Learning to Bargain: Last-Iterate Convergence of Follow-the-Regularized-Leader in Games with a Discontinuity

MICHIGAN ENGINEERING
COMPUTER SCIENCE AND ENGINEERING
UNIVERSITY OF MICHIGAN

Serafina Kamp, Reese Liebman, Ben Fish

University of Michigan – Ann Arbor, Department of Computer Science Engineering

Motivation

Bargaining is a *fundamental strategic interaction* in many domains

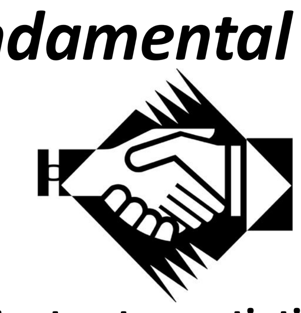
Contract negotiation


Price negotiation


Salary negotiation


Farmer's market shopping

Despite the existence of unique bargaining solutions, there is still an *equilibrium selection problem*, especially in *repeated bargaining games*

- Multiple Nash equilibria in simple bargaining games, multiple subgame perfect equilibria in bargaining with outside options
- Experiments from behavioral economics show people often deviate from unique SPE in the ultimatum game
- Folk theorem - there need not be a unique equilibrium when a stage game is repeated

Model

Ultimatum game

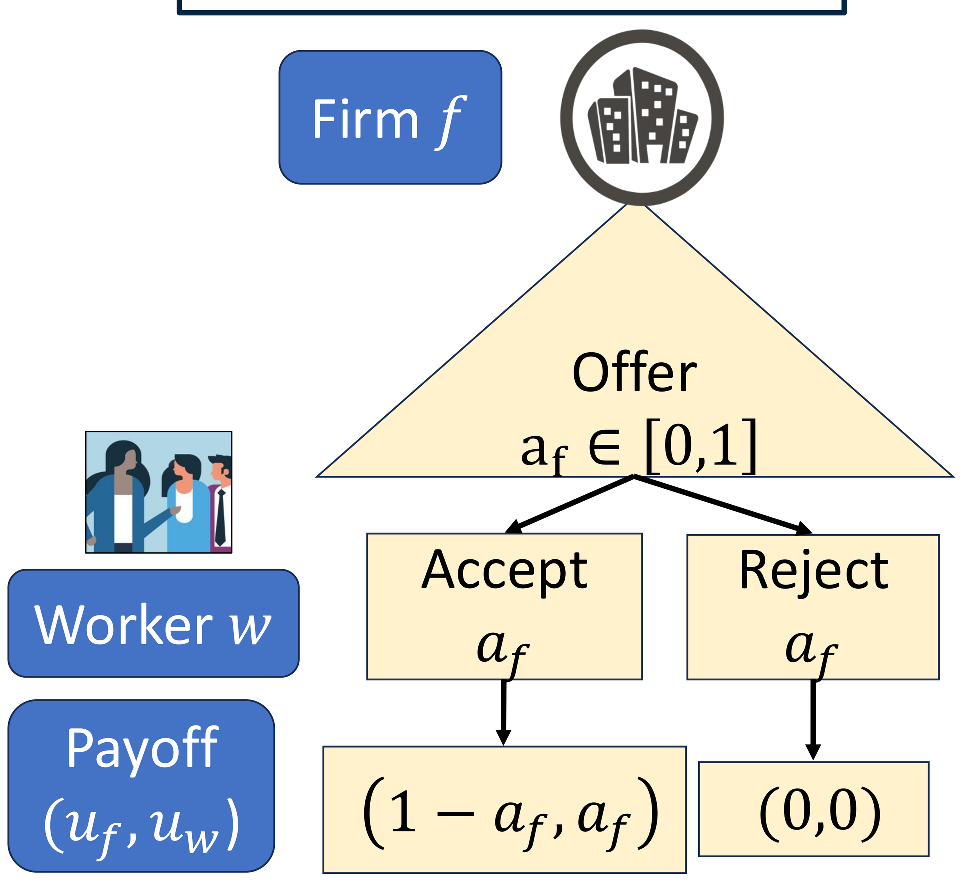

Assumptions

- Normal-form games with common action set [0,1]
- u_f and u_w are generalizations of payoffs of the ultimatum game, see figure.
- a_f is an "offer", a_w is an "acceptance threshold"

FTRL parameters:

- Action set $\mathcal{A}$ is D-discretization of [0,1], i.e., $\mathcal{A} = \{0, \frac{1}{D}, \frac{2}{D}, \ldots, 1\}$
- Learning rate η
- Reference point ρ

Find-the-Discontunity games, $\mathcal{D}$

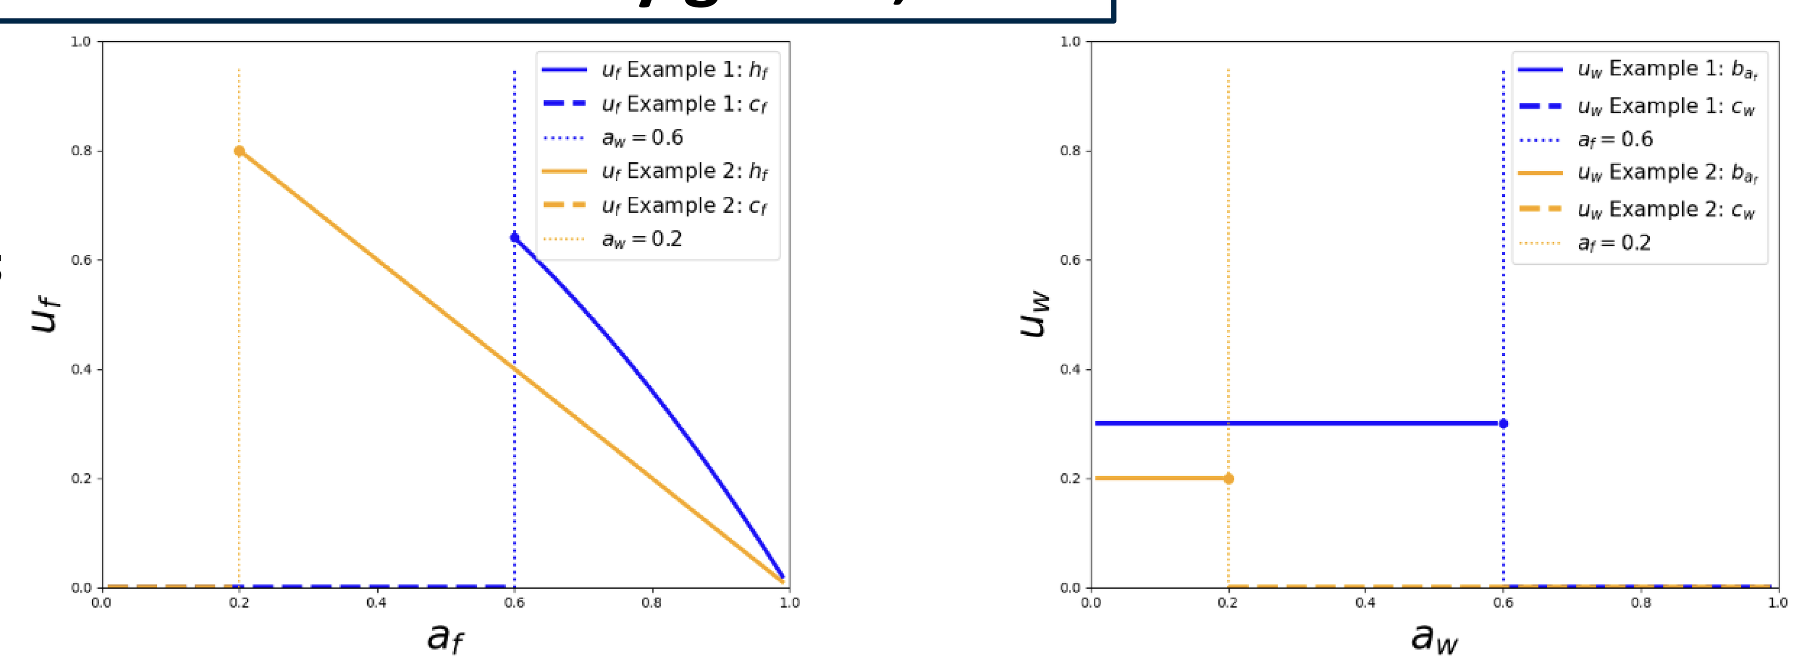

FTRL update step

$$x_i^{(t+1)} = \arg\max_{x\in\Delta(\mathcal{A}_i)} \eta\langle U_i^{(t)}, x\rangle - \frac{1}{2}\|x-\rho_i\|_2^2.$$

Mixed strategy

Cumulative expected utility vector

Results

Last-iterate convergence guarantees

Theorem (*informal*). Suppose agents learn strategies for a game that belongs to $\mathcal{D}$ using FTRL with $\eta > 0, \rho = \mathbf{0}$, and any initial mixed strategy. Then, the agents will converge to an ϵ-Nash equilibrium in finite time.

Mixed Nash equilibrium at convergence

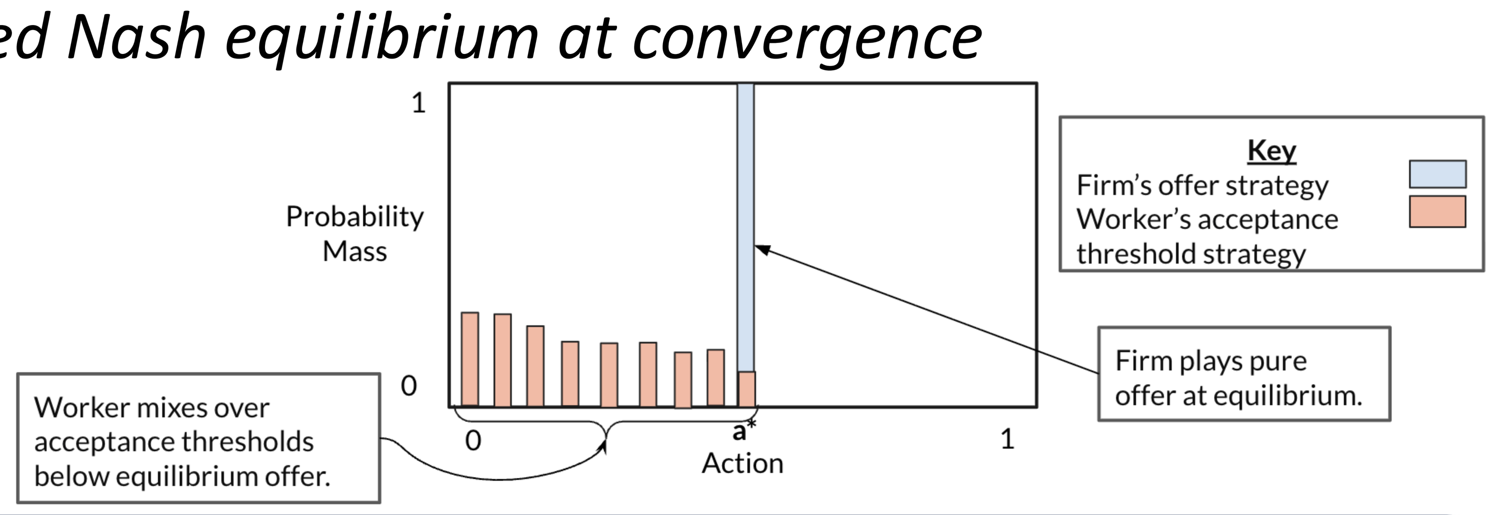

Theorem (*informal*). Suppose agents learn strategies for the ultimatum game using FTRL with $\eta > 0, \rho = \mathbf{0}$, and any initial mixed strategy. Then, the time to convergence to ϵ-Nash equilibrium is $O\left(\underbrace{\frac{D^{\frac{D}{\eta}^{\cdot^{\cdot^{D}}}}}{\eta}}_{3\cdot D} + \log\left(\frac{D^3}{\epsilon}\right)\cdot\frac{D^2}{\eta}\right)$.

Convergence proof sketch

Important quantities

- $w_{max} = max\{a | x_{w,a} > 0\}$, largest acceptance threshold with non-zero probability
- $f_{min} = \min\{a | x_{f,a} > 0\}$, smallest offer with non-zero probability
- For each $w_{max} = \frac{k}{D}$ for k ∈ $\{0,1,2,\ldots,D\}$, there exists a threshold γ_k such that, for f, the offer w_{max} is
  - *strictly dominant* when $x_{w,w_{max}} > \gamma_k$
  - *weakly dominant* when $x_{w,w_{max}} = \gamma_k$
  - *strictly dominated* when $x_{w,w_{max}} < \gamma_k$

State machine description of dynamics per w_{max} value

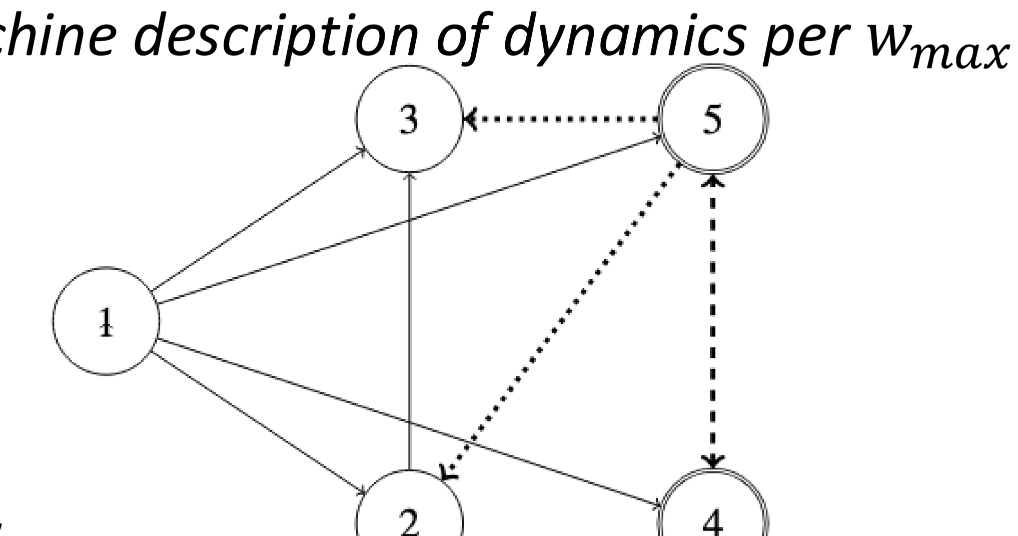

Arrow key

- Finite time transition
- Finitely many transitions, if convergence
- Finite time transition, if no convergence

State key

1	$f_{\min} > w_{\max}$
2	$f_{\min} = w_{\max}$, $x_{w,w_{\max}} < \gamma_k$
3	$f_{\min} < w_{\max}$, $x_{w,w_{\max}} < \gamma_k$
4	$f_{\min} = w_{\max}$, $x_{w,w_{\max}} \geq \gamma_k$
5	$f_{\min} < w_{\max}$, $x_{w,w_{\max}} \geq \gamma_k$

Numerical experiments

Different colored pixel <=> **different Nash equilibrium outcome**

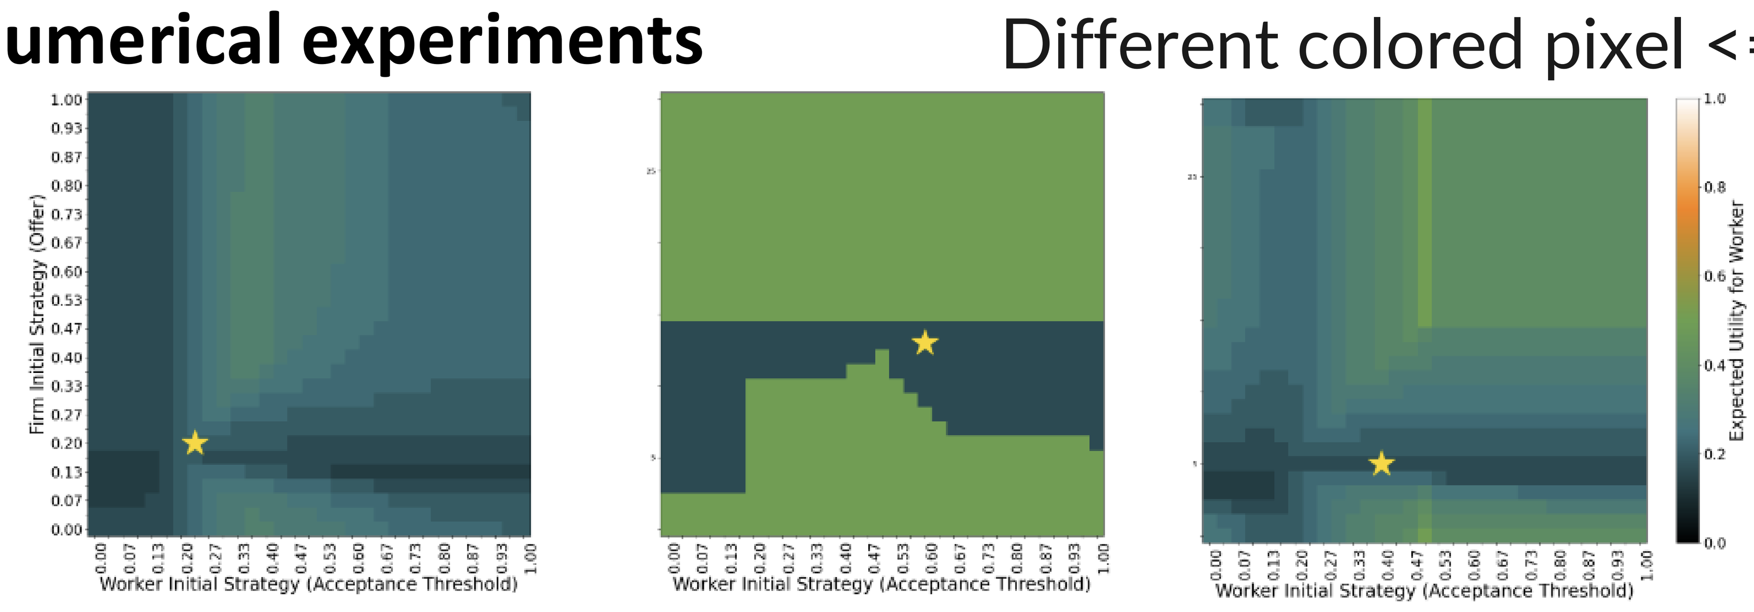

(a) $\rho_f = \rho_w = 0$. (b) $\rho_f = \frac{1}{6}, \rho_w = \frac{1}{2}$. (c) $\rho_f = \frac{1}{2}, \rho_w = \frac{29}{30}$.

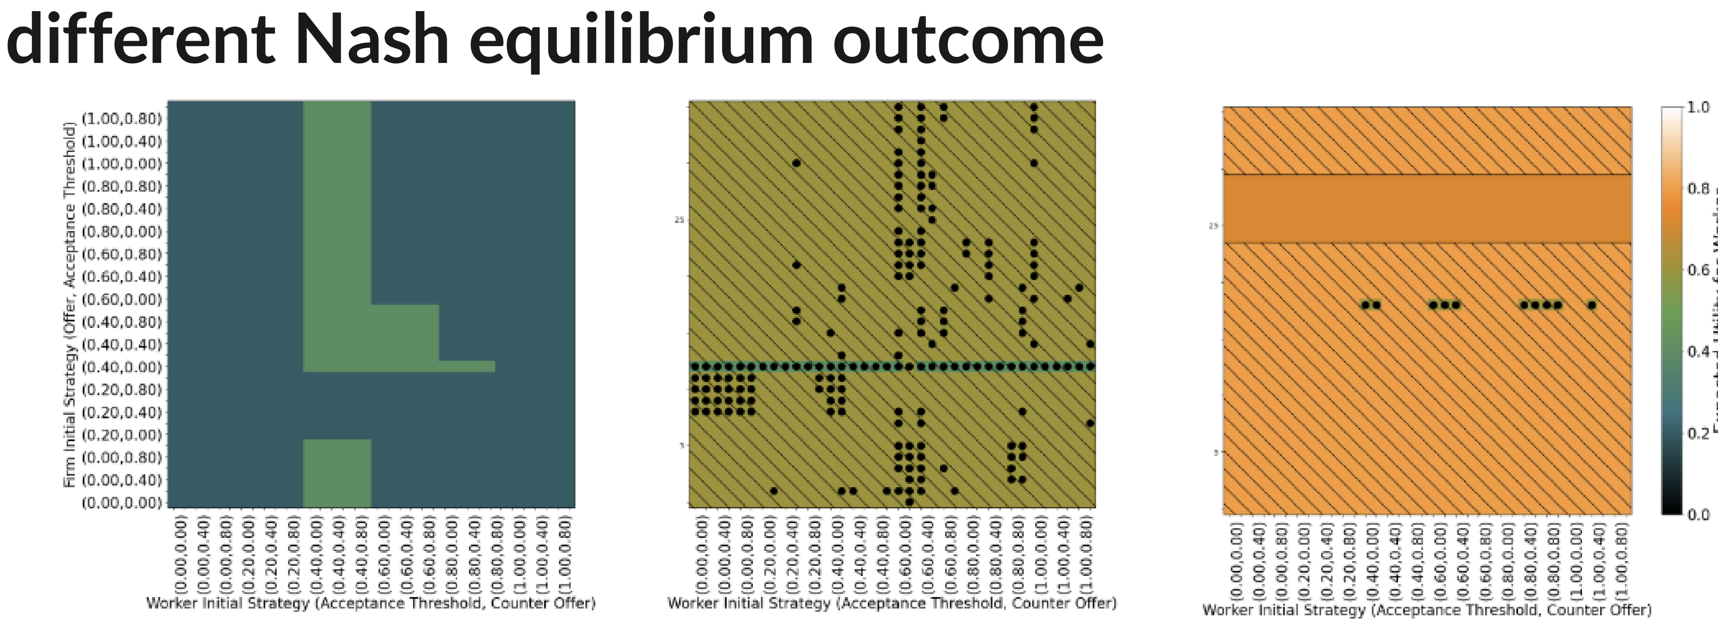

(a) $\delta = 0.1$. (b) $\delta = 0.55$. (c) $\delta = 0.9$.

Ultimatum game. Gold star is minimax equilibrium point of game over initial strategy choice. **2-round game.** *Credible* (diagonal hatching) and *non-credible* (circles) threats in the equilibrium strategy.