Mean Field Games in Economics

Benjamin Moll

Princeton

IMA Mean Field Games and Applications Course
November 13, 2012

The Theme of my Talk

... TO HELP ME OUT WITH MY MATH PROBLEMS

Outline

(1) A class of Mean Field Games in economics: heterogeneous agent models

(2) Mean Field Games with aggregate shocks

(3) Surpassing language barriers when talking to economists

(4) "Boltzmann Mean Field Games"

Why Heterogeneous Agent Models?

- Many interesting questions involve some kind of **distribution**:

  (1) Does high income and wealth **inequality** lead to faster or slower GDP **growth**?

  (2) Does high income and wealth **inequality** increase or decrease GDP **volatility** (severity of recessions)?

  (3) Where does the **firm size distribution** come from? Why does it follow a power law?

- To theorize about these, need heterogeneous agent models.
- Will argue: mathematics are somewhat underdeveloped, potentially high payoff from well-trained mathematicians working on these.

A Prototypical Heterogeneous Agent Model

- Based on Rao Aiyagari (1994), "Uninsured Idiosyncratic Risk and Aggregate Saving", Quarterly Journal of Economics
- Taught in first year of every self-respecting economics PhD program
- Original paper: discrete time; here: continuous time
- Two types of agents: households and firms

Households

- are heterogeneous in their wealth k and work ability z (joint distribution $g(k, z, t)$), solve

$$\max_{c_t} \mathbb{E}_0 \int_0^\infty e^{-\rho t} u(c_t) dt \qquad \text{s.t.}$$

$$dk_t = [w_t z_t + r_t k_t - c_t] dt$$

$$dz_t = \mu^z(z_t) dt + \sigma^z(z_t) dW_t$$

$$k_t \geq \underline{k}$$

  - c_t: consumption
  - u: utility function, $u' > 0, u'' < 0$.
  - ρ: discount rate
  - w_t: wage
  - r_t : interest rate
  - W_t: standard Brownian motion, independent across households
  - $\underline{k}$: borrowing limit, e.g. if $\underline{k} = 0$, can only save

Households: Simplified Version

- Previous slide: taken one-for-one from Aiyagari.
- Here: simplify further
- Assume $w_t z_t dt$ iid normal over time with mean $w_t dt$ and variance $\sigma^2 dt$
- Households solve

$$\max_{c_t} \mathbb{E}_0 \int_0^\infty e^{-\rho t} u(c_t) dt \qquad \text{s.t.}$$

$$dk_t = [w_t + r_t k_t - c_t]dt + \sigma dW_t$$

$$k_t \geq \underline{k}.$$

- Advantage: only one state variable k
- (bit weird: unbounded labor income shocks, even negative)

Firms

- All identical, solve

$$\max_{K_t, L_t} \; AF(K_t, L_t) - (r_t + \delta)K_t - w_t L_t$$

  - A: productivity, fixed for now
  - K_t: capital
  - L_t: labor
  - F: production function, increasing and concave
  - δ: depreciation rate, $r_t + \delta$: user cost of capital

- In equilibrium

$$L_t = 1, \qquad K_t = \int kg(k, t)dk$$

Equilibrium

- Prices are

$$r_t = AF_K(K_t, 1) - \delta, \quad w_t = AF_L(K_t, 1), \quad K_t = \int kg(k,t)dk$$

- Households' HJB equation:

$$\rho v(k,t) = \max_c \ u(c) + v_k(k,t)[w_t + r_t k - c] + \frac{\sigma^2}{2} v_{kk}(k,t) + v_t(k,t)$$

- Kolmogorov-Forward (Fokker-Planck) eq. for wealth dist.

$$\frac{\partial g(k,t)}{\partial t} = -\frac{\partial}{\partial k}[\mu^k(k,t)g(k,t)] + \frac{\sigma^2}{2}\frac{\partial^2}{\partial k^2} g(k,t)$$

where $c(k,t)$ is optimal consumption choice from HJB and

$$\mu^k(k,t) = w_t + r_t k - c(k,t)$$

- A mean field game! Interactions through prices.
- Here: prices only depend on mean. But easy to think of other setups where also other features of distribution matter.

Solution in Practice

- Assume some parametric functional forms for u and F
- Most commonly used are "CRRA utility" and "Cobb-Douglas production function"

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma}, \quad \gamma > 0, \qquad F(K, L) = K^{\alpha} L^{1-\alpha}, \quad \alpha \in (0, 1)$$

- Therefore prices entering household's HJB eq. are

$$r_t = \alpha A K_t^{\alpha-1} - \delta, \quad w_t = (1-\alpha) A K_t^{\alpha}, \quad K_t = \int k g(k, t) dk$$

- Note: different from linear-quadratic examples mathematicians like to use
- Borrowing constraint: $\underline{k} = 0$.
- Solve numerically. More on this momentarily.

Open Questions

(1) Existence of equilibrium

(2) Uniqueness of equilibrium

(3) Also some scope for improving numerical solution.

Open Questions

- Most existing approaches focus on stationary equilibrium only.
- Typical algorithm for solving stationary equilibrium uses "Monte Carlo method":
  1. Guess interest rate, r
  2. Compute households' optimal decision rules from dynamic programming problem
  3. Simulate N (say 50,000) individuals and compute aggregate capital demand and supply
  4. If demand=supply, stop. Otherwise, update r using bisection method and go back to 1.
  5. Cross fingers, pray and hope that algorithm converges
- Actually works surprisingly well in practice. Still unsatisfactory. Non-stationary equilibria even trickier.
- Some simple transparent MFG FD method much preferred.

Mean Field Games with Aggregate Shocks

Mean Field Games with Aggregate Shocks

- This is where I think the real payoff could be
- Consider Aiyagari (1994) with shocks to aggregate productivity, A.
- First studied by Per Krusell and Tony Smith (1998), "Income and Wealth Heterogeneity in the Macroeconomy", Journal of Political Economy

Mean Field Games with Aggregate Shocks

- Households:

$$\max_{c_t} \mathbb{E}_0 \int_0^\infty e^{-\rho t} u(c_t) dt \qquad \text{s.t.}$$

$$dk_t = [w_t + r_t k_t - c_t]dt + \sigma dW_t$$

$$k_t \geq \underline{k}.$$

- Firms:

$$\max_{K_t, L_t} A_t F(K_t, L_t) - (r_t + \delta) K_t - w_t L_t$$

$$dA_t = \mu^A(A_t)dt + \sigma^A(A_t)dZ_t$$

- Equilibrium:

$$L_t = 1, \qquad K_t = \int k g(k,t) dk$$

Mean Field Games with Aggregate Shocks

- Question: What is the appropriate state variable?
- Answer: it's really the entire wealth distribution (plus aggregate productivity)

$$\text{state} = (g, A)$$

- Problem: g is infinite-dimensional.
- Why not enough to keep track of first moment $K_t = \int kg(k,t)dk$? Answer: evolution dK_t depends on entire g, not just K_t (unless all individual dk_t's are linear in k_t).
- Nicely explained in Victor Rios-Rull (1997), "Computation of Equilibria in Heterogeneous Agent Models"

Mean Field Games with Aggregate Shocks

- Krusell-Smith get around this with an approximation: households care only about some finite set of moments of wealth distribution (in practice, only first moment):

$$v(k; g, A) \approx v(k; K, A)$$

- $\Rightarrow$ convert infinite-dimensional problem into finite-dim. one.
- Makes some sort of sense: mean field type interaction only through prices (r, w). Assumption is that households only take into account future price movements that are due to movements of the aggregate capital stock (and A). Bounded rationality interpretation.
- But obviously still unsatisfactory ▸ KS on accuracy of their algorithm

What I'm Hoping For

- Someone comes up with a tractable, MFG version of Krusell-Smith in continuous time
- tractable = tractable enough to solve globally without approximations/bounded rationality assumptions. Even better: tractable enough to prove existence and uniqueness
- Can publish that anywhere!
- But even true for MFG version of Aiyagari (1994), i.e. without aggregate shocks

Language Barriers

Surpassing Language Barriers when Talking to Economists

- **Economists don't know ∇, Δ !**
- You like to write

$$\frac{\partial u}{\partial t} + \frac{\sigma^2}{2}\Delta u + H(x, m, \nabla u) - \rho u = 0 \qquad \text{(HJB)}$$

$$\frac{\partial m}{\partial t} = \nabla \cdot \left(m \frac{\partial H}{\partial p}(x, m, \Delta u) \right) + \frac{\sigma^2}{2}\Delta m \qquad \text{(KFE)}$$

where $H(x, m, p) = \max_a h(x, m, a) + pa$.

Talking to Economists

- If you talk to economists, may want to write

$$\rho u(x,t) = H(x, m, u_x(x,t)) + \frac{\sigma^2}{2} u_{xx}(x,t) + \frac{\partial u(x,t)}{\partial t} \qquad \text{(HJB)}$$

$$\frac{\partial m(x,t)}{\partial t} = -\frac{\partial}{\partial x}[m(x,t) H_p(x, m, u_x(x,t))] + \frac{\sigma^2}{2} m_{xx}(x,t) \qquad \text{(KFE)}$$

- Or even better get rid of the Hamiltonian

$$\rho u(x,t) = \max_a \; h(x, m, a) + u_x(x,t)a + \frac{\sigma^2}{2} u_{xx}(x,t) + \frac{\partial u(x,t)}{\partial t} \qquad \text{(HJB)}$$

$$\frac{\partial m(x,t)}{\partial t} = -\frac{\partial}{\partial x}[m(x,t) a^*(x,m)] + \frac{\sigma^2}{2} m_{xx}(x,t) \qquad \text{(KFE)}$$

where $a^*(x,m)$ is maximand in (HJB)

Another Observation

- MFG literature seems to often focus on problems with a separable Hamiltonian

$$H(x,p) + V(m)$$

- This is of very limited use to economists because most problems in economics do not have separable Hamiltonians
- E.g. in Aiyagari and Krusell-Smith

$$H(x,m,p) = \max_c \; u(c) + p[w(m) + r(m)x - c]$$

which is not separable.

Boltzmann Mean Field Games

A Theory of Growth from Knowledge Diffusion

- Based on Lucas and Moll (2012), "Knowledge Growth and the Allocation of Time"
- Consider economy with continuum of infinitely-lived agents
- Productivity of each is a random variable, $a \sim$ density g, cdf G. Entire distribution is state variable of economy.
- Convenient to work with "cost" rather than productivity
- Let $a = z^{-\theta}, \ \theta \in (0, 1)$, where $z \sim$ density f, cdf F
- Can always back out

$$G(a, t) = 1 - F(a^{-1/\theta}, t)$$

- Computations use z; figures use a.

A Technology of Learning

- $s(z,t)$: fraction of time searching.
- $1 - s(z,t)$: fraction of time working
- Earnings:

$$y(z,t) = [1 - s(z,t)]z^{-\theta}$$

- Per capita GDP:

$$Y(t) = \int_0^\infty [1 - s(z,t)]z^{-\theta} f(z,t)dz$$

- How do individual productivities evolve?

Evolution of Productivity Distribution

- Allocate $s(z,t)$ to search.
- Over $(t, t+h)$, meet one other agent with probability $\alpha\,[s(z,t)]\,h$.
- Meeting = draw $\tilde{z}$ from $f(\tilde{z}, t)$.
- $z(t+h) = \min\{\tilde{z}, z(t)\}$.
- Meetings completely *asymmetric*.
- Number of agents at z:

$$\frac{\partial f(z,t)}{\partial t} = \left.\frac{\partial f(z,t)}{\partial t}\right|_{\text{out}} + \left.\frac{\partial f(z,t)}{\partial t}\right|_{\text{in}}$$

Evolution of Productivity Distribution

- Outflow: Agent z adopts lower cost if he meets another with cost $y \leq z$.
- Happens w/ prob $\alpha(s(z,t))f(y,t)$. Hence

$$\left.\frac{\partial f(z,t)}{\partial t}\right|_{\text{out}} = -\alpha(s(z,t))f(z,t)\int_0^z f(y,t)dy$$

- Inflow: agent $y \geq z$ adopts z if he meets agent at z.
- Happens w/ prob $\alpha(s(y,t))f(z,t)$. Hence

$$\left.\frac{\partial f(z,t)}{\partial t}\right|_{\text{in}} = f(z,t)\int_z^\infty \alpha(s(y,t))f(y,t)dy$$

Evolution of Productivity Distribution

- Combining ins and outs we have

$$\frac{\partial f(z,t)}{\partial t} = -\alpha(s(z,t))f(z,t)\int_0^z f(y,t)dy + f(z,t)\int_z^\infty \alpha(s(y,t))f(y,t)dy$$

- A Boltzmann equation.
- Many possible generalizations, present two later.

Decentralized Time Allocation Decisions

- Assume that agents maximize expected PV or earnings, discounted at given $\rho > 0$
- No risk aversion, no theory of real interest rates: a kind of partial equilibrium
- Individual preferences:

$$V(z,t) = \mathbb{E}_t \left\{ \int_t^{\infty} e^{-\rho(\tau - t)} \left[1 - s(\tilde{z}(\tau), \tau)\right] \tilde{z}(\tau)^{-\theta} d\tau \middle| z(t) = z \right\}$$

- Poisson arrival rate $\alpha(s) \geq 0$ with

$$\alpha'(s) > 0, \quad \alpha''(s) < 0, \quad \text{all } s.$$

Equilibrium

- Bellman equation:

$$\rho V(z,t) - \frac{\partial V(z,t)}{\partial t} = \max_{s \in [0,1]} \left\{ (1-s) z^{-\theta} + \alpha(s) \int_0^z [V(y,t) - V(z,t)] f(y,t) dy \right\}$$

- Boltzmann equation:

$$\frac{\partial f(z,t)}{\partial t} = -\alpha(s(z,t)) f(z,t) \int_0^z f(y,t) dy + f(z,t) \int_z^\infty \alpha(s(y,t)) f(y,t) dy.$$

- A **"Boltzmann mean field game"**!

Balanced Growth Path

Definition

A *balanced growth path (BGP)* is a number γ and a triple of functions (ϕ, σ, v) on $\mathbb{R}_+$ such that

$$f(z,t) = e^{\gamma t}\phi(ze^{\gamma t}),$$

$$V(z,t) = e^{\theta\gamma t}v(ze^{\gamma t}),$$

$$s(z,t) = \sigma(ze^{\gamma t})$$

- Intuitively: all z-quantiles shrink at γ ($\Rightarrow$ all $z^{-\theta}$-quantiles grow at $\theta\gamma$).
- CDF $F(z,t) = \Phi(ze^{\gamma t}) \Rightarrow q$th quantile defined by

$$\Phi(z_q(t)e^{\gamma t}) = q \quad \Leftrightarrow \quad z_q(t) = e^{-\gamma t}\Phi^{-1}(q).$$

Balanced Growth Path

- Restate (BE), (LM) for BGP only. Use $x = ze^{\gamma t}$

$$(\rho - \theta\gamma)\, v(x) - v'(x)\gamma x = \max_{\sigma \in [0,1]} \left\{ (1-\sigma)x^{-\theta} + \alpha(\sigma) \int_0^x [v(y) - v(x)]\phi(y)\, dy \right\}$$

$$\phi(x)\gamma + \phi'(x)\gamma x = \phi(x) \int_x^\infty \alpha(\sigma(y))\phi(y)dy - \alpha(\sigma(x))\phi(x) \int_0^x \phi(y)dy.$$

- Growth rate?

$$\phi(0)\gamma = \phi(0) \int_0^\infty \alpha(\sigma(y))\phi(y)dy.$$

$$\gamma = \int_0^\infty \alpha(\sigma(y))\phi(y)dy$$

Balanced Growth Path

- BGP with $\gamma > 0$ only if $\phi(0) = f(0,0) > 0$.

Lemma

$\phi(0) = f(0,0) > 0 \Leftrightarrow G(a,0)$ *has Pareto tail with parameter* $1/\theta$.

- Taken literally: all knowledge already exists at $t = 0$.
- Alternative interpretation: $G(a,0)$ is bounded but new knowledge arrives at arbitrarily low frequency ("innovation").
- Observationally equivalent, see section 6.3 in paper

Balanced Growth Path

- Next slide presents results of single simulation of BGP
  - Initial cost distribution is exponential: $\phi_0(x) = \lambda e^{-\lambda x}$
  - Function $\alpha(\sigma)$ given by $\alpha(\sigma) = k\sigma^{\eta}$
- Display time allocation function $\sigma(x)$, density $\phi(x)$, and two Lorenz curves
  - income Lorenz curve, based on $(1 - \sigma(x))\, x^{-\theta}$
  - value Lorenz curve, based on calculated value function, $v(x)$

Balanced Growth Path

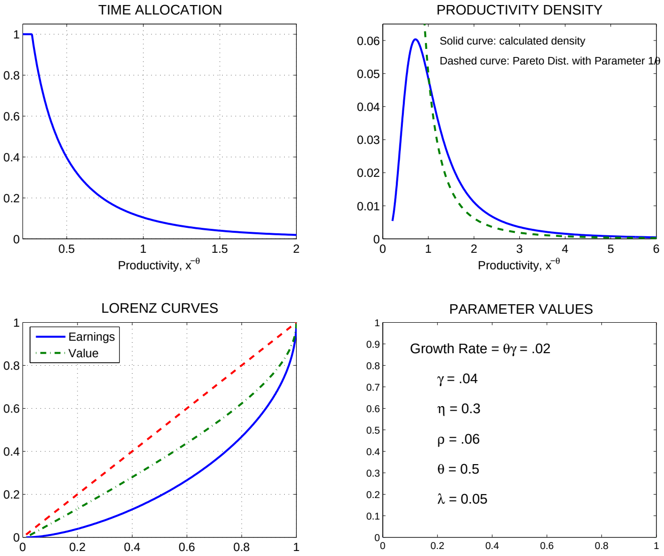

Two Sample Paths

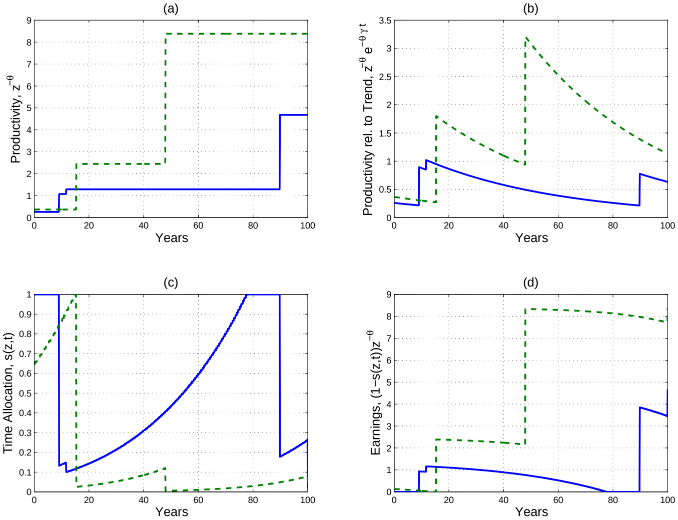

Alternative Learning Technologies

- Learning technology described above involves
  - probabilistic model of agents' meetings
  - description of effects of meetings on agents' knowledge
- Easy to think of modifications, other Boltzman equations
- Algorithm very flexible. Beginning to explore possible variations

Limits to Learning

- Impose an order to learning, limits on intellectual range
  - If y meets $z < y$ at t, he can adopt z with given probability $k(y, z, t)$
  - w/prob. $1 - k(y, z, t)$ he cannot do this; retains cost y
- Natural assumption: k unchanged if z and y at same *quantiles*. Along BGP this requires

$$k(y, z, t) = k(ye^{\gamma t}, ze^{\gamma t}, 0).$$

- Convenient to work with

$$k(y, z, 0) = e^{-\kappa |y-z|}$$

- Alternative interpretation: *meeting* probabilities depend on $|y - z|$; socioeconomic stratification or segregation.

Balanced Growth Path

- BGP equations become

$$(\rho - \theta\gamma)\, v(x) - v'(x)\gamma x = \max_{\sigma \in [0,1]} \left\{ (1-\sigma)x^{-\theta} + \alpha(\sigma) \int_0^x [v(y) - v(x)]\phi(y) e^{-\kappa(x-y)} dy \right\}$$

$$\phi(x)\gamma + \phi'(x)\gamma x = \phi(x) \int_x^\infty \alpha(\sigma(y))\phi(y) e^{-\kappa(y-x)} dy$$

$$- \alpha(\sigma(x))\phi(x) \int_0^x \phi(y) e^{-\kappa(x-y)} dy$$

$$\gamma = \int_0^\infty \alpha(\sigma(y))\phi(y) e^{-\kappa y} dy$$

Optimal Time Allocation for Various κ Values

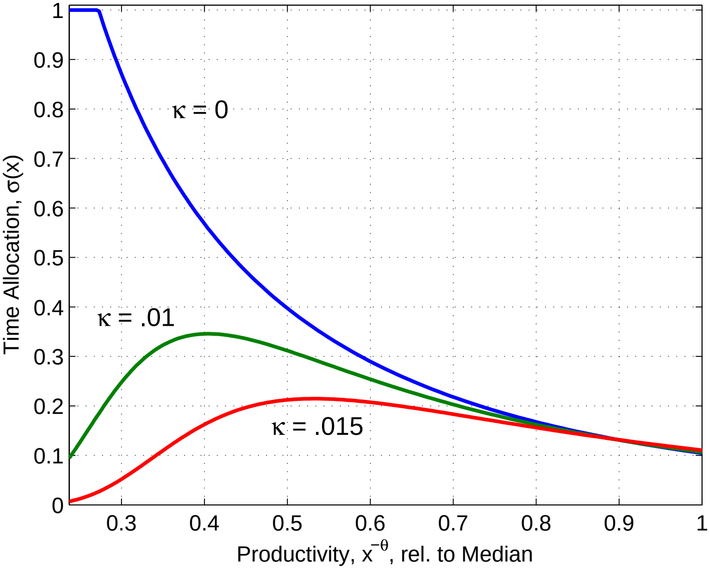

Limits to Learning Reduce Equality and Growth

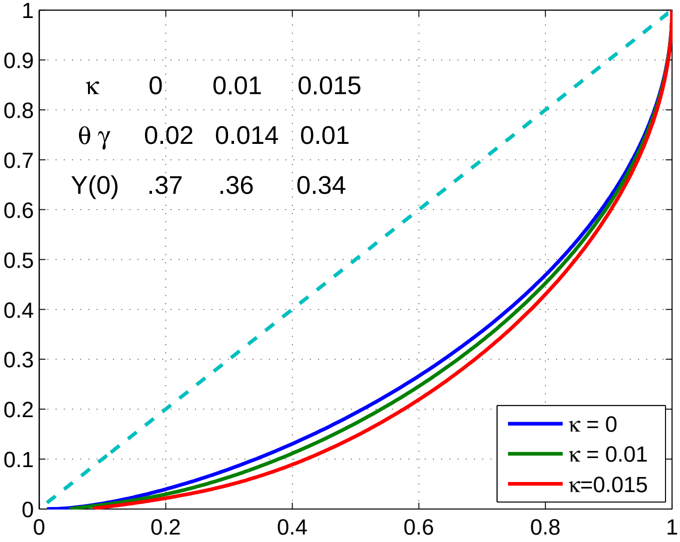

Limits to Learning Reduce Equality and Growth

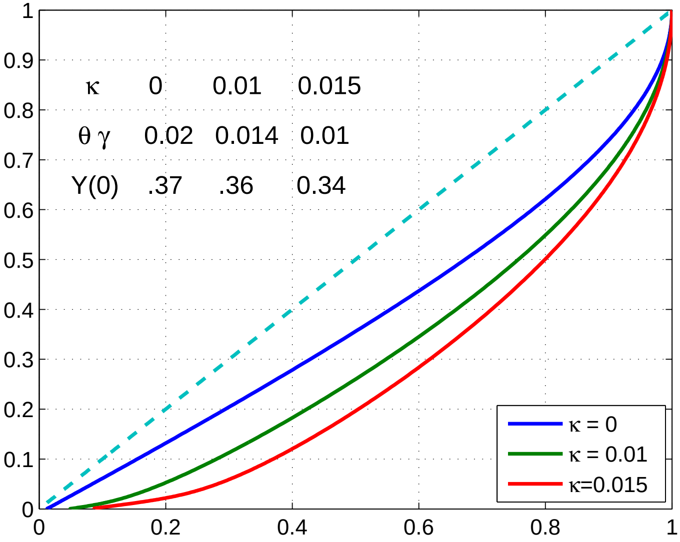

Conclusion

- Mean field games potentially extremely useful in economics...
- ... lots of exciting questions involve mean field type interactions...
- ... but mathematics often pretty challenging, at least for the average economist.
- Potentially high payoff from mathematicians working on this!

Krusell-Smith on Accuracy of their Algorithm ▸ Back

- p.897: *"Like most numerical procedures, the present one does not provide bounds on how far the approximate equilibrium deviates from an exact equilibrium. In particular, one might imagine that there are self-fulfilling approximate equilibria: because agents perceive a simple law of motion, they behave accordingly."*