Tullock Contests (Contests with Proportional Allocation)

Learning Dynamics …

Abheek Ghosh

Joint work with Edith Elkind and Paul Goldberg

UNIVERSITY OF OXFORD

Contests

games where

- a set of agents compete
- by putting costly and irreversible effort
- to win valuable prizes

e.g., sports (more later)

Tullock contest

- n agents
- prize $= 1$ (normalized)
- effort of agent i: $x_i \geq 0$
- effort profile $\boldsymbol{x} = (x_1, x_2, \ldots, x_n)$
- proportional allocation
- non-negative, continuous, increasing, (weakly) convex cost
- expected utility

$$u_i(\boldsymbol{x}) = \frac{x_i}{\sum_j x_j} - c_i(x_i)$$

Example

Agents	Effort x_i	Reward (proportional)	Cost Function $c_i(x_i)$	Cost	Utility
	0.2	$\frac{2}{6} = 0.33$	$c_1(x_1) = \frac{x_1}{2}$	0.1	$0.33 - 0.1 = 0.23$
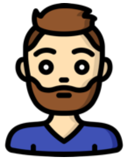	0.1	$\frac{1}{6} = 0.17$	$c_2(x_2) = x_2$	0.1	$0.17 - 0.1 = 0.07$
	0.3	$\frac{3}{6} = 0.5$	$c_3(x_3) = x_3^2$	0.09	$0.5 - 0.09 = 0.41$

Some applications

- proof-of-work (stake) cryptocurrencies like Bitcoin (Etherium)
  - effort (stake): x_i
  - probability of creating the block: $\frac{x_i}{\sum_j x_j}$
  - computational (opportunity) cost: $c_i(x_i)$
- rent-seeking (work by Tullock)
- political lobbying and donation
- research & development races
- extensions (discussed later)
  - parallel contests: crowdsourcing (including in blockchains)
  - group contests

Properties

strictly concave utility function $u_i(\boldsymbol{x}) = \frac{x_i}{\sum_j x_j} - c_i(x_i)$

$\Longrightarrow$ unique best response (BR) $BR_i(\boldsymbol{x}_{-i}) = \underset{z \geq 0}{\operatorname{argmax}}\ u_i(z, \boldsymbol{x}_{-i})$

unique pure-strategy Nash equilibrium (non-trivial)

Talk overview

- learning dynamics
  - best-response dynamics
    - linear cost functions a.k.a. lottery contest
    - convex cost functions
  - continuous best-response dynamics, fictitious play, …
- extensions/variations of Tullock contests
  - parallel contests
  - group contests
  - discrete action spaces

Learning dynamics

Why study learning dynamics?

- predict agents' behavior
- equilibrium analysis assumes all agents
  - know the rules of the game
  - know everyone's utility function
  - are fully rational

  conceptually and empirically these assumptions may not hold
- learning dynamics
  - agents respond to the incentives provided by their environment
  - … in a decentralized manner
  - e.g., best-response, fictitious play, no-regret dynamics

Best-response (BR) dynamics

- BR dynamics
  - initial state:
$$\boldsymbol{x}(0) = (x_1(0), x_2(0), \dots, x_n(0))$$
  - update:
$$x_i(t+1) = BR_i\big(\boldsymbol{x}_{-i}(t)\big)$$

- a random agent moves at each time step (say picked uniformly for this talk)

Linear costs

$$u_i(\boldsymbol{x}) = \frac{x_i}{\sum_j x_j} - c_i(x_i)$$

linear

Results – homogeneous agents

homogeneous: same cost function

convergence to an ϵ-approx equilibrium in

- 2 agents

$$\log\log\left(\frac{1}{\epsilon}\right) + \log\log\left(\frac{1}{\gamma}\right) + \Theta(1)$$

- $n \geq 3$ agents, with high probability

$$O\left(n\log\left(\frac{n}{\epsilon}\right) + \log\log\left(\frac{1}{\gamma}\right)\right), \qquad \Omega\left(n\log(n) + \log\left(\frac{1}{\epsilon}\right) + \log\log\left(\frac{1}{\gamma}\right)\right)$$

γ is function of initial state (in most cases, the smallest positive effort)

Results – non-homogeneous

non-convergence of BR dynamics

- instances that lead to a cycle
  - *generic*: set of instances (and starting points) that lead to cycle have positive measure

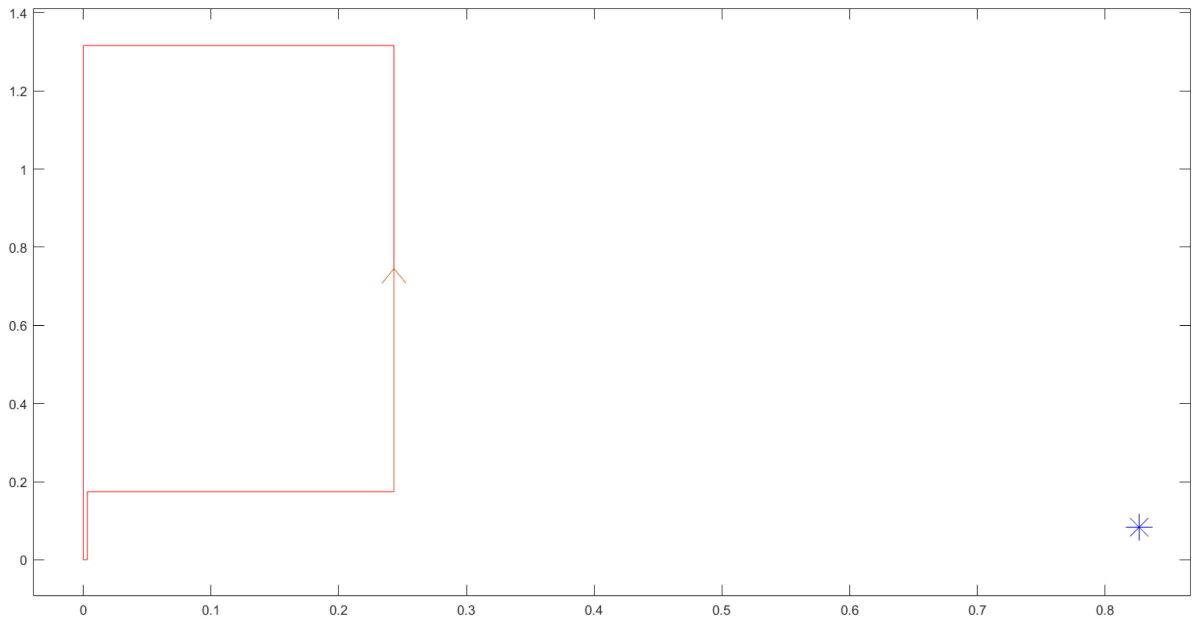

$$c_1 x_1 = x_1$$
$$c_2 x_2 = \frac{x_2}{10}$$

cycle of 6 steps

Proof idea for $\mathbf{n} \geq 3$ homogeneous agents

combine:

- coordinate descent
- smooth and strongly convex potential function (close to equilibrium)
- Markov chains (away from equilibrium)

Coordinate descent

an arbitrary convex function $g(\boldsymbol{x})$

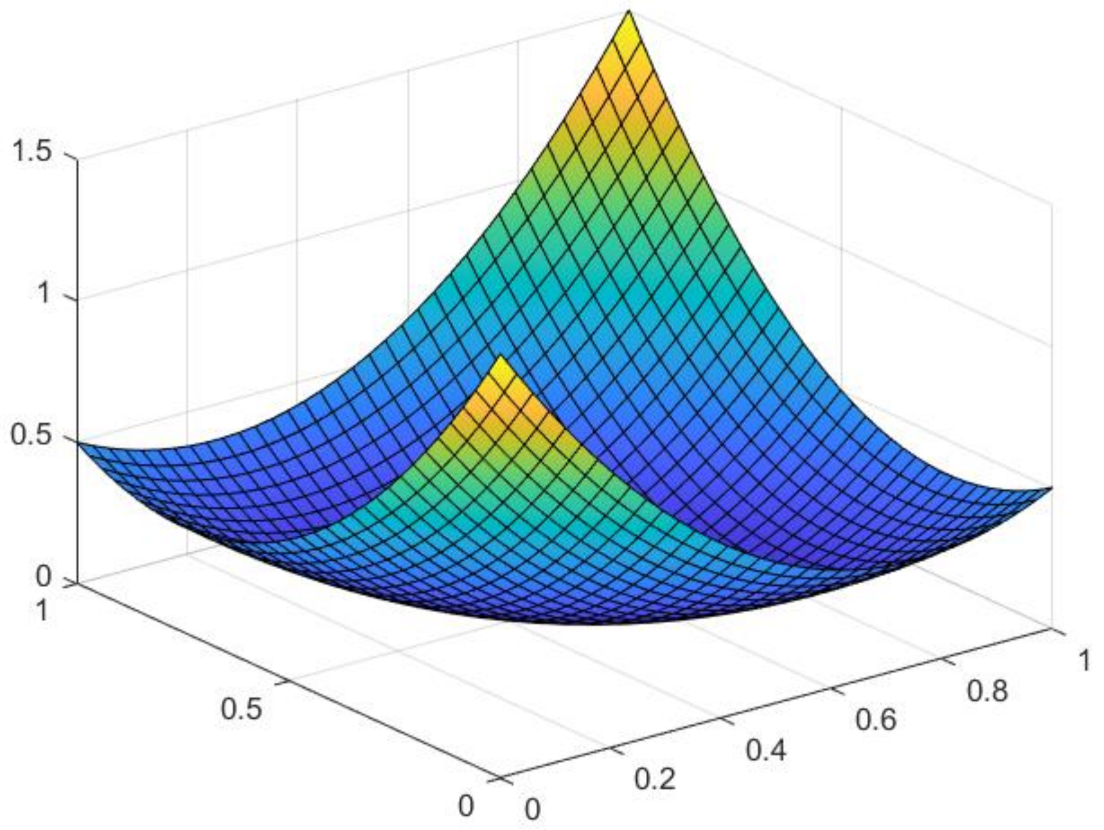

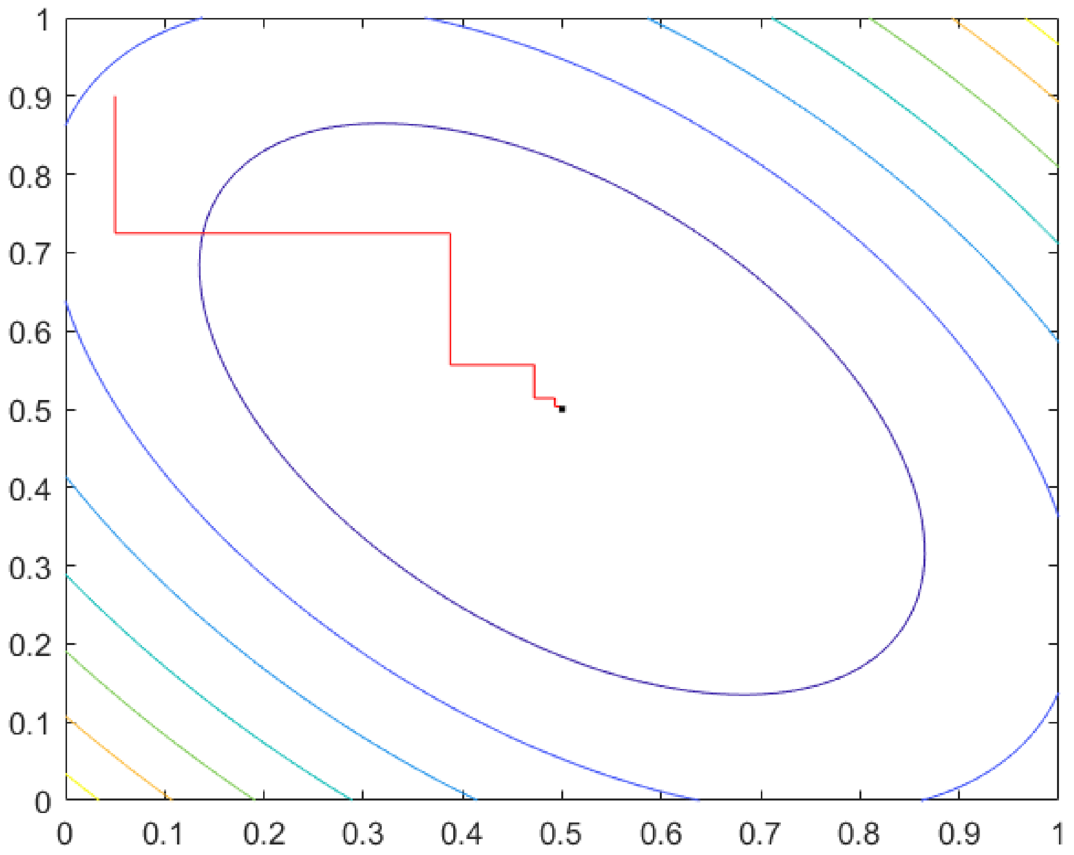

Smoothness and strong-convexity

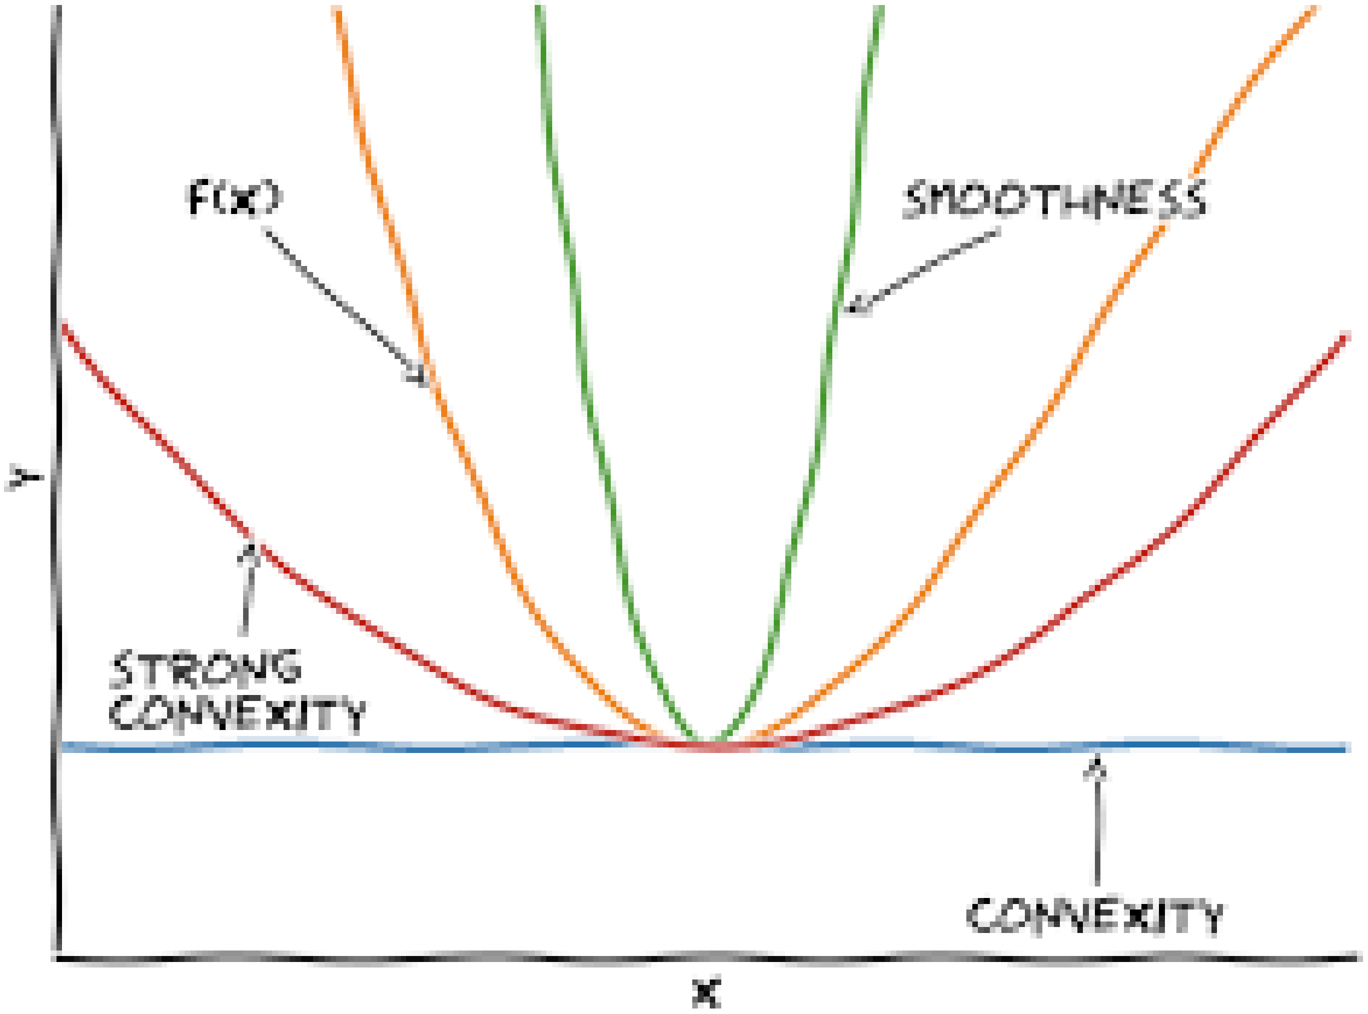

Potential function

- smooth and strongly convex function + coordinate descent
  - $O\left(\log\left(\frac{1}{\epsilon}\right)\right)$ convergence to ϵ-approx. minima

- construct such a potential near the equilibrium
- away from the equilibrium: alternative techniques

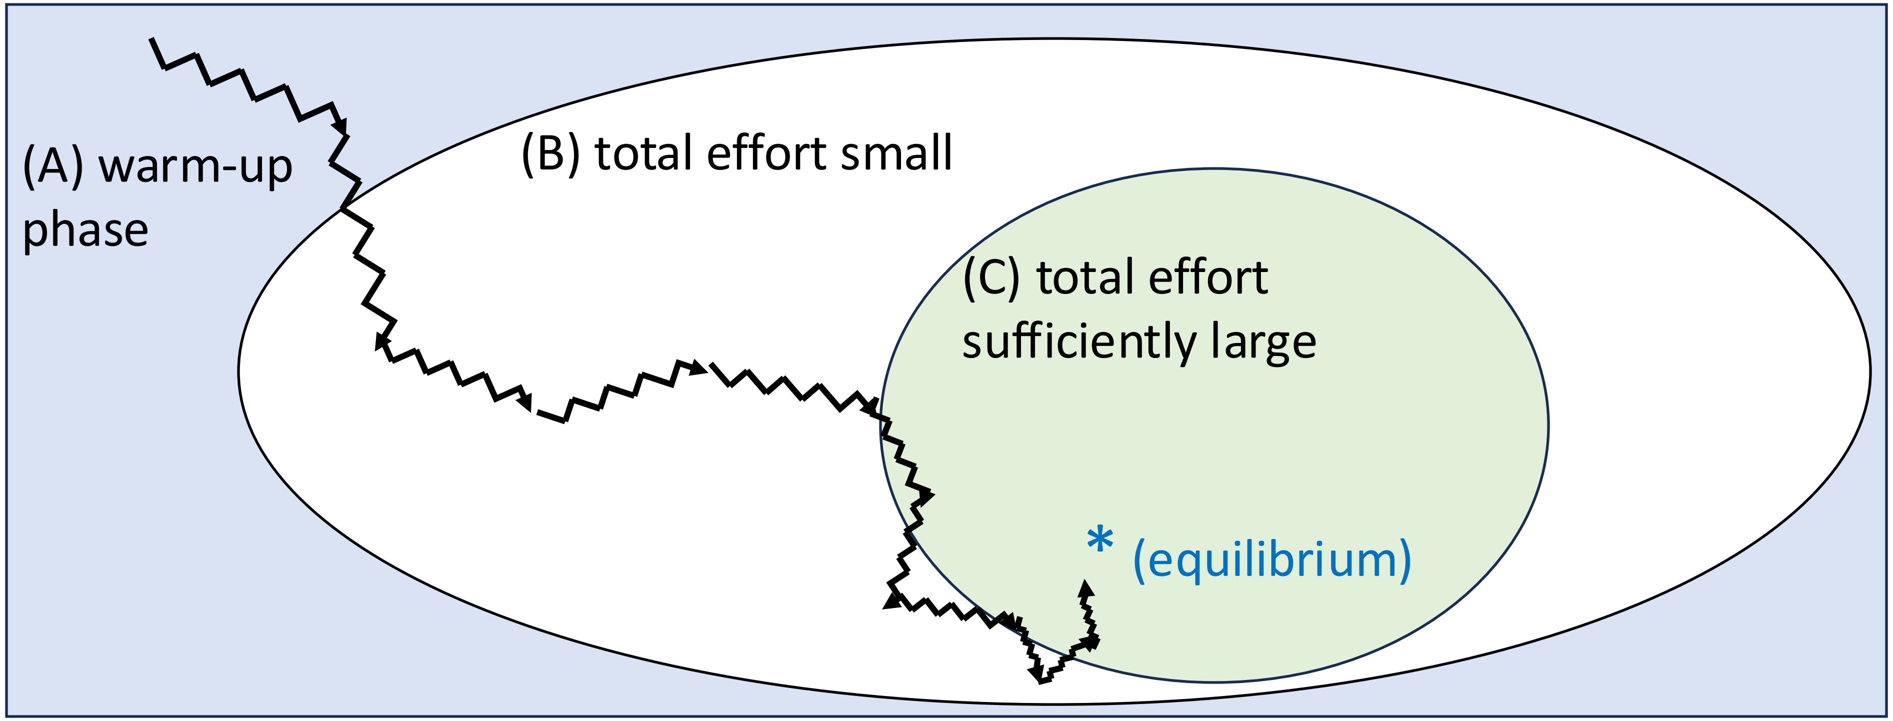
(A) warm-up phase
(B) total effort small
(C) total effort sufficiently large
* (equilibrium)

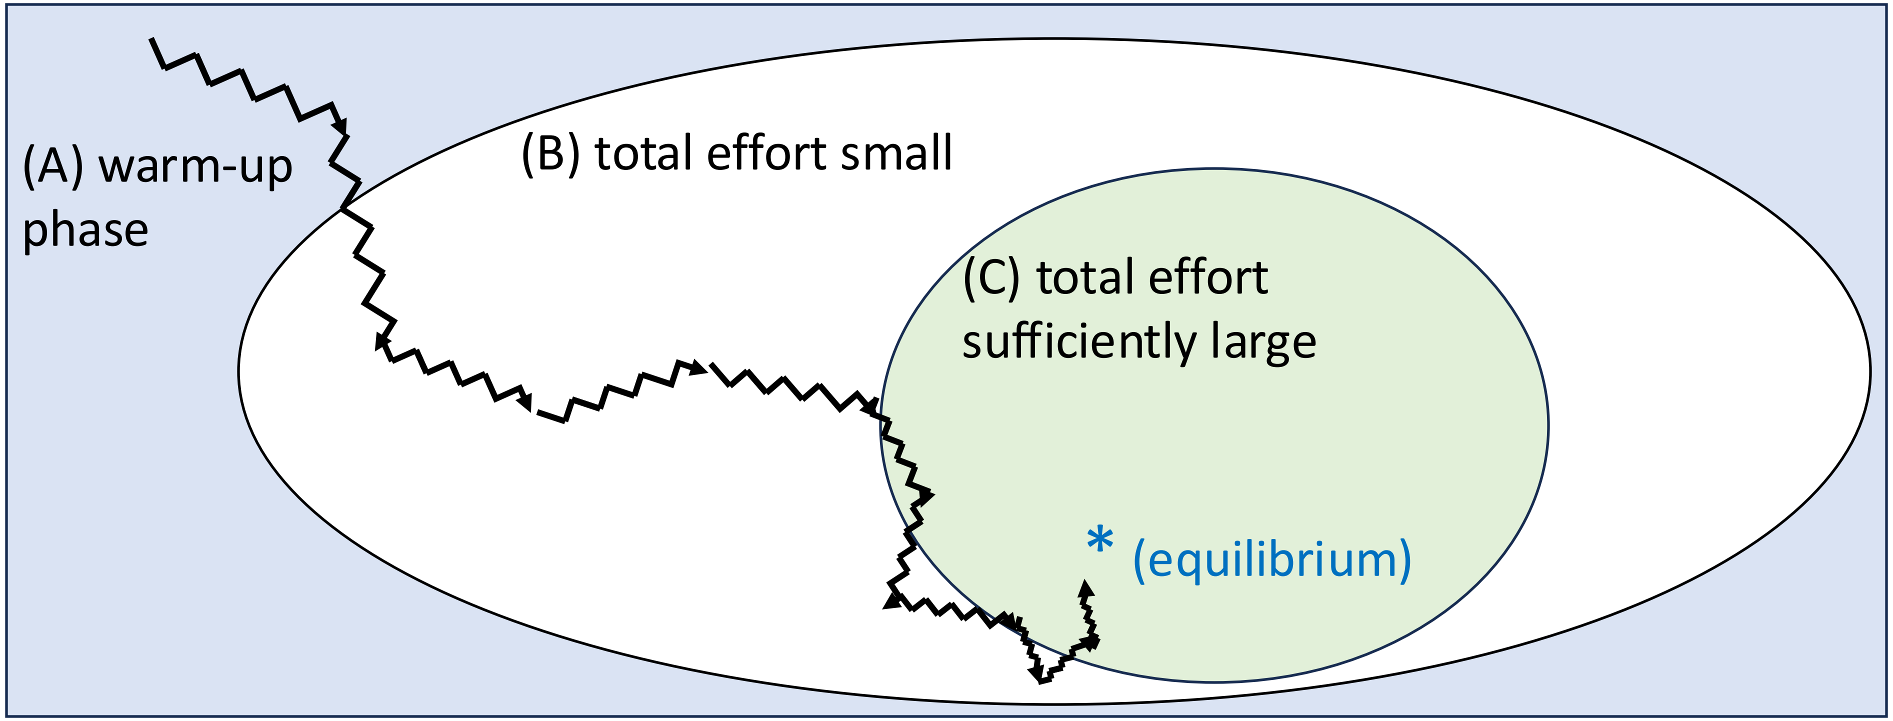

Region (C): high total effort

- smooth and strongly convex potential
- potential decreases rapidly (coordinate descent ≡ BR dynamics)

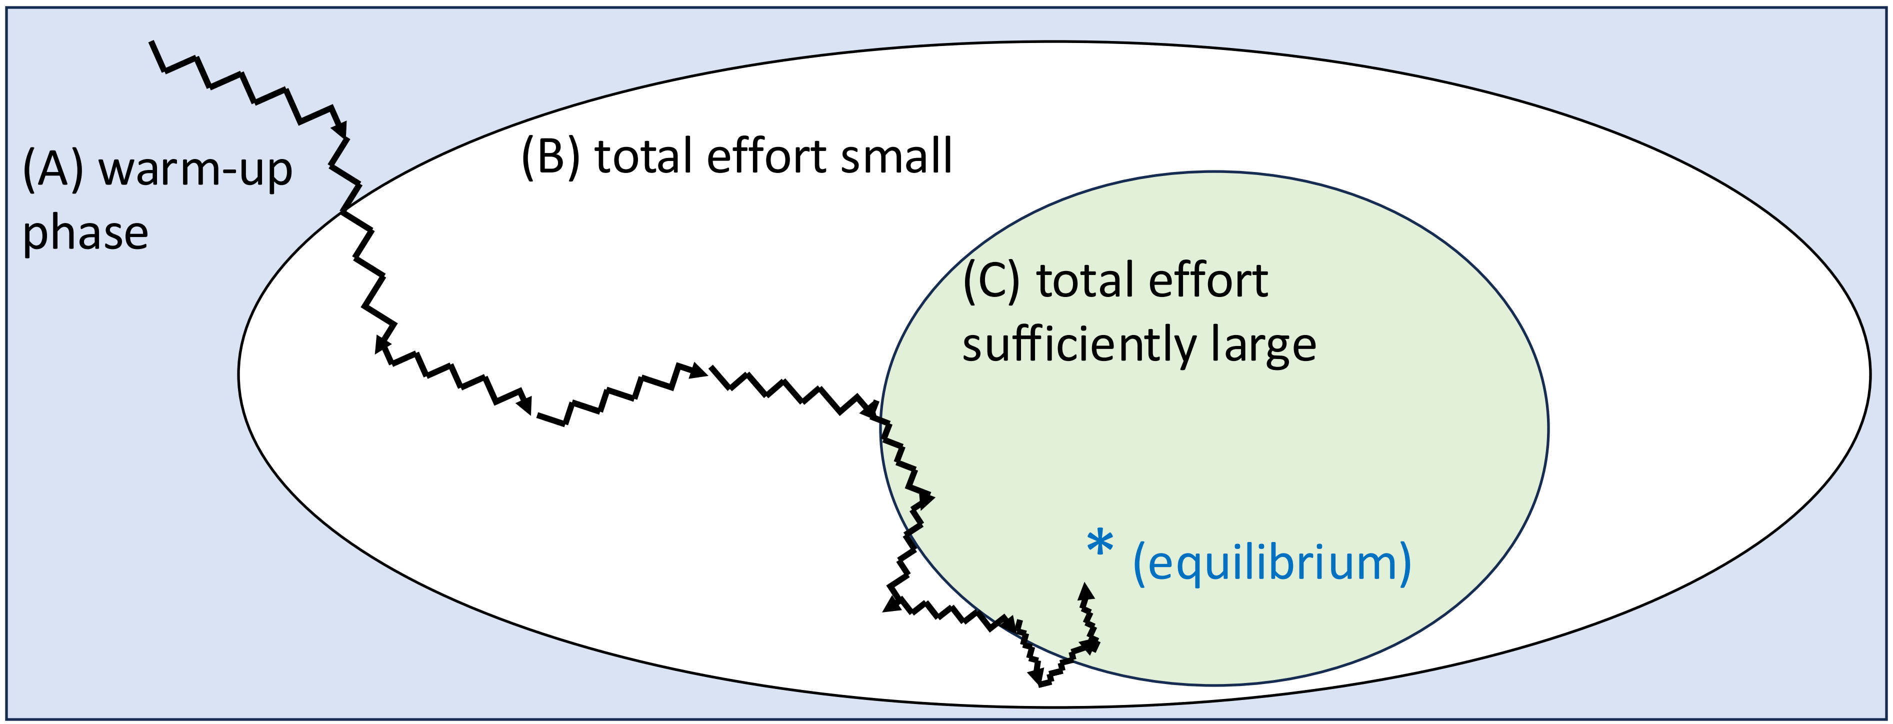

Region (B): double-exponentially decreasing Markov chain on total effort

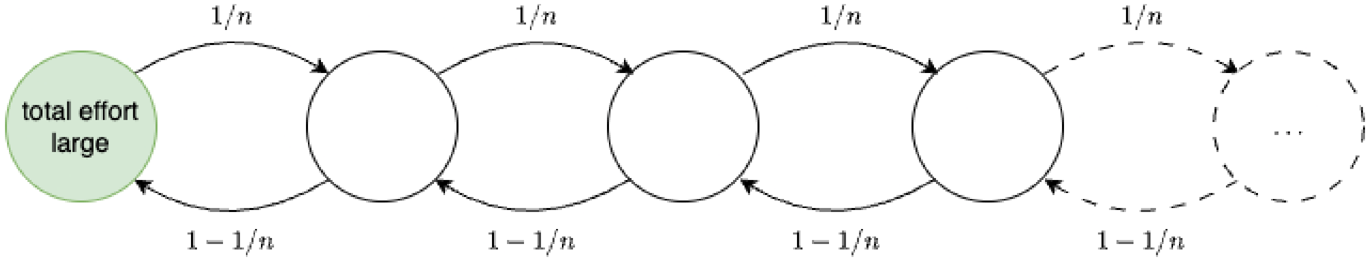

Convex costs

$$u_i(\boldsymbol{x}) = \frac{x_i}{\sum_j x_j} - c_i(x_i)$$

weakly convex

Convex costs

(homogeneous agents) convergence to an ϵ-appx equilibrium

- 2 agents (same bound as linear costs)

$$\log\log\left(\frac{1}{\epsilon}\right) + \log\log\left(\frac{1}{\gamma}\right) + \Theta(1)$$

- $n \geq 3$ agents (weaker bound in n than linear costs)

$$O\left(n^2\log(n)\log\left(\frac{n}{\epsilon}\right) + \log\log\left(\frac{1}{\gamma}\right)\right)$$

Analysis

- analysis for the linear case doesn’t extend. Some challenges:
  - no closed-form formula for BR
  - no potential function

- in the analysis
  - an adversarial/approximate linearization of the BR dynamics
  - discounted-sum dynamics

Convex cost

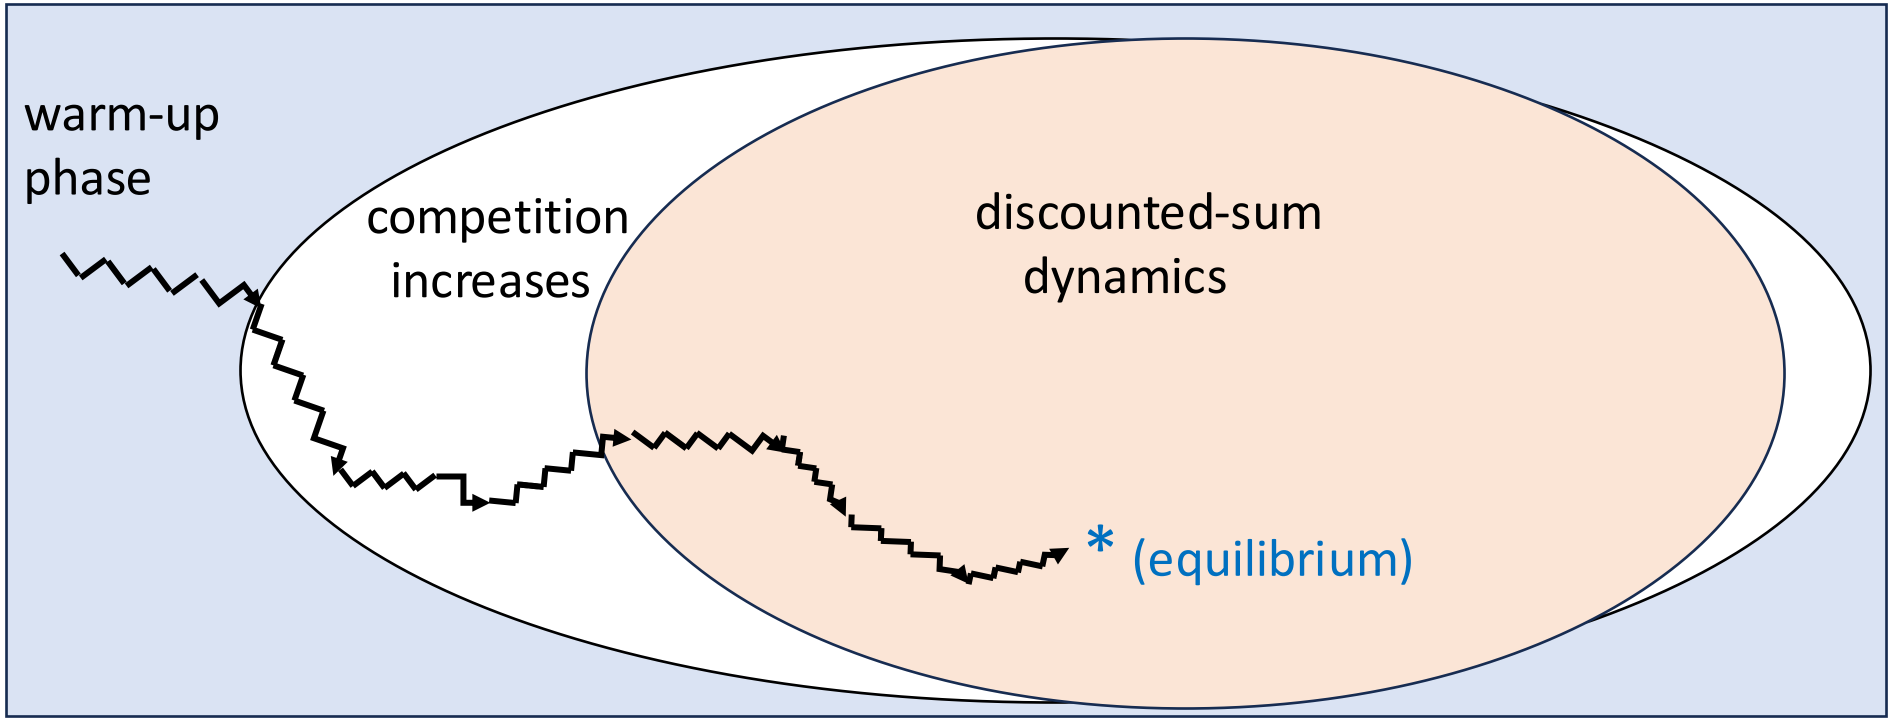
warm-up
phase
competition
increases
discounted-sum
dynamics
* (equilibrium)

Discounted-sum dynamics

- initial state:

$$\boldsymbol{z}(0) = \left(z_1(0), z_2(0), \ldots, z_n(0)\right) \in \mathbf{R}^{\mathrm{n}}$$

- update (for random i):

$$z_i(t+1) = -\beta_t \cdot \sum_{j \neq i} z_j(t)$$

β_t picked adversarially in $[0, B]$ where $0 \leq B < 1$.

- we show that this dynamics converges to 0 rapidly

Convergence

- potential function $\mathbf{1}(z_j < 0)$ is the indicator function

$$\phi(\boldsymbol{z}) = \max\left(\sum_j z_j \cdot \mathbf{1}(z_j > 0), \sum_j -z_j \cdot \mathbf{1}(z_j < 0)\right)$$

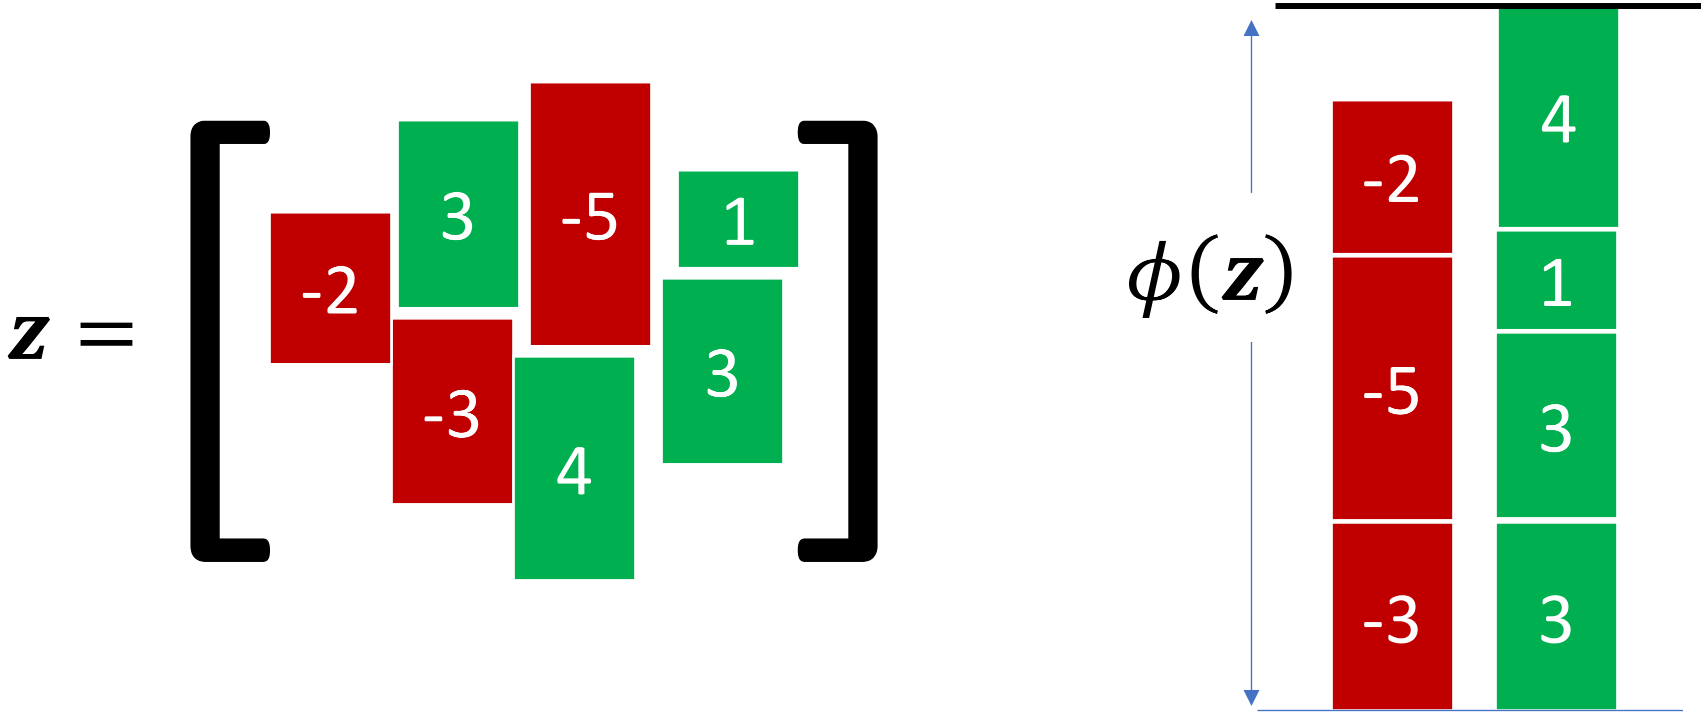

Case 1: smaller side

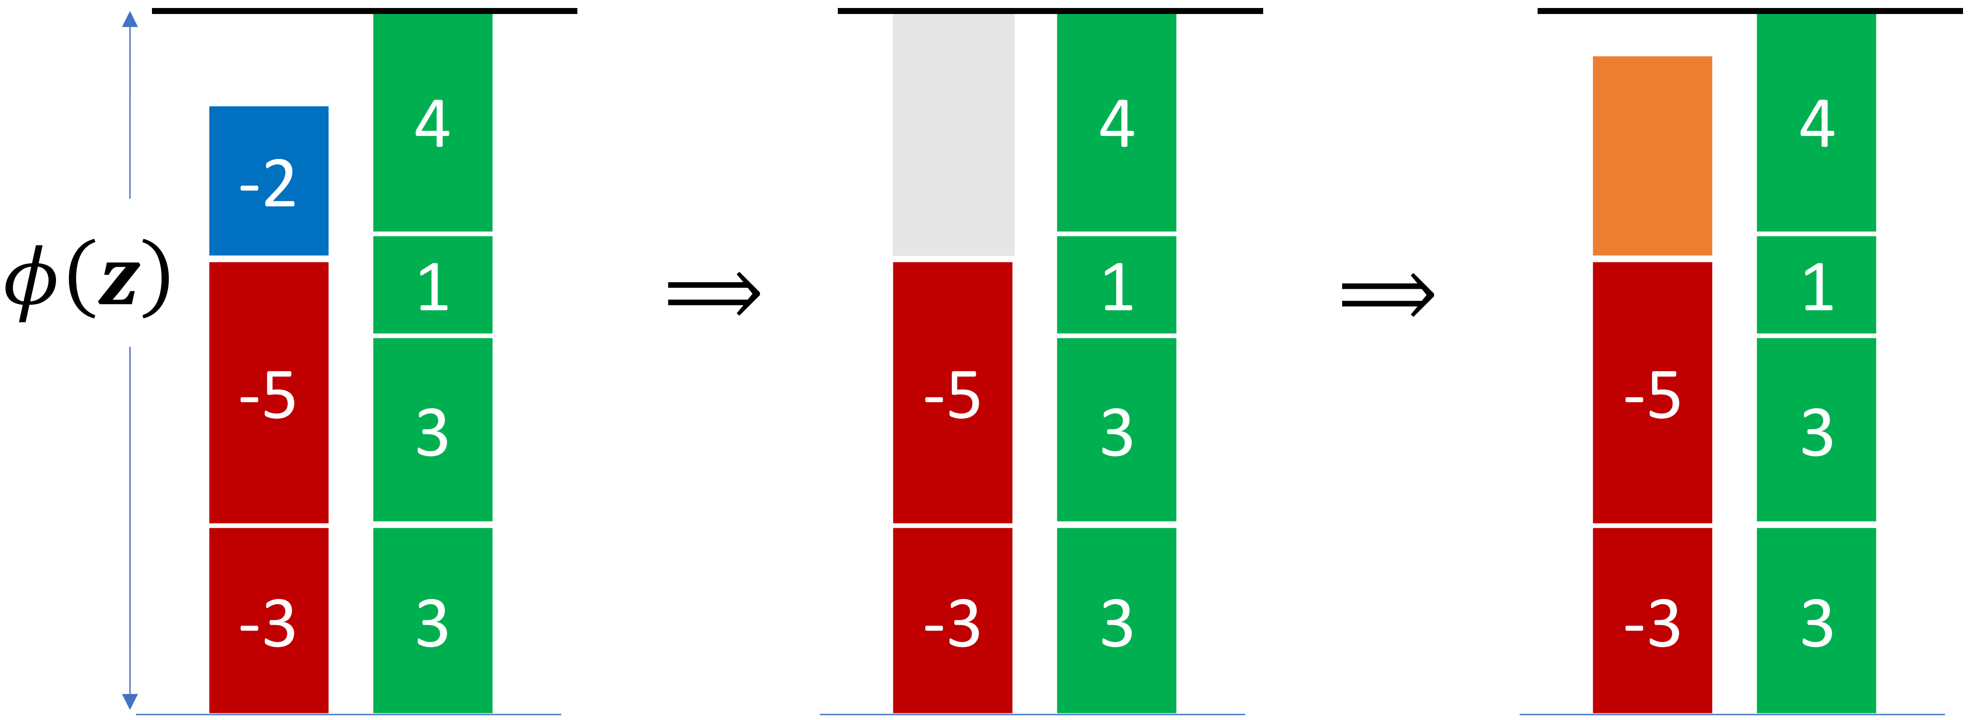

Case 2: larger side, large element

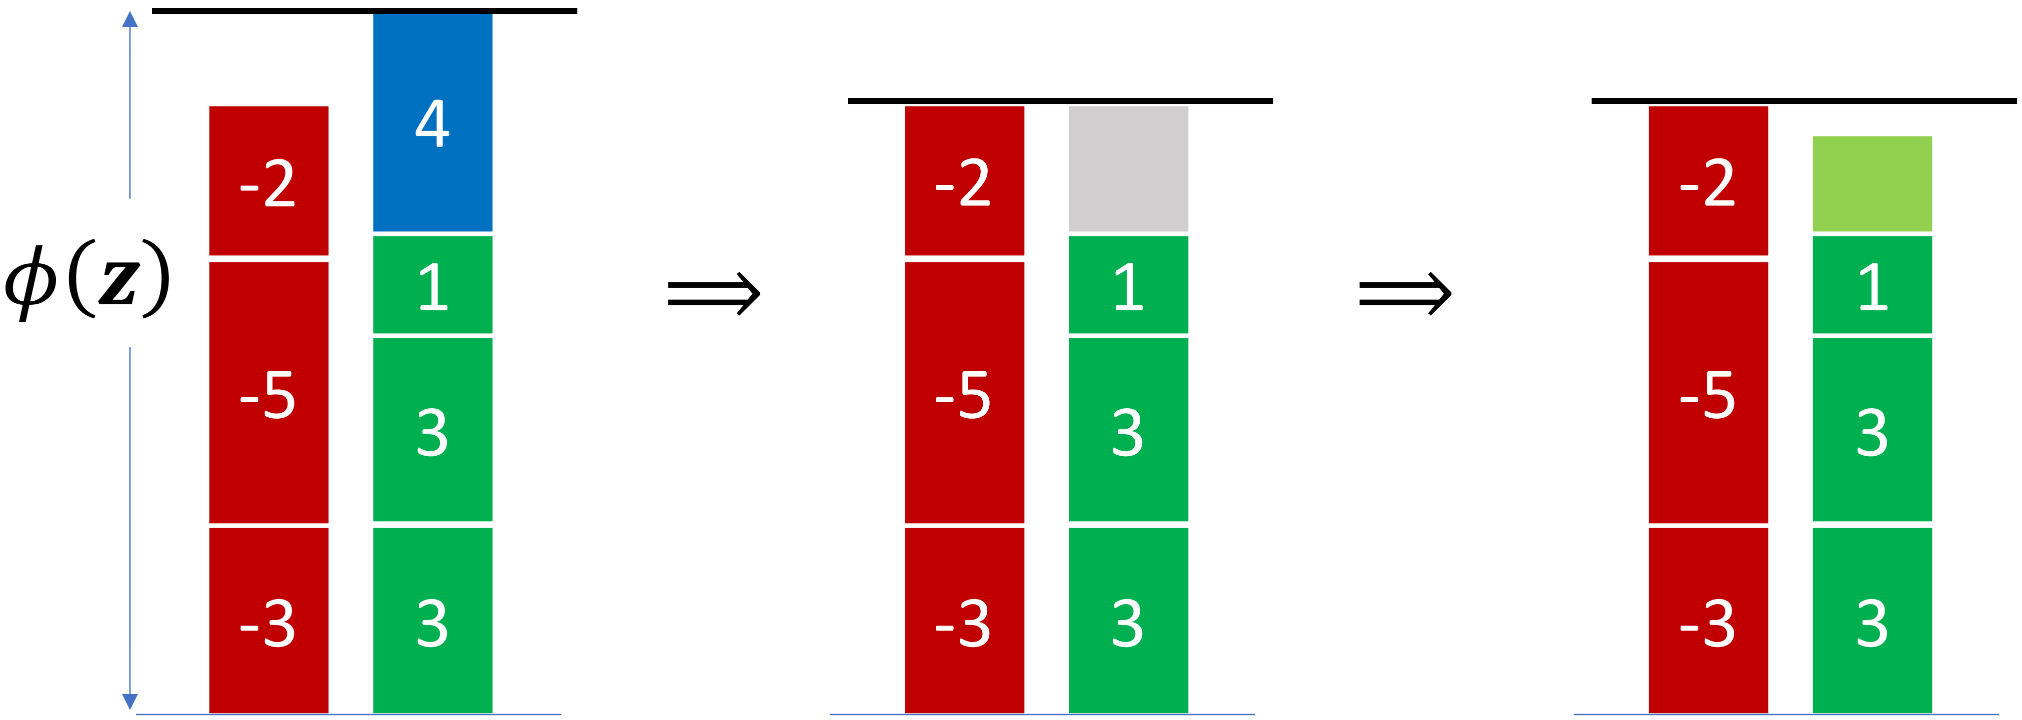

Case 3: larger side, small element

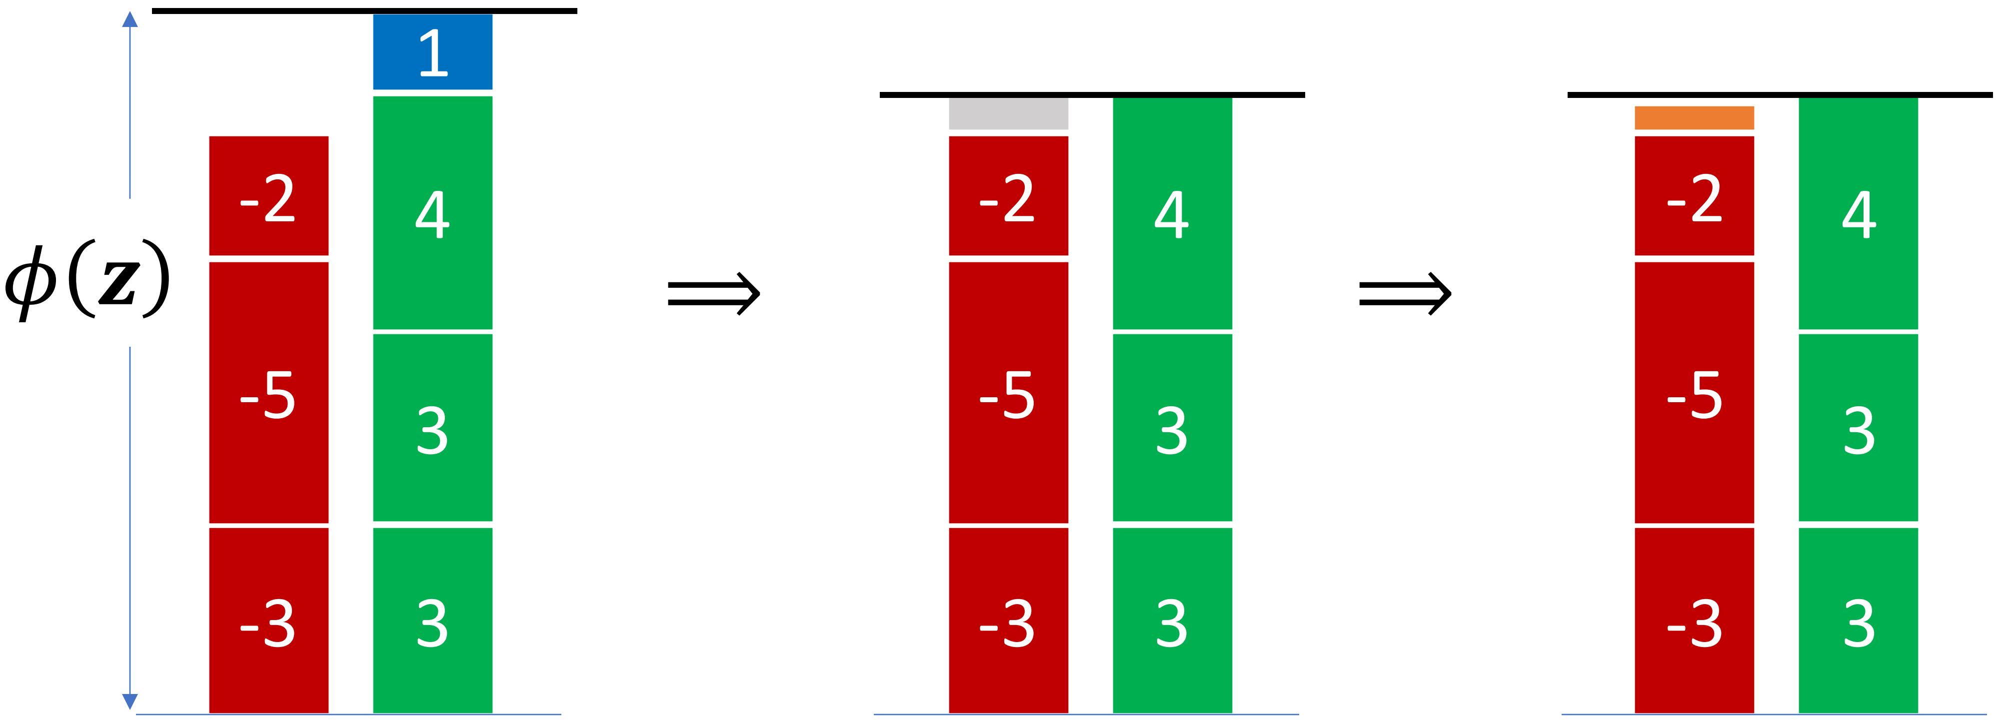

Continuous BR dynamics

- (usual) BR dynamics:

$$x_i(t+1) = BR_i\big(\boldsymbol{x}_{-i}(t)\big)$$
$$x_i(t+1) = x_i(t) + \Big(BR_i\big(\boldsymbol{x}_{-i}(t)\big) - x_i(t)\Big)$$

- small Δt steps

$$x_i(t+\Delta t) = x_i(t) + \Delta t \cdot \Big(BR_i\big(\boldsymbol{x}_{-i}(t)\big) - x_i(t)\Big)$$

- continuous BR dynamics

$$\frac{dx_i(t)}{dt} = \lim_{\Delta t \to 0} \frac{x_i(t+\Delta t) - x_i(t)}{\Delta t} = BR_i\big(\boldsymbol{x}_{-i}(t)\big) - x_i(t)$$

Results

- continuous BR dynamics

$$\frac{dx_i(t)}{dt} = BR_i\big(\boldsymbol{x}_{-i}(t)\big) - x_i(t)$$

Result: converges to an ϵ-approximate equilibrium
for non-homogeneous agents in $O\left(\log\left(\frac{1}{\epsilon}\right)\right)$ time

Idea: Lyapunov function argument

- small Δt steps

$$x_i(t+\Delta t) = x_i(t) + \Delta t \cdot \Big(BR_i\big(\boldsymbol{x}_{-i}(t)\big) - x_i(t)\Big)$$

Result: always converges if $\Delta t \leq \frac{\kappa_1}{n}$
instances that cycle if $\Delta t \geq \frac{\kappa_2}{n}$ for constants κ_1, κ_2

Idea: Taylor expansion of the Lyapunov function

Fictitious play and generalizations

- $x_i(t+1) = BR_i\left(\frac{1}{t}\sum_{\tau=1}^{t} \boldsymbol{x}_{-i}(\tau)\right)$
- generalization – weighted average

$$x_i(t+1) = BR_i\left(\frac{\sum_{\tau=1}^{t} w_\tau \, \boldsymbol{x}_{-i}(\tau)}{\sum_{\tau=1}^{t} w_\tau}\right)$$

- converges if $\frac{w_t}{\sum_{\tau=1}^{t} w_\tau} \to 0$ and $\sum_t w_t \to \infty$ as $t \to \infty$

Open problems (related to learning dynamics)

- other dynamics: best response to moving average, etc.
- other games (generalizations of Tullock contests)
  - aggregative games
  - Cournot games
  - diagonally strictly concave games [Rosen 65]
- learning with bandit feedback
  - agents know whether they win or not (i wins with probability $\frac{x_i}{\sum_j x_j}$)
  - but don't see others' actions
  - study Bayesian/statistical learning models

Extensions of Tullock contests

less well understood – open problems

Parallel contests

- m Tullock contests run in parallel
- each agent can play only one
- if agent i picks contest j, her utility

$$u_{i,j}(\ldots) = v_{i,j} \frac{x_{i,j}}{\sum_k x_{k,j}} - c_{i,j}(x_{i,j})$$

  - $x_{i,j}$: agent i's effort for contest j
  - $v_{i,j}$: agent i's value for contest j
  - $c_{i,j}$: agent i's cost function for contest j
- e.g., crowdsourcing, etc. (applications where agents have multiple options)
- upcoming work: existence/non-existence and computation of pure-strategy Nash equilibrium

Group contests

- (back to only one Tullock contest)
- partition agents into k groups: $G_1, G_2, \dots, G_k$
- if agent $i \in G_l$ her utility

$$u_i(\boldsymbol{x}) = \frac{\sum_{j \in G_l} x_j}{\sum_j x_j} - c_i(x_i)$$

- opportunity for free riding
- e.g., political party donation, upkeep of a blockchain system

Discrete action spaces

- instead of any effort $x_i \in \mathbf{R}_{\geq 0}$
- discrete actions $x_i \in X_i$, where X_i is a finite set of $\mathbf{R}_{\geq 0}$
- complexity of computing an equilibrium: open

Refernences

- Best-Response Dynamics in Lottery Contests
  Abheek Ghosh, Paul W. Goldberg. EC '23.
- Best-Response Dynamics in Tullock Contests with Convex Costs
  Abheek Ghosh. WINE ‘23.
- Continuous-Time Best-Response and Related Dynamics in Tullock Contests with Convex Costs
  Edith Elkind, Abheek Ghosh, Paul W. Goldberg. WINE '24.

Remark

general Tullock model (with concave utility), $r \le 1$

$$u_i(\boldsymbol{x}) = \frac{x_i^r}{\sum_j x_j^r} - c_i(x_i) \equiv \frac{y_i}{\sum_j y_j} - \bar{c}_i(y_i) = u_i(\boldsymbol{y})$$

equivalent to the model mentioned earlier by change of variables

- $x_i^r \to y_i$
- $c_i(x_i) \to c_i\left(y_i^{\frac{1}{r}}\right) \to \bar{c}_i(y_i)$ also convex as $r \le 1$