Convergence of Incentive-Driven Dynamics in Fisher Markets

Yuval Rabani
The Hebrew University of Jerusalem

joint work with:
Krishnamurthy Dvijotham and Leonard J. Schulman (Caltech)

Decisions in Markets

- “selfish” agents create, allocate, and exchange limited resources in a state of equilibrium
- knowledge, control distributed; no common goal; system is not engineered
- how is the system driven towards equilibrium?
- how is equilibrium restored in reaction to “shocks”?
- these are questions about the dynamics / stability

Game Theoretic Perspective

- dynamics: given a stage game, play it repeatedly
- this is just a longer game
- extensive form (perfect) equilibria -
  - compute in advance optimal moves in all situations
  - players play an equilibrium, thus a static notion
  - often intractable [Borgs et al. 2010; but Halpern et al. 2014]
  - unrealistically prescient players
  - what happens if the game changes unexpectedly?

Dynamics Wish List

- empirical evidence for moves and outcome
- strategic justification of moves and outcome
- moves undamped, not tailored for convergence
- convergence, to an attractive outcome, quickly
- adaptation to unpredictable evolution

Learning Dynamics

- no-regret dynamics -
  - converge to a correlated equilibrium [Hart & Mas-Collel 2000]
  - but reactions need to be damped carefully
  - what's the strategic justification? explicit form?
- best-response -
  - bounded rationality (myopic players)
  - desirable outcome? - in some cases [Awerbuch et al. 2008, Chien & Sinclair 2011]
- fictitious play [Brown 1951]
- logit dynamics [Blume 1993, …, Auletta et al. 2015]
- level k model (receding horizon control) -
  - extensive empirical evidence [Stahl & Wilson 1994-5; see Crawford et al. 2013]
  - adaptive and sophisticated learning [Milgrom & Roberts 1991]

Fisher Markets

- m perfectly divisible goods, each owned by a seller
- the endowment of each good is scaled to 1
- the demand is generated by n buyers with budgets and concave utility functions

Demand

- every good j has a varying price p_j
- every buyer i has a fixed budget b_i
- i's utility function $u_i(x_i)$ assigns a value to every basket of goods x_i (x_{ij} is the quantity of good j)
- i demands $x_i(p) = \text{argmax } u_i(x_i)$ subject to the budget constraint $\sum_j p_j x_{ij} \leq b_i$
- the demand for j at prices p is $x_{\cdot j}(p) = \sum_i x_{ij}(p)$
- equilibrium price vector p^*: $\forall j$, $x_{\cdot j}(p^*) = 1$ (the market clears)

unique in our setting, but not in general

Elasticity of Demand

- the price elasticity of the demand for j with respect to p_k is $\dfrac{\partial \log x_{\cdot j}(p)}{\partial \log p_k}$
- own elasticity: $j = k$; cross elasticity: $j \neq k$

- we assume (elastic and bounded demand):
  - cross price elasticities ≥ 0 (e.g., WGS utilities)
  - own price elasticities < -1 (incentive to clear market)
  - all prices $\leq p_{min} \Rightarrow$ all demands > 1
  - all prices $\geq p_{max} \Rightarrow$ all demands < 1

  } boundedness

$$0 < p_{min} < p_{max} < \infty$$

Price Updates

- each round: announce new prices and observe demand
  - synchronous - simultaneous price updates
  - asynchronous - arbitrary update schedule
  - ongoing markets [CF08] - also irregularly delayed market signals
- (discrete time) proportional tâtonnement [Walras 1874, Samuelson 1941]:

$$p_j^{t+1} = p_j^t \left(1 + \varepsilon \cdot (x_{\cdot j}(p^t) - 1)\right)$$

- empirical evidence (even when it doesn’t converge) [Plott et al. 2000-2011]
- In Fisher markets with CES+Cobb-Douglas+Leontief utilities:
  - converges (quickly) to equilibrium [Cole et al. 2008-2016]
  - equivalent to gradient/coord. descent [CCD13, CC16]
  - thus, regret minimizing
  - but, this requires careful choice of the damping factor ε

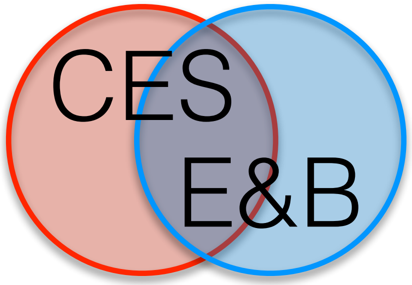

Best Response

- each seller j acts as follows:
    - predicts that the other sellers will stay put at the current prices $p_{-j} = (p_1, \ldots, p_{j-1}, \cdot, p_{j+1}, \ldots, p_m)$
    - sets its own price to a value $F_j(p)$ that maximizes predicted revenue

- own price elasticity $< -1 \Rightarrow x_{\cdot j}(p_{-j}, F_j(p)) = 1$

A General Framework

- A price update $p_j \mapsto F_j(p)$ is:
    - monotone, iff $p \geq q$ (coord.-wise) $\Rightarrow F_j(p) \geq F_j(q)$
    - sub-homogeneous, iff $\lambda \in (0,1) \Rightarrow \forall p$, $F_j(\lambda p) \geq \lambda F_j(p)$
    - strictly so, iff the inequality is strict for all $p > 0$
    - $[p_{min}, p_{max}]$-price bounded, iff $p \in [p_{min}, p_{max}]^m \Rightarrow F_j(p) \in [p_{min}, p_{max}]$
- F is … iff $\forall j$, F_j is …
- F is stable iff $F(p^*) = p^*$

Belief Formation

- $\mathscr{F}$ = finite set of (single seller) price updates

  for seller j, $f \in \mathscr{F}$ sets j's price given the other prices p_{-j}

- j's level 0 update: keep current price p_j
- j's level 0 belief on $s \neq j$: s uses a level 0 update.
- a level 1 update of j: use $f(p_{-j})$ for some $f \in \mathscr{F}$ (p_{-j} are level 0 beliefs of j)
- a level 1 belief of j on $s \neq j$: s uses a level ≤ 1 update.
- a level k update of j: use $f(q_{-j})$ for some $f \in \mathscr{F}$; q_{-j} are level $< k$ beliefs of j
- a level k belief of j on $s \neq j$: s uses a level $\leq k$ update.

Belief-Based Updates

Thm: Suppose that all $f \in \mathcal{F}$ are monotone, strictly sub-homogeneous, and bounded. Then $\forall k$ every level k update satisfies the same properties.

Our Main Theorem

Consider an update of prices

$$p \mapsto F(p) = (F_1(p), F_2(p), \ldots, F_m(p))$$

<u>Thm</u>: If F is monotone, strictly sub-homogeneous, and bounded, then it is a strict contraction under the Thompson metric.

<u>Thompson metric</u>:

$$d(p,q) = \| (\log(p_1 / q_1), \ldots, \log(p_m / q_m)) \|_\infty$$

Consequences

dynamic: $\mathscr{F}$ = monotone, strictly sub-homogeneous, bounded, and stable prices updates (finite set).

$$p^{t+1} = F^{\beta(t)}(p^t) = (F_1(p^t), F_2(p^t), \ldots, F_m(p^t))$$

F_j is a level $k_j(t) \geq 1$ update

$\beta_j(t)$ are the level $< k_j(t)$ beliefs of j that determine F_j

Thm: $\exists \xi_{max} < 1$ such that $d(p^t, p^*) < (\xi_{max})^t \cdot d(p^0, p^*)$

Corollary: $\| p^t - p^* \|_\infty < ((p_{max})^2 / p_{min}) \cdot (\xi_{max})^t \cdot d(p^0, p^*)$

$\| p^t - p^* \|_2 < \sqrt{n} \cdot ((p_{max})^2 / p_{min}) \cdot (\xi_{max})^t \cdot d(p^0, p^*)$

Concrete Markets

Thm: If the demand is elastic and bounded, then $F^{\beta(t)}$ is monotone, strictly sub-homogeneous, bounded, and stable.

$\Rightarrow$ the dynamic converges quickly to equilibrium

- beliefs not assumed to and cannot be consistent
- each seller believes: “I’m slightly smarter than the others” (and they believe the same thing about me)

Concluding Remarks

- applies also to asynchronous updates
  - t measures epochs
  - epoch = interval of ≥ 1 update of every price
- the worst case is best-response (to level 0 beliefs)
- more general applicability? (games?)
- using (noisy?) information to update beliefs?