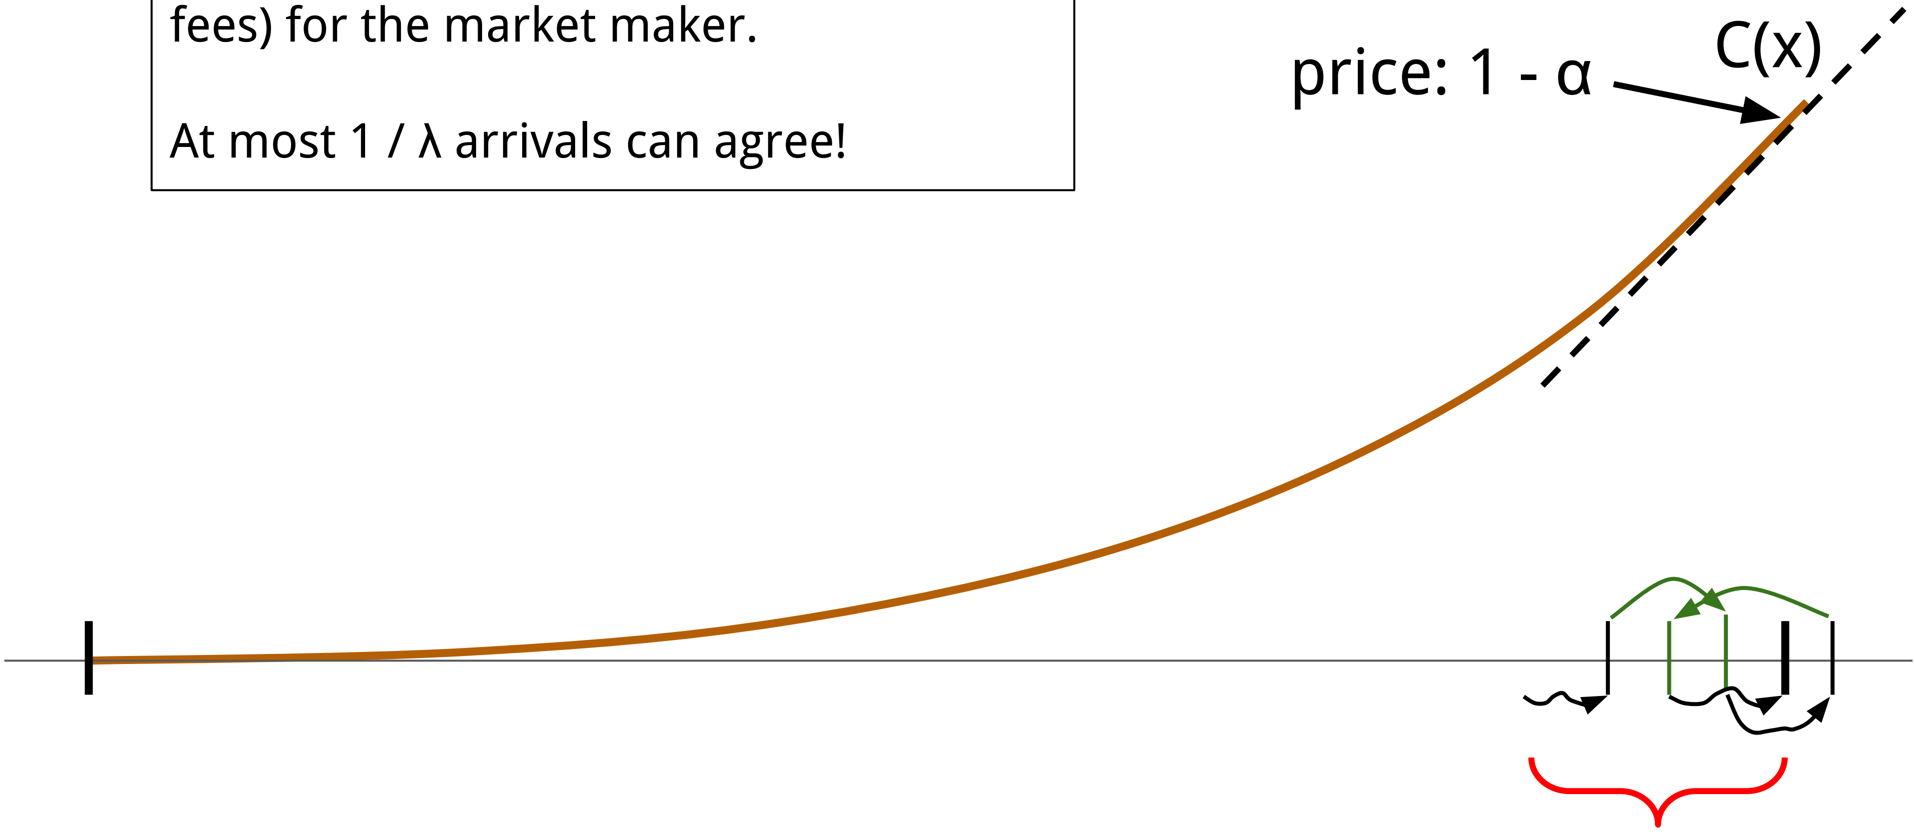

… mixed with lots of **disagreement!**

Iterative market construction

1. Set $T^1 = O(1)$ depending on privacy, accuracy parameters.
   Set $\lambda^1 = \Theta(1 / \log^2 T^1)$ and run this private market.

2. If not all participants arrive, done.

3. Set initial price = final price of above market.
   Se $T^2 = 4T^1$.
   Halve the accuracy parameters.
   Set $\lambda^2 = \Theta(1 / \log^2 T^2)$.
   Run this private market.

4. If not all participants arrive, done. Else, set $T^3 = 4T^2$ and continue….

Iterative market construction

Theorem

The iterative market satisfies all the above privacy, precision, incentive constraints as well as **worst case loss bounded by O(1)** regardless of number of arrivals.

Iterative market construction

Theorem
The iterative market satisfies all the above privacy, precision, incentive constraints as well as **worst case loss bounded by O(1)** regardless of number of arrivals.

Proof idea.
Each market either completes, or stops early.

Each market that completes makes enough profit to subsidize the $O(1/\lambda)$ loss of the next market.

Only the last market stops early; it is either already subsidized (net profit), or the first market (constant-size loss).

Future directions

- Other (more elegant) constructions?
- Any helpful light shed on adaptive-volume (liquidity) markets?
- Interactions between privacy and information aggregation
  seem to be opposed…
- More broadly: **value of information**, purchasing information

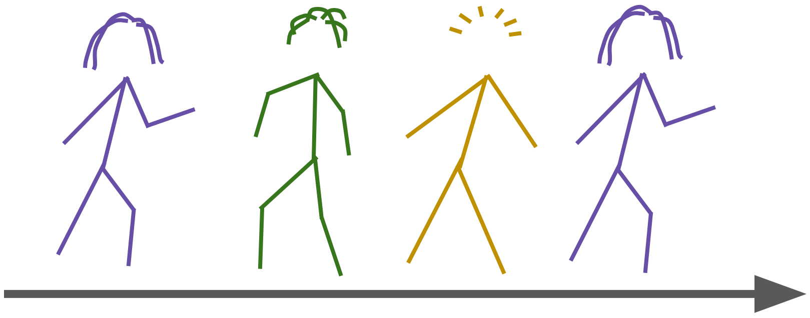

Thanks!