Mean Field Consumption-Accumulation Optimization with HARA Utility

Minyi Huang

School of Mathematics and Statistics
Carleton University, Ottawa

Workshop on Mean Field Games and Related Topics 2

Padova, September 4-7, 2013

Some background

- Mean field games: Competitive decision with a large number of agents
  - The fast growing literature for MFGs
- Stochastic growth theory: Consumption and investment optimization
  - Optimal control of a whole sector of an economy
  - More generally: Nash games of N producers
  - The pioneering work (Brock and Mirman, J. Econ. Theory, 1972); a nice survey (Olson and Roy, 2006)

Our plan

- Endogenize the mean field effect into the production dynamics
  - Thus generalize the conventional modeling
- Decentralized optimization, closed-form solution
- Illustrate long time nonlinear behaviour
  - The success of mean field games so far relies on the fact: *Agents can rationally anticipate the mean field behaviour*
  - Can the above be assumed without questioning?
  - The model here shows new challenges. Very different from ergodic stochastic control

Note from the computational aspect

- Except the LQG/LQEG cases (see e.g. the work of Bardi; Bensoussan et al; Caines, Huang and Malhamé; Li and Zhang; Tembine at al.; ......), closed form solutions in MFGs are rare.
- Guéant, Lasry and Lions (2011) obtained closed form solutions in a human capital optimization model by considering certain linear control classes.
- We look for closed-form solutions in a new class of models.
  - Stepping on this, we can address deeper issues.

Classical stochastic one-sector growth model: Review

The one-sector economy at stage t involves two basic quantities:

- κ_t: the capital stock (used for investment)
- c_t: consumption

The next stage output y_{t+1}:

$$y_{t+1} = f(\kappa_t, r_t), \quad t = 0, 1, \ldots,$$

- $f(\cdot, \cdot)$: called the production function
- r_t: random disturbance
- y_0: the initial output, given
- If the output remaining after investment is all consumed, one has the constraint $\kappa_t + c_t = y_t$

Objective: maximize the utility functional

$$E \sum_{t=0}^{\infty} \rho^t \nu(c_t),$$

- $\nu(c_t)$: utility from consumption, usually concave on $[0, \infty)$

Issues of interest: sustainability, extinction (like whale hunting), business cycles and other oscillatory behaviors, etc.

The model is closely related to the theory of optimal savings and portfolio optimization which usually consider linear production functions.

Brock and Mirman (J. Econ. Theory, 1972) pioneered stochastic growth theory.

Notation in the mean field model

Keep track of the notation (for the main part):

u_t^i:	control (investment)
X_t^i:	state (production output)
N:	number of players in the game
c_t^i:	consumption
$V_i(x,t)$:	value function
$G(p,W), g$:	growth coefficient in production
W:	white noise
p:	aggregate investment
γ:	HARA utility exponent

Mean field production dynamics of N agents

- X_t^i: output (or wealth) of agent i, $1 \le i \le N$
- $u_t^i \in [0, X_t^i]$: investment (so no borrowing)
- $c_t^i = X_t^i - u_t^i$: amount for consumption
- $u_t^{(N)} = (1/N)\sum_{j=1}^N u_t^j$: aggregate investment level

The next stage output, measured by the unit of capital, is

$$X_{t+1}^i = G(u_t^{(N)}, W_t^i)u_t^i, \qquad t \ge 0, \tag{3.1}$$

Regard $u_t^{(N)}$ as a quantity measured according to a macroscopic unit.

Motivation for the mean field production dynamics

- Some literature addresses negative externalities.
- For example, Barro and Sala-I-Martin (Rev. Econ. Stud., 1992) considered the so-called congestion effect

$$y = Ak(G/K)^{\alpha}, \quad \alpha \in (0,1)$$

  - y: output of a private producer; k: its input;
  - K: aggregate input of n producers;
  - G: total resource provided by the government.

For similar models, see (Liu and Turnovsky, J. Pub. Econ., 2005)

The utility functional

The utility functional is

$$J_i(u^i, u^{-i}) = E\sum_{t=0}^{T} \rho^t v(X_t^i - u_t^i),$$

- $\rho \in (0, 1]$: the discount factor
- $c_t^i = X_t^i - u_t^i$: consumption, $\quad u^{-i} = (\cdots, u^{i-1}, u^{i+1}, \cdots)$

We take the HARA utility

$$v(z) = \frac{1}{\gamma} z^\gamma, \quad z \geq 0, \qquad \gamma \in (0, 1).$$

Assumptions

(A1) (i) Each sequence $\{W_t^i, t \in \mathbb{Z}_+\}$ consists of i.i.d. random variables with support D_W and distribution function F_W. The N sequences $\{W_t^i, t \in \mathbb{Z}_+\}$, $i = 1, \ldots, N$ are i.i.d. (ii) The initial states $\{X_0^i, 1 \le i \le N\}$ are i.i.d. positive random variables with distribution F_{X_0} and mean m_0, which are also independent of the N noise sequences.

(A2) (i) The function G: $[0, \infty) \times D_W \to [0, \infty)$ is continuous; (ii) for a fixed $w \in D_W$, $G(z, w)$ is a decreasing function of $z \in [0, \infty)$.

(A3) (iii) $EG(0, W) < \infty$ and $EG(p, W) > 0$ for each $p \in [0, \infty)$.

(A2) implies congestion effect: when the aggregate investment level increases, the production becomes less efficient.

Example

Example. Suppose $G(z, w) = \frac{\alpha w}{1+\delta z^\eta}$, where $\alpha > 0, \delta > 0, \eta > 0$ are parameters.

- The congestion effect is indicated by the monotonic decreasing property of G

Further interpretation of (A2):

- Diminishing return of the sector (although no notable d.r. to individual scale)
- Illustration: Suppose the production structure

$$G(z, W) = g(z)W, \qquad g \text{ decreasing.}$$

Let $X_t^{(N)} = \frac{1}{N}\sum_{i=1}^N X_t^i$. Then

$$X_t^{(N)} = g(u_t^{(N)})u_t^{(N)}EW_0^1 + \frac{g(u_t^{(N)})}{N}\sum_{i=1}^N u_t^i(W_t^i - EW_t^i)$$
$$= g(u_t^{(N)})u_t^{(N)}EW_0^1 + \text{“small fluctuation”}$$

To do (in the Nash certainty equivalence approach)

- Step 1 – Approximate the aggregate investment $(u_t^{(N)})_{t=0}^{T-1}$ by a sequence $(p_0, \ldots, p_{T-1})$.
- Step 2 – Solve an optimal control problem with a representative agent
  - Obtain control $\hat{u}_t^i = \alpha_t \hat{X}_t^i$ (linear feedback).
  - Determine closed loop state $\hat{X}_t^i$.
- Step 3 – Consistency requirement

$$\lim_{N\to\infty}(1/N)\sum_{i=1}^{N}\hat{u}_t^i = \lim_{N\to\infty}(1/N)\sum_{i=1}^{N}\alpha_t\hat{X}_t^i = \alpha_t E\hat{X}_t^i = p_t.$$

- Step 4 – Show ε_N-Nash property for $(\hat{u}_t^1, \ldots, \hat{u}_t^N)$.

Step 1: mean field limit

Now agent i considers the optimal control problem with dynamics

$$X_{t+1}^i = G(p_t, W_t^i)u_t^i, \qquad t \geq 0, \tag{3.2}$$

where $u_t^i \in [0, X_t^i]$.

The utility functional is now written as

$$\bar{J}_i(u^i, (p_t)_0^{T-1}, 0) = E \sum_{t=0}^{T} \rho^t v(X_t^i - u_t^i), \tag{3.3}$$

Step 2: optimal control

Further denote the utility functional with initial time t

$$\bar{J}_i(u^i, (p_l)_t^{T-1}, t) = E\sum_{s=t}^{T} \rho^{s-t} v(X_s^i - u_s^i),$$

which is affected only by $(p_l)_t^{T-1}$. The value function is defined as

$$V_i(x,t) = \sup_{\{u_s^i\}_{s=t}^{T}} E[\bar{J}_i(u^i, (p_l)_t^{T-1}, t) | X_t^i = x].$$

We have the dynamic programming equation

$$V_i(x,t) = \max_{0\le u_i \le x} \left[v(x-u_i) + \rho E V_i(G(p_t, W_t^i)u_i,\ t+1) \right],$$

where $t = 0, 1, \ldots, T-1$.

Step 2

We look for a solution

$$V_i(x,t) = \frac{1}{\gamma} D_t^{\gamma-1} x^{\gamma}. \quad (3.4)$$

Denote $\Phi(z) = \rho E G^{\gamma}(z, W)$ and $\phi(z) = \Phi^{\frac{1}{\gamma-1}}(z)$.

Theorem *(i)* The value function $V_i(x,t)$ takes the form (3.4), where

$$D_t = \frac{\phi(p_t) D_{t+1}}{1 + \phi(p_t) D_{t+1}}, \quad t \le T-1, \quad D_T = 1. \quad (3.5)$$

(ii) The optimal control has the feedback form

$$u_t^i = \frac{X_t^i}{1 + \phi(p_t) D_{t+1}}, \quad t \le T-1, \quad u_T^i = 0. \quad (3.6)$$

Step 3: consistency

We have the closed-loop

$$X_t^i = G(p_t, W_t^i)\frac{X_t^i}{1+\phi(p_t)D_{t+1}}.$$

By symmetry, $\lim_{N\to\infty} Eu_t^{(N)} = Eu_t^i =: \Lambda_t(p_0, \ldots, p_{T-1})$.

Define the operator Λ (an example to follow)

$$\begin{array}{ccc} (p_0, \ldots, p_{T-1}) & \longmapsto & \Lambda(p_0, \ldots, p_{T-1}) := (\Lambda_0, \ldots, \Lambda_{T-1})(p_0, \ldots, p_{T-1}) \\ \text{assumed mean field} & \longmapsto & \text{actual mean field of closed-loop} \end{array}$$

Theorem Λ has a fixed point in a rectangle region.

By this fixed point, the initially assumed mean field is regenerated by the agents taking individual optimal responses.

Step 3: consistency

Example:
We take $T = 2$. Let (p_0, p_1) be given. Then

$$\Lambda_0(p_0, p_1) = \frac{(1+\phi_1)EX_0^i}{1+\phi_1+\phi_1\phi_0}, \quad \Lambda_1(p_0, p_1) = \frac{EG(p_0, W_0^i)EX_0^i}{1+\phi_1+\phi_1\phi_0}.$$

where $\phi_t = \phi(p_t)$ and we recall $\phi(z) = \Phi^{\frac{1}{\gamma-1}}(z)$.

Step 4: ε-Nash

Theorem The set of NCE based strategies $\{\check{u}_t^i, 0 \le t \le T, 1 \le i \le N\}$ is an ε_N-Nash equilibrium, i.e., for any $i \in \{1, \ldots, N\}$,

$$\sup_{u^i} J_i(u^i, \check{u}^{-i}) - \varepsilon_N \le J_i(\check{u}^i, \check{u}^{-i}) \le \sup_{u^i} J_i(u^i, \check{u}^{-i}), \qquad (3.7)$$

where u^i is a centralized strategy and $0 \le \varepsilon_N \to 0$ as $N \to \infty$.

- I did not achieve $\varepsilon_N = O(1/\sqrt{N})$ since no Lipschitz condition is assumed on G (so no usual linear inequality for error estimate and then Gronwall's lemma)
- The performance estimate only uses certain dominated convergence argument

Discussion

- One might think the generalization to infinite horizon is straightforward.
  - Unfortunately this is not the case.
  - We can see rich collective behavior by developing out-of-equilibrium analysis.

The model

Consider the model

$$X_{t+1}^i = G(p, W_t^i)u_t^i, \quad t \geq 0. \tag{4.1}$$

Let the utility functional be

$$J_{[0,\infty)} = \frac{1}{\gamma} E \sum_{i=0}^{\infty} \rho^t (X_t^i - u_t^i)^\gamma. \tag{4.2}$$

Proposition Assume $X_0^i > 0$. Then $\sup J_{[0,\infty)} < \infty$ if and only if (H) $\rho E G^\gamma(p, W) < 1$.

- Suppose "by luck" p is already an equilibrium state (i.e., optimal responses of many agents exactly replicate it).
- Will use this model to diagnose dynamic properties of the mean field system under self-optimizing behavior.

Assume (H). The stationary dynamic programming equation

$$V_i(x) = \max_{u^i} \left[v(x - u^i) + \rho E V_i(G(p, W)u^i) \right].$$

We try a solution of the form

$$V_i(x) = \frac{1}{\gamma} D^{\gamma-1} x^\gamma, \quad D > 0.$$

The optimal control is

$$u_t^i = \frac{X_t^i}{1 + \phi(p)D}, \tag{4.3}$$

where $D > 0$ satisfies

$$D = \frac{\phi(p)D}{1 + \phi(p)D}, \tag{4.4}$$

$\phi(p) = [\rho E G^\gamma(p, W)]^{\frac{1}{\gamma-1}} > 1$ due to (H).

The NCE approach gives

$$\begin{cases} 1+\phi(p)D = \phi(p), \\ EG(p,W) = 1+\phi(p)D \qquad (\text{so} = \phi(p)) \\ EX_0^i = p(1+\phi(p)D) \end{cases} \tag{4.5}$$

where $\phi(p) = [\rho EG^\gamma(p,W)]^{\frac{1}{\gamma-1}}$.

- The first equation is due to optimal reponse/control (recall $D = \frac{\phi(p)D}{1+\phi(p)D}$)
- The second and third equations are the consistency condition on p
  - The second equation means the mean of the aggregate investment is preserved
  - The third equation means the initial state mean is "right"

question

We repeat

$$\begin{cases} 1+\phi(p)D = \phi(p), \\ EG(p,W) = 1+\phi(p)D \qquad (\text{so} = \phi(p)) \\ EX_0^i = p(1+\phi(p)D) \end{cases} \tag{4.6}$$

Suppose the above NCE based u_i is applied by an infinite population which has its initial mean EX_0^i just slightly different from the "right one" as above.

Can the "right mean" be restored when time goes to infinity? Or at least, this system can tolerate the small error?

Definition The pair (p, D) is called a relaxed stationary mean field (RSMF) solution of (4.1)-(4.2) if

$$\begin{cases} 1 + \phi(p)D = \phi(p) \\ EG(p, W) = 1 + \phi(p)D \quad (\mathrm{so} = \phi(p)) \end{cases} \tag{4.7}$$

holds and satisfies $D > 0$, $p \geq 0$ and $\rho EG^\gamma(p, W) < 1$.

Recall $\phi(p) = [\rho EG^\gamma(p, W)]^{\frac{1}{\gamma-1}}$

Here the pair (p, D) is so called since the restriction on the initial mean is removed.

We give an existence result for a class of models with multiplicative noise, i.e.,

$$G(p, W) = g(p)W \tag{4.8}$$

for some function $g > 0$ on $[0, \infty)$, $W \geq 0$ and $EW = 1$.

Theorem Assume (i) $0 < \rho < 1$, (ii) g is continuous and strictly decreases on $[0, \infty)$, (iii) $\rho E(W^\gamma)g(0) \geq 1$, and $\rho E(W^\gamma)g(\infty) < 1$ where $g(\infty) = \lim_{p\to\infty} g(p)$.

Then there exists a unique RSMF solution.

The mean field dynamics with RMFS

The closed-loop state equation is given by

$$X_{t+1}^i = \frac{G(u_t^{(N)}, W_t^i) X_t^i}{\phi(\hat{p})}, \quad 1 \le i \le N.$$

We obtain the relation

$$u_{t+1}^i = \frac{G(u_t^{(N)}, W_t^i)}{\phi(\hat{p})} u_t^i, \quad t \ge 0. \tag{4.9}$$

The mean of $\hat{u}_t^i$ should coincide with p_t, and hence we derive the mean field dynamics

$$p_{t+1} = \frac{p_t EG(p_t, W)}{\phi(\hat{p})}, \quad t \ge 0, \tag{4.10}$$

where $p_0 = Eu_0^i$.

Stable equilibrium

For cases 1)-3), we take $g(p) = \frac{1}{\rho E(W^\gamma)} \cdot \frac{C}{1+(C-1)p^3}$.

Case 1). $C = 1.2$, so $p_{t+1} = \frac{1.2 p_t}{1+0.2p_t^3} =: Q_C(p_t)$, $Q'_C(\hat{p}) = 0.5$ (slope).

Case 2). If $C = 2$, so $p_{t+1} = \frac{2p_t}{1+p_t^3}$, $Q'_C(\hat{p}) = -0.5$.

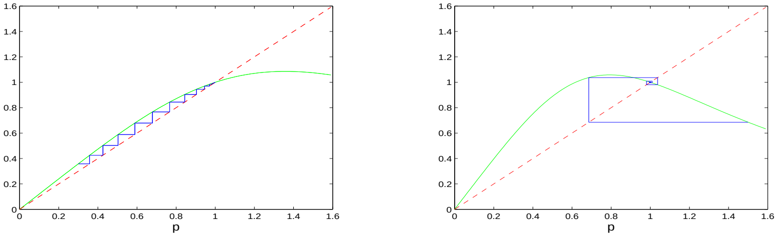

Figure : The iteration of $Q_C(p)$ and stable equilibria. Left: $C = 1.2$; if $0 < p_0 < 1$, the iteration converges to 1 monotonically. Right: $C = 2$.

Limit cycle

Case 3). Take $C = 4$ (i.e., now more sensitive to the MF) to obtain

$$p_{t+1} = \frac{4p_t}{1+3p_t^3} =: Q_C(p_t), \quad Q'_C(\hat{p}) = -1.25,$$

which has a limit cycle determined by the two points on the graph.

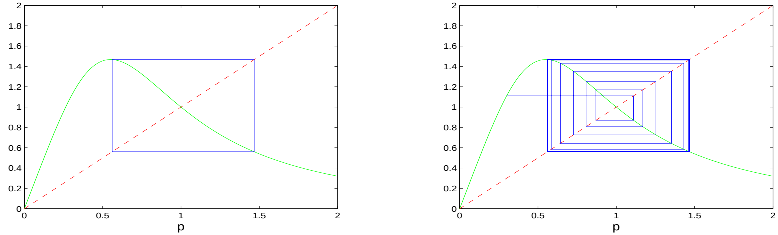

Figure : The iteration of $Q_C(p)$, $C = 4$. Left: the limit cycle. Right: the convergence to the limit cycle with an initial value $p_0 = 0.3$.

Chaotic mean field dynamics

Define $g(p) = \frac{h(p)}{\rho EW^{\gamma}}$ for the production function (see (4.8)), where $EW = 1$. (See graph of h below)
The sharp decrease near $p = 1$ models a "pollution effect": when the aggregate investment level approaches a certain reference level (here equal to 1), the production becomes very inefficient.

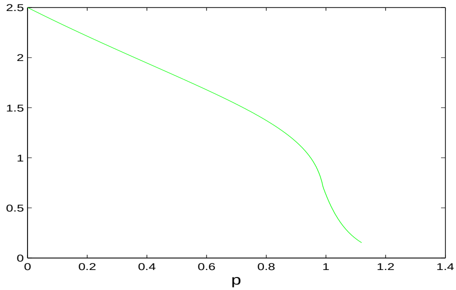

Chaotic mean field dynamics

$$p_{t+1} = p_t h(p_t) =: Q(p_t).$$

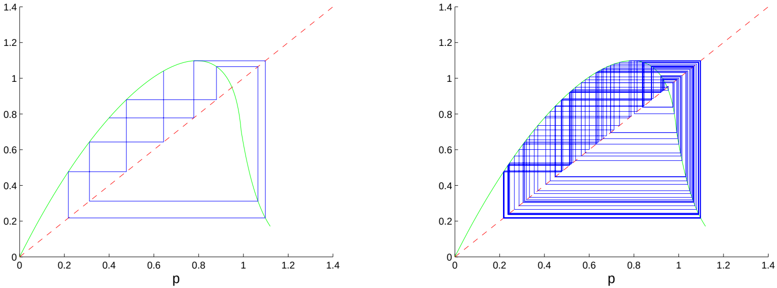

Figure : Initial condition $p_0 = 0.4$. Left: 8 iterates; right: 100 iterates.

See (Li and Yorke, 1975) for definition of chaos. Also see (J. Benhabib ed., Cycles and Chaos in Economic Equil., Princeton Univ. Press, 1992)

Discussion

- The RSMF solution can be used as a classifier of the model.
- When the RSMF solution detects a stable equilibrium, one may propose an asymptotically constant mean field and do consistent MF approximations even if the initial state mean is not the "right one" (Huang, CDC'13).

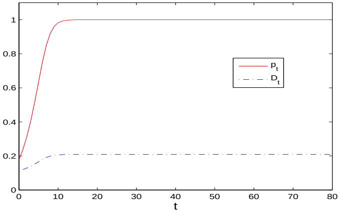

Figure : Infinite horizon mean field approx. with transient behavior.

Discussion

- When oscillatory behavior is diagnosed, is there any hope to develop consistent MF approximations from a certain class of functions? No answer known.
- This weird situation is similar to the El Farol Bar Problem (initially examined by W.B. Arthur, Amer. Econ. Rev., 1994)

Continuous time (CT) modeling

- A mean field version of the AK model

$$dk_t^i = (A(k_t^{(N)})k_t^i - \mu k_t^i - c_t^i)dt - \sigma k_t^i dw_t^i$$

where $-(\mu k_t^i dt + \sigma k_t^i dw_t^i)$ models stochastic depreciation.

- c_t^i is treated as the control variable

Note:

- The above mean field variant is not a strict CT counterpart of the previous discrete time model

The utility functional

$$J_i(0, k_0^i, c(\cdot)) = E\left[\int_0^T e^{-\rho t} U(c_t^i) dt + e^{-\rho T} S(k_T^i)\right],$$

where we take the HARA utility

$$U(z) = \frac{z^\gamma}{\gamma}, \quad 0 < \gamma < 1. \tag{4.11}$$

for $z \geq 0$. We further take

$$S(x) = \frac{\eta x^\gamma}{\gamma}$$

for a small $\eta > 0$. We need this small η to regularize the control near T (well known trick in portfolio selection literature).

The limiting problem

Approximate $A(k_t^{(N)})$ by $A_t := A(\bar{k}_t)$ (becoming a deterministic function of t).

The dynamics of agent i are now approximated by

$$dk_t^i = (A_t k_t^i - \mu k_t^i - c_t^i)dt - \sigma k_t^i dw_t^i, \tag{4.12}$$

The cost is rewritten as

$$\bar{J}_i(0, k_0^i, c^i) = E\left[\int_0^T e^{-\rho t} U(c_t^i)dt + e^{-\rho T} S(k_T^i)\right],$$

which is affected only by c_t^i once A_t, $t \in [0, T]$ is fixed.

Optimal control of the limiting problem

The optimal control law has the explicit form

$$\hat{c}_t^i = \frac{k_t^i}{\eta^{\frac{1}{1-\gamma}} e^{-\int_t^T \lambda_s ds} + \int_t^T e^{-\int_t^\tau \lambda_s ds} d\tau}.$$

where

$$\lambda_t := \lambda(\bar{k}_t) = \frac{\rho + \gamma\mu - \gamma A(\bar{k}_t)}{1-\gamma} + \frac{\sigma^2\gamma}{2}.$$

Replicating $\bar{k}_t$

By averaging

$$dk_t^i = (A(\bar{k}_t)k_t^i - \mu k_t^i - \hat{c}_t^i)dt - \sigma k_t^i dw_t^i,$$

we obtain

$$d\bar{k}_t^{\rm new} = (A(\bar{k}_t) - \mu - \frac{1}{\eta^{\frac{1}{1-\gamma}} e^{-\int_t^T \lambda(\bar{k}_s)ds} + \int_t^T e^{-\int_t^\tau \lambda(\bar{k}_s)ds} d\tau})\bar{k}_t^{\rm new} dt$$

For a given $\bar{k}_0$, we obtain

$$\bar{k}_\bullet^{\rm new} = \Gamma(\bar{k}_\bullet), \qquad t \in [0, T]$$

The consistent mean field approximation reduces to find a fixed point of Γ.

Concluding remarks

- For the continuous time case,
  - No chaos in this case
  - Limit cycles are still possible
  - A close counterpart of the discrete time model is of interest (Introduce continuous time multiplicative noise?)
- Mean field effect depending on distribution instead of just mean (empirical first order moment)