Statistical physics of agent-based systems: Learning dynamics and complex co-operative behaviour in Minority Games

Tobias Galla

The Abdus Salam International Centre for Theoretical Physics
and CNR/INFM SISSA Unit
Trieste, Italy

Acknowledgements

- Ginestra Bianconi (ICTP), Damien Challet (Oxford), Ton Coolen (London), Andrea De Martino (Rome), Matteo Marsili (ICTP), David Sherrington (Oxford), Yi-Cheng Zhang (Fribourg)
- support by European Community’s Human Potential Programme, Research Training Network STIPCO

  also EVERGROW and COMPLEX MARKETS

Advertisement

European Conference on Complex Systems
Saïd Business School, Oxford UK, 25-29 September 2006

Satellite workshop on
’Complex Adaptive Systems and Interacting Agents’

organised by
Andrea De Martino, Enzo Marinari, David Sherrington and myself

`http://chimera.roma1.infn.it/CASIA/`

Minority Game

... originally a simple model for inductive decision making of agents (El-Farol bar problem)

Interest by

- economists
  - simple model of a market, stylised facts ...
- physicists
  - phase transitions, ergodicity breaking, spin glass problem, off-equilibrium dynamics
- mathematicians
  - exact solutions

Stock market

- traders
- they observe a price time-series (and other information)
- based on this they buy/sell
- price is formed based on their actions
- they learn and adapt (some better than others maybe)

Stock market

- traders

  particles, spins, microscopic degrees of freedom

- they observe a price time-series (and other information)

  externally and/or internally generated information, history, can be non-Markovian

- based on this they buy/sell

  decision making (noise ...)

- price is formed based on their actions

  global interaction, macroscopic observable, mean-field

- they learn and adapt (some better than others maybe)

  dynamics, update rules, equations of motion

The model: Minority Game

[Challet, Zhang 1997]

- N traders $i = 1, \ldots, N$
- given signal $\mu(t) \in \{1, \ldots, P\}$ at each time-step

  here: random external information
- then every player has to make a binary trading decision $b_i(t) \in \{-1, 1\}$
- all players in minority are successful, players in majority unsuccessful
- if $A(t)$ is the total bid $A(t) = \sum_i b_i(t)$, then payoff for i is

$$-b_i(t)A(t)$$

The model: Minority Game

How do players make trading decisions ?

- everybody has S trading strategies $\vec{a}_{i,s}$, $s = 1, \ldots, S$ mapping μ onto $a_i^\mu \in \{-1, 1\}$ (buy or sell)

Strategy is a table mapping μ onto binary decision

μ	1	2	3	4	...	P
a_i^μ	-1	1	1	-1	...	-1

Given history μ a strategy table tells me to play a_i^μ.

The model: Minority Game

Consider case $S = 2$ strategies per player in the following

strategy $s = +1$

μ	1	2	3	4	...	P
$a^{\mu}_{i,s=+1}$	-1	1	1	-1	...	-1

strategy $s = -1$

μ	1	2	3	4	...	P
$a^{\mu}_{i,s=-1}$	1	1	-1	-1	...	1

Then what this player has to decide at time t is which of the two tables to use.

Assign scores to each strategy to measure their success.

The model: Minority Game

- aim: to be in the minority
- which strategy to use ? The one which has performed best so far !
- to assess performance keep a score for each strategy:

$$u_{i,s}(t+1) = u_{i,s}(t) + \underbrace{(-a_{i,s}^{\mu(t)} A(t))}_{\text{minority game payoff}}$$

- strategies generated randomly before start of the game

MG dynamics

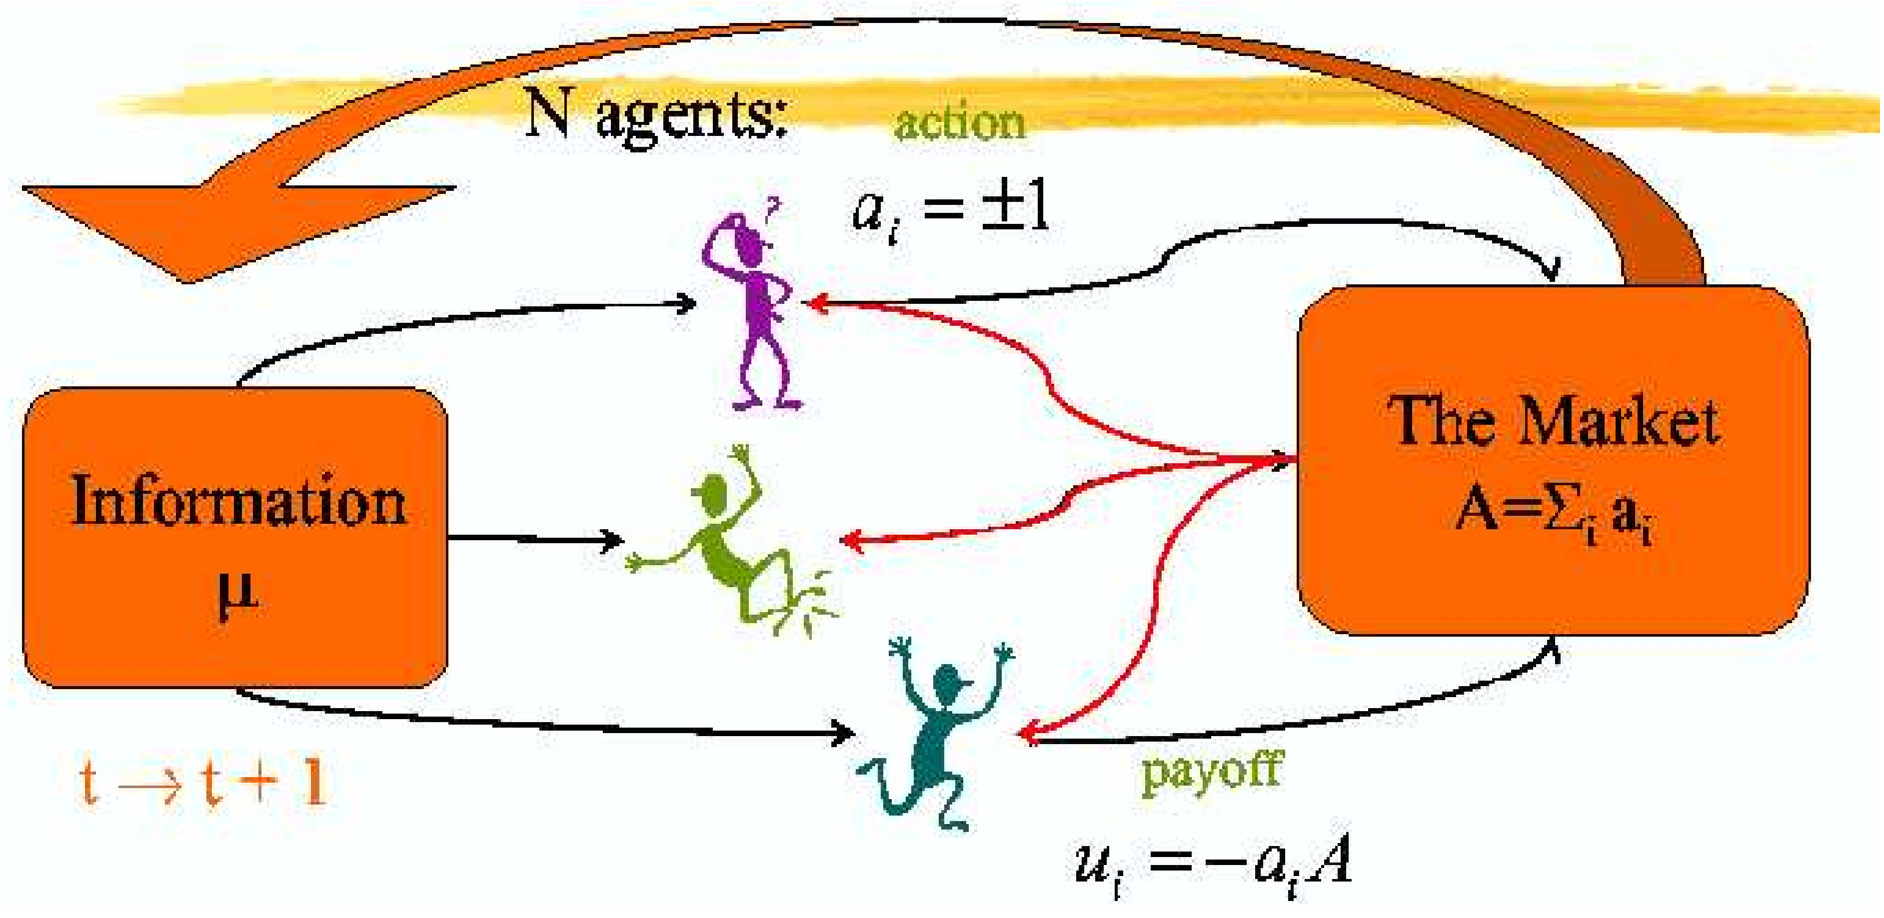

[Marsili’s slide]

MG for physicists

The phenomenology of the basic MG

What are the interesting observables ?

And what are the model parameters ?

The phenomenology of the basic MG

Model parameters ... just one.

$$\alpha = \frac{\text{number of values information can take}}{\text{number of agents}} = \frac{P}{N}$$

- i.e. α high: large information space and/or small market
- low α means the opposite: large market and/or small information space

Observables

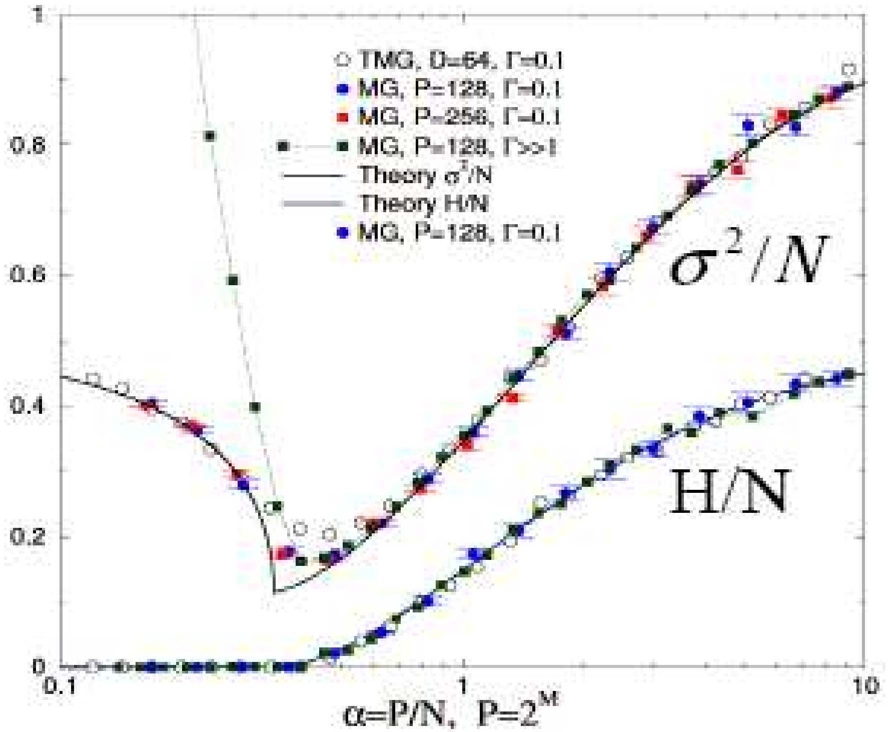

[Challet, Marsili, Zecchina]

Predictability

$$H = \frac{1}{P} \sum_{\mu=1}^{P} \langle A|\mu \rangle^2$$

$H > 0 \Rightarrow \langle A|\mu \rangle \neq 0$ statistically predictabl

$H = 0 \Rightarrow \langle A|\mu \rangle = 0$ predictability zero

global performance/volatility

$\sigma^2 = \langle A^2 \rangle = -$total gain

Observables

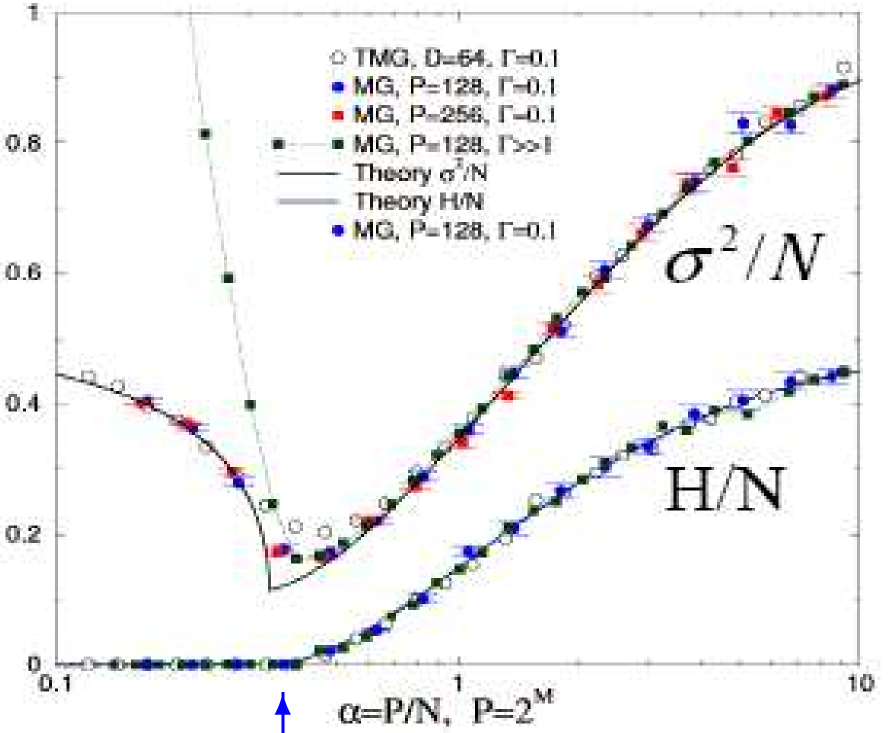

Predictability

$$H = \frac{1}{P}\sum_{\mu=1}^{P}\langle A|\mu\rangle^2$$

$H > 0 \Rightarrow \langle A|\mu\rangle \neq 0$ statistically predictabl

$H = 0 \Rightarrow \langle A|\mu\rangle = 0$ predictability zero

global performance/volatility

$\sigma^2 = \langle A^2\rangle = -$total gain

Phase transition between a predictable and an unpredictable phase

Ergodicity breaking

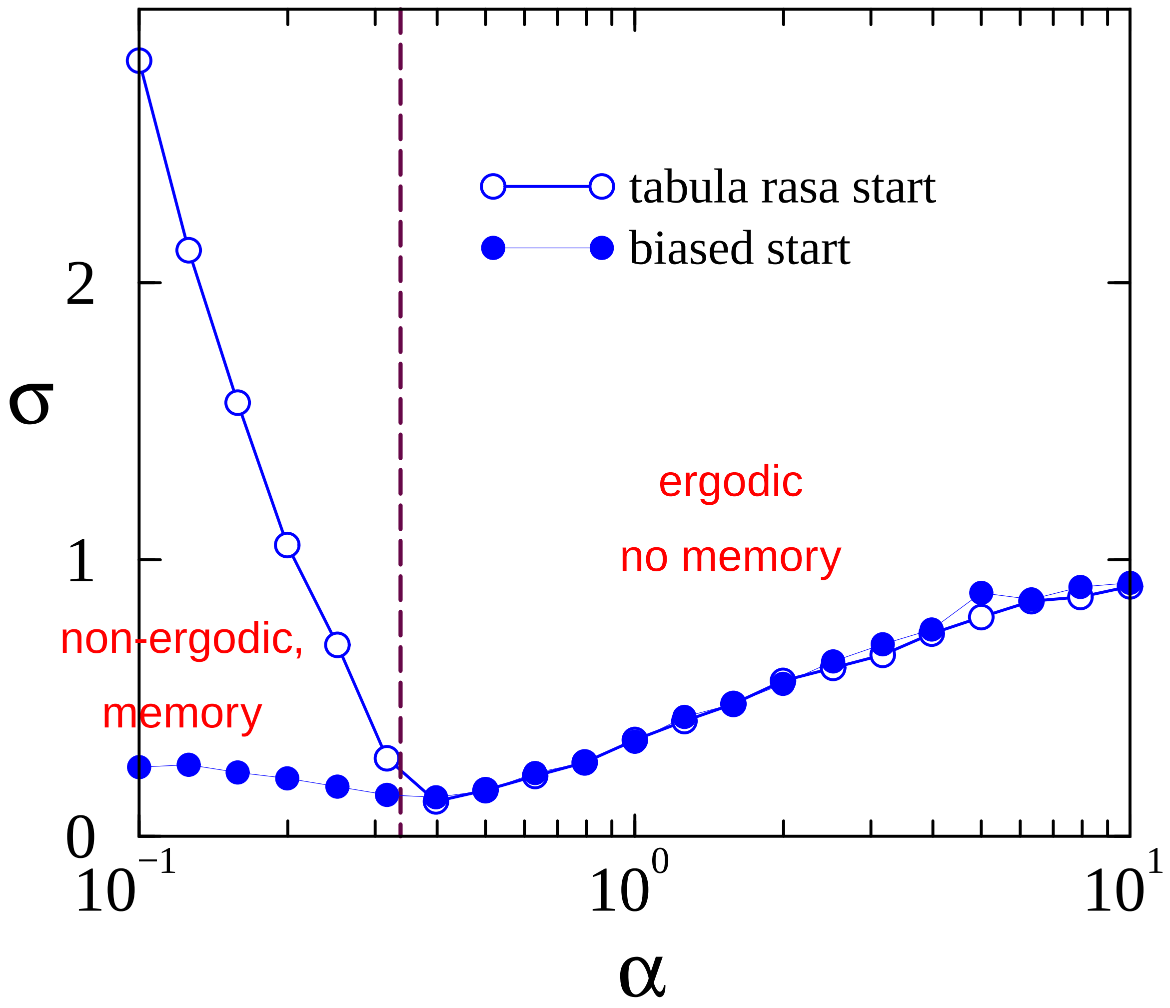

Ergodicity breaking

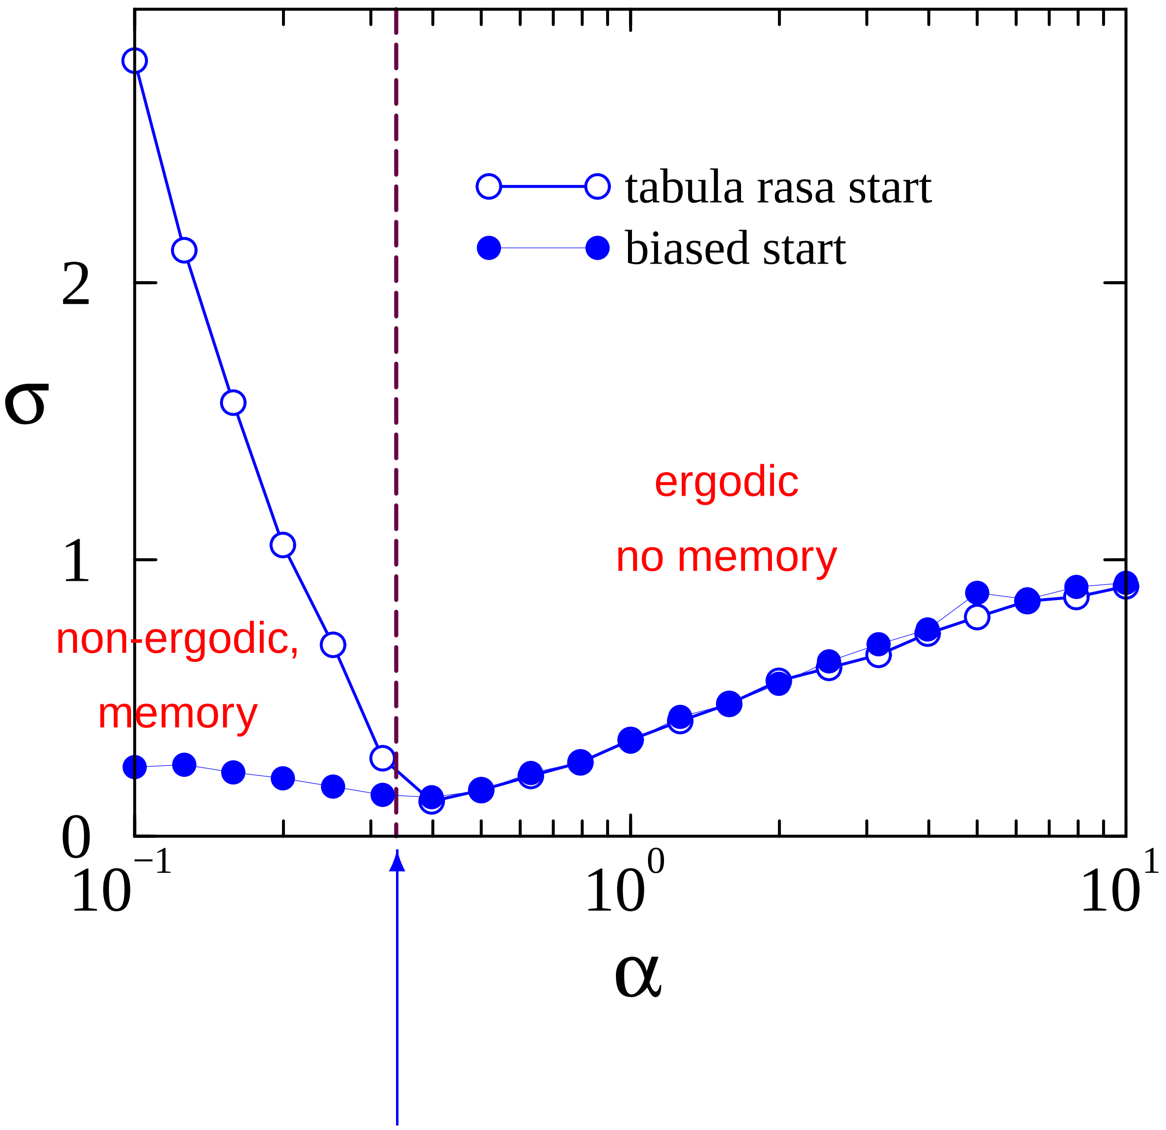

Phase transition between a non-ergodic and an ergodic phase

Ergodicity breaking

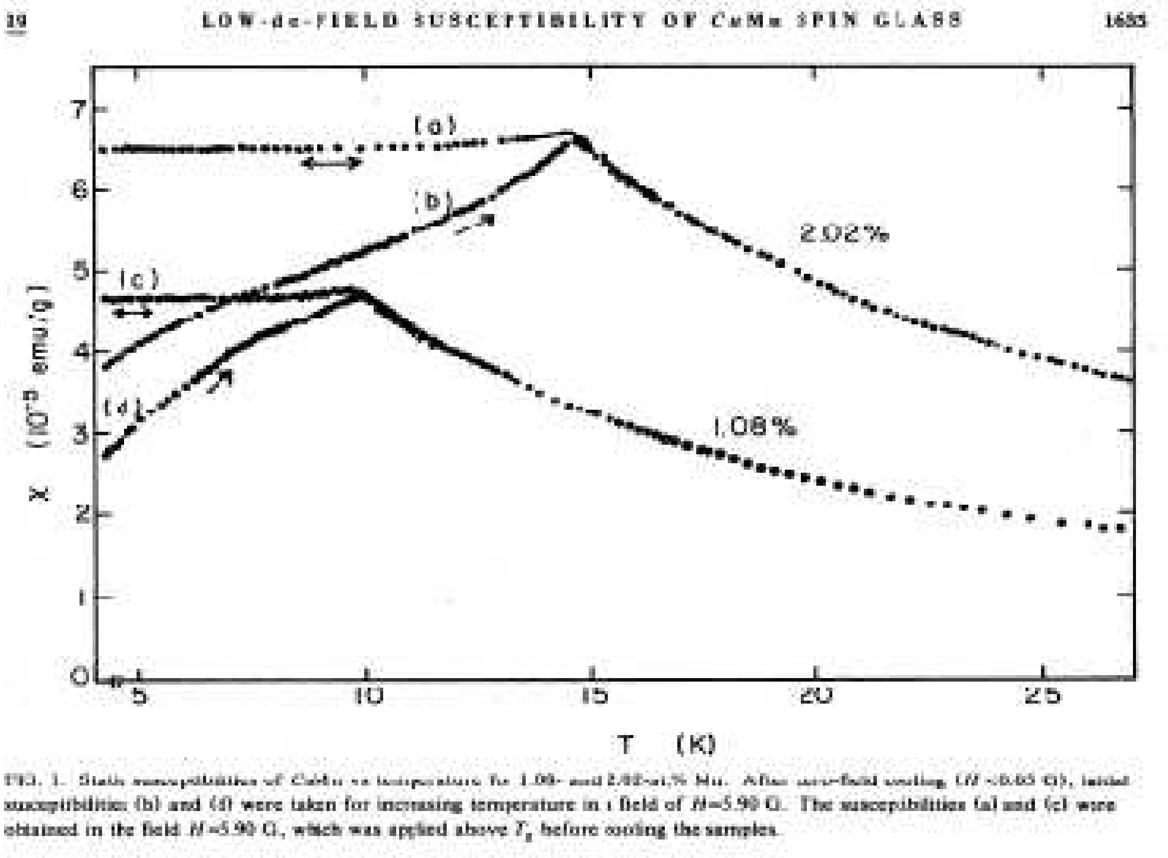

19 LOW-dc-FIELD SUSCEPTIBILITY OF CuMn SPIN GLASS 1633

FIG. 1. Static susceptibilities of CuMn vs temperature for 1.08- and 2.02-at.% Mn. After zero-field cooling ($H<0.05$ G), initial susceptibilities (b) and (d) were taken for increasing temperature in a field of $H=5.90$ G. The susceptibilities (a) and (c) were obtained in the field $H=5.90$ G, which was applied above T_g before cooling the samples.

static susceptibilities of CuMn, field-cooling versus zero-field cooling

MG as spin glass model

MG shares many features with spin-glass models

$$\mathcal{H}_{SK} = \sum_{ij} J_{ij} s_i s_j, \quad \overline{J_{ij}^2} = \frac{1}{N}$$

[Sherrington-Kirkpatrick model, SK 1975]

- frustration (not everybody can win)
- quenched disorder (random strategy assignments)
- mean-field interactions (interaction with ev’body else)

Remember

$S = 2$ strategies per player:

$s_i = +1$, score u_{i+}

μ	1	2	3	4	...	P
$a^{\mu}_{i,s=+1}$	-1	1	1	-1	...	-1

$s_i = -1$, score u_{i-}

μ	1	2	3	4	...	P
$a^{\mu}_{i,s=-1}$	1	1	-1	-1	...	1

Then what this player has to decide at time t is which of the two tables to use.

$$s_i(t) = \text{sgn}[u_{i+}(t) - u_{i-}(t)]$$

Learning dynamics

proposed action total action

$$u_{i,+}(t+1) = u_{i,+}(t) - a_{i,+}^{\mu(t)} A(t)$$

$$u_{i,-}(t+1) = u_{i,-}(t) - a_{i,-}^{\mu(t)} A(t)$$

Evolution of score difference ($q_i = u_{i+} - u_{i-}$):

$$q_i(t+1) = q_i(t) - \left[a_{i,+}^{\mu(t)} - a_{i,-}^{\mu(t)} \right] A(t)$$

Learning dynamics

On-line update for score difference ($q = u_+ - u_-$):

$$q_i(t+1) = q_i(t) - \left[a_{i,+}^{\mu(t)} - a_{i,-}^{\mu(t)}\right] A(t)$$

and

$$A(t) = \sum_j f(\text{sgn}[q_j(t)] | \text{strategies of } j)$$

Batch update for score difference (average over μ):

$$q_i(t+1) = q_i(t) - \sum_j J_{ij} \text{sgn}[q_j(t)] - h_i$$

quenched disorder, spin glass problem

$$J_{ij} = \underbrace{\frac{1}{P} \sum_{\mu=1}^{P} \frac{(a_{i+}^{\mu} - a_{i-}^{\mu})}{2} \frac{(a_{j+}^{\mu} - a_{j-}^{\mu})}{2}}_{\text{Hebbian}}, \quad h_i = \frac{1}{P} \sum_{j=1}^{N} \sum_{\mu=1}^{P} \frac{(a_{i+}^{\mu} - a_{i-}^{\mu})}{2} \frac{(a_{j+}^{\mu} + a_{j-}^{\mu})}{2}$$

Dynamics

$$q_i(t+1) - q_i(t) \quad = \quad -\sum_i J_{ij}\mathsf{sgn}[q_j(t)] - h_i$$

but not

$$q_i(t+1) - q_i(t) \quad = \quad -\sum_i J_{ij} q_j(t) - h_i = -\frac{\partial H[\mathbf{q}]}{\partial q_i}$$

No gradient-descent. No detailed balance. Still pseudo-Hamiltonian:

$$\mathcal{H}(\mathbf{s}) = \frac{1}{2}\sum_{ij} J_{ij} s_i s_j + \sum_i h_i s_i$$

MG as an anti-Hopfield model

$$\mathcal{H}(\mathbf{s}) = \frac{1}{2}\sum_{ij} J_{ij} s_i s_j + \sum_i h_i s_i, \qquad J_{ij} = \frac{1}{\alpha N}\sum_{\mu} \xi_i^\mu \xi_j^\mu$$

Hopfield model has

$$\mathcal{H}(\mathbf{s}) = -\frac{1}{2}\sum_{ij} J_{ij} s_i s_j$$

MG is an ‘unlearning’ game.

MG for mathematicians

Generating functional analysis

[Heimel, Coolen PRE 2001]

$$q_i(t+1) - q_i(t) \quad = \quad -\sum_j J_{ij}\text{sgn}[q_j(t)] - h_i + \underbrace{\vartheta(t)}_{perturbation field}$$

Dynamical partition function

$$\begin{aligned}
Z[\psi] \quad &= \quad \int D\mathbf{q}\,\delta(\text{eq of motion})\ \exp\left(i\sum_{it}\psi_i(t)\text{sgn}[q_i(t)]\right) \\
&= \quad \int D\mathbf{q}D\widehat{\mathbf{q}}\exp\left(\sum_{it}\widehat{q}_i(t)[q_i(t+1)-q_i(t)+\sum_j J_{ij}\text{sgn}[q_j(t)]+h_i-\vartheta(t)]\right) \\
&\quad\ \times\exp\left(i\sum_{it}\psi_i(t)\text{sgn}[q_i(t)]\right)
\end{aligned}$$

Then path integrals, disorder-average, saddle-point equations ...

Generating functional analysis

$$q(t+1) = q(t) + \vartheta(t) - \alpha \sum_{t'} [\mathbb{I} + G]^{-1}_{tt'} \text{sgn}[q(t')] + \sqrt{\alpha}\eta(t)$$

with noise covariance

$$\begin{aligned} < \eta(t)\eta(t') > &= [(\mathbb{I} + G)^{-1} D (\mathbb{I} + G^T)^{-1}]_{tt'} \\ D_{tt'} &= 1 + C_{tt'} \end{aligned}$$

Dynamical order parameters:

$$C_{tt'} = < \text{sgn}[q(t)]\text{sgn}[q(t')] >, \qquad G_{tt'} = \frac{\partial}{\partial \vartheta(t')} < \text{sgn}[q(t)] >$$

[Heimel/Coolen PRE 2001]

[Coolen/Heimel J Phys A 2001]

[Coolen J Phys A 2005]

Basic MG

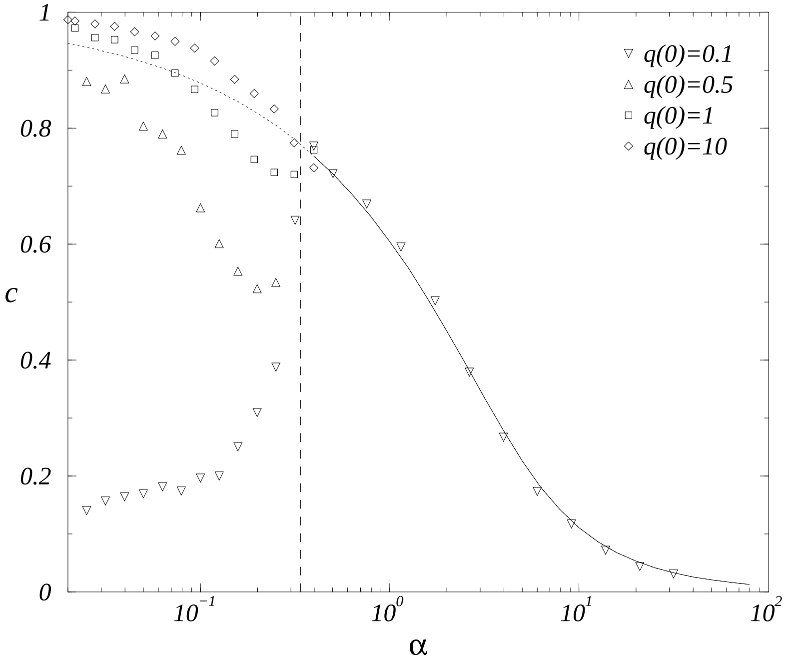

EA parameter

exact result [Heimel/Coolen]

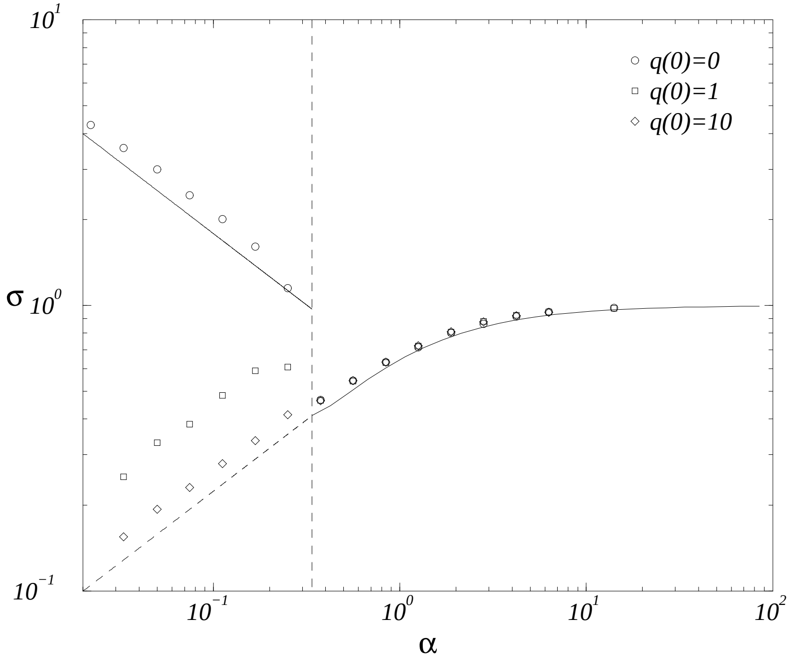

volatility

approximation: drop transients [Heimel/Coolen]

Spherical MG

Replace

$$q_i(t+1) - q_i(t) = -\sum_i J_{ij} \underbrace{\text{sgn}[q_j(t)]}_{\text{Ising}} - h_i$$

by

$$q_i(t+1) - q_i(t) = -\sum_j J_{ij} \underbrace{\phi_j(t)}_{\text{continuous}} - h_i,$$

with

$$\phi_i = \frac{q_i}{\lambda}, \quad \sum_i \phi_i^2 = N$$

[Galla, Coolen, Sherrington J Phys A 2003]

[Galla, Sherrington JSTAT 2005]

Spherical MG

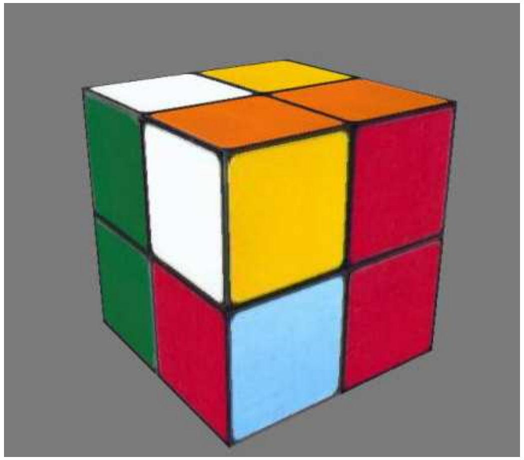

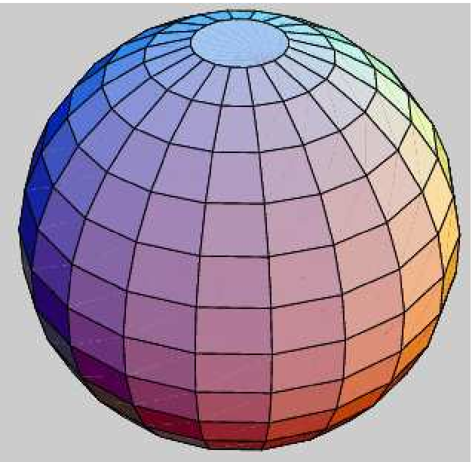

$$Z = \sum_{\{s_i=\pm 1\}} \exp(-\beta H) \qquad Z = \int d\vec{\phi}\, \delta(\vec{\phi}^{\,2} - N) \exp(-\beta H)$$

[Kac, Berlin ‘The Spherical Model of a Ferromagnet’, Phys. Rev. 86, 821-835 (1952)]

Spherical MG

conventional MG:

spherical MG:

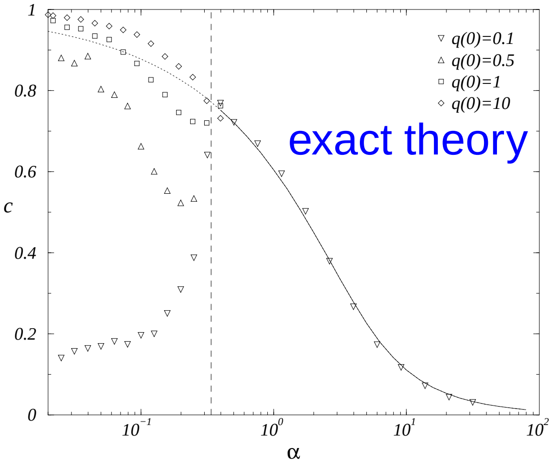

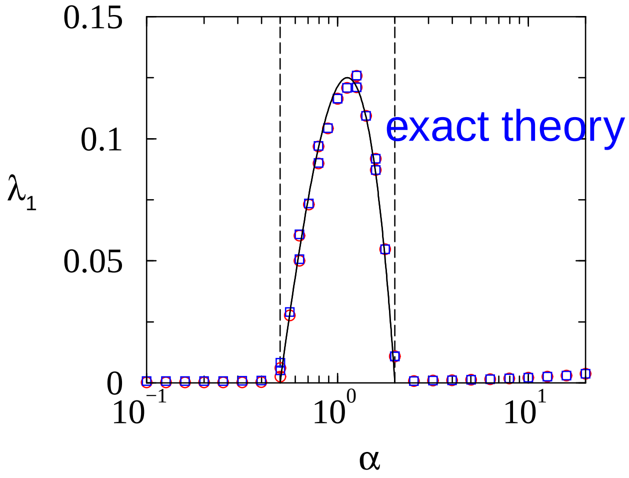

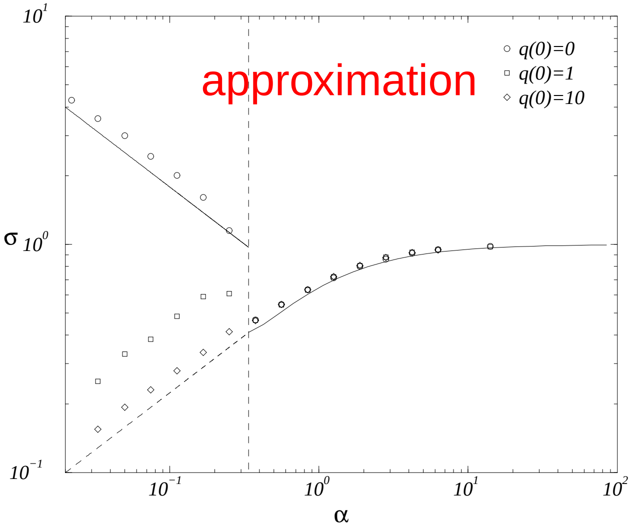

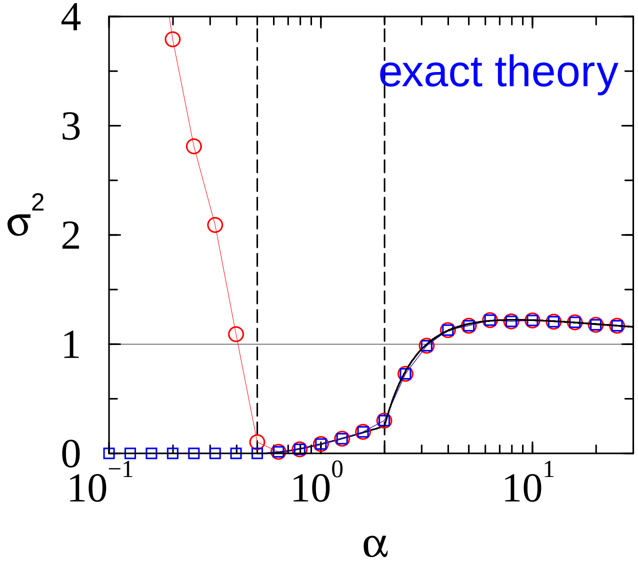

Back to physics

Batch versus on-line learning

on-line learning: strategy switches allowed at every step

$$u_{i,s}(\ell+1) = u_{i,s}(\ell) - \underbrace{a_{i,s}^{\mu(\ell)} A^{\mu(\ell)}(\ell)}_{\text{minority game payoff}}$$

batch learning: strategy switches allowed only after $\mathcal{O}(\alpha N)$ steps

$$u_{i,s}(t+1) = u_{i,s}(t) - \frac{1}{\alpha N}\sum_{\mu=1}^{\alpha N} a_{i,s}^{\mu} A^{\mu}(t)$$

Does it make a difference ?

Timing of adaptation

Not in the standard MG:

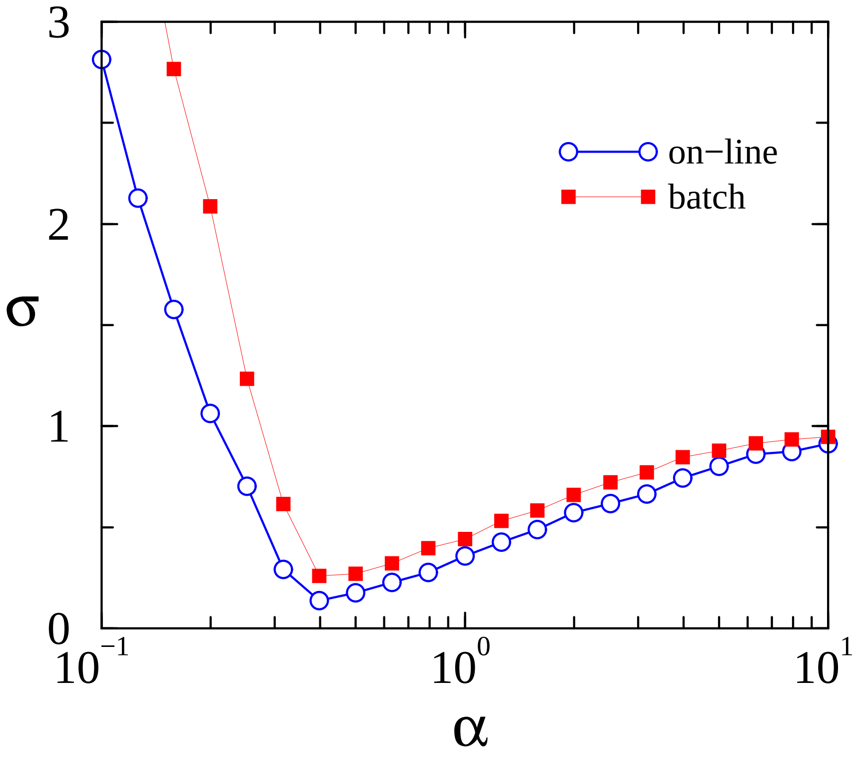

Timing of adaptation

But in an MG with anti-correlated strategy assignments it does:

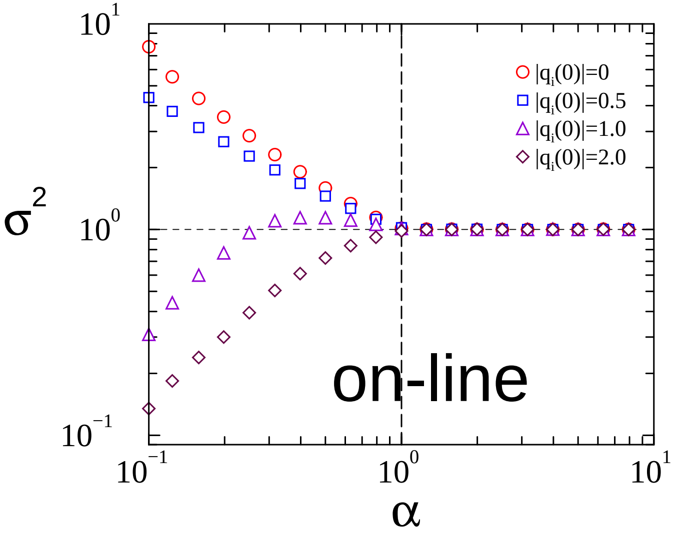

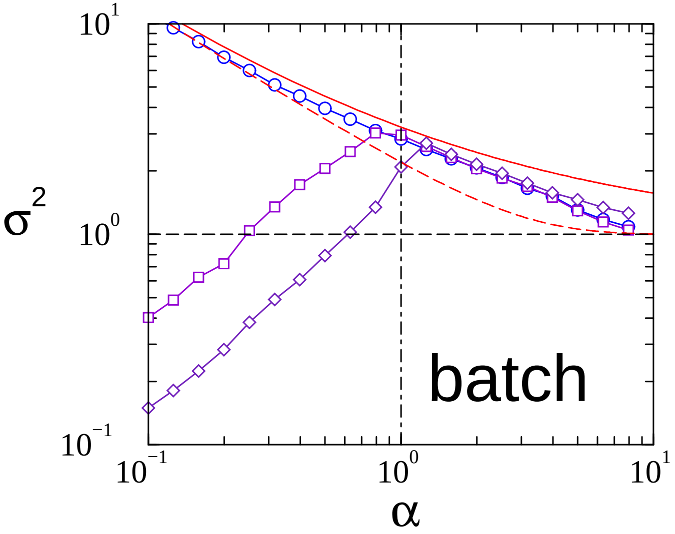

[Sherrington, Galla Physica A 2003]

[Galla, Sherrington EPJB 2005]

Timing of adaptation

Interpolation between on-line and batch:
updates every M time-steps

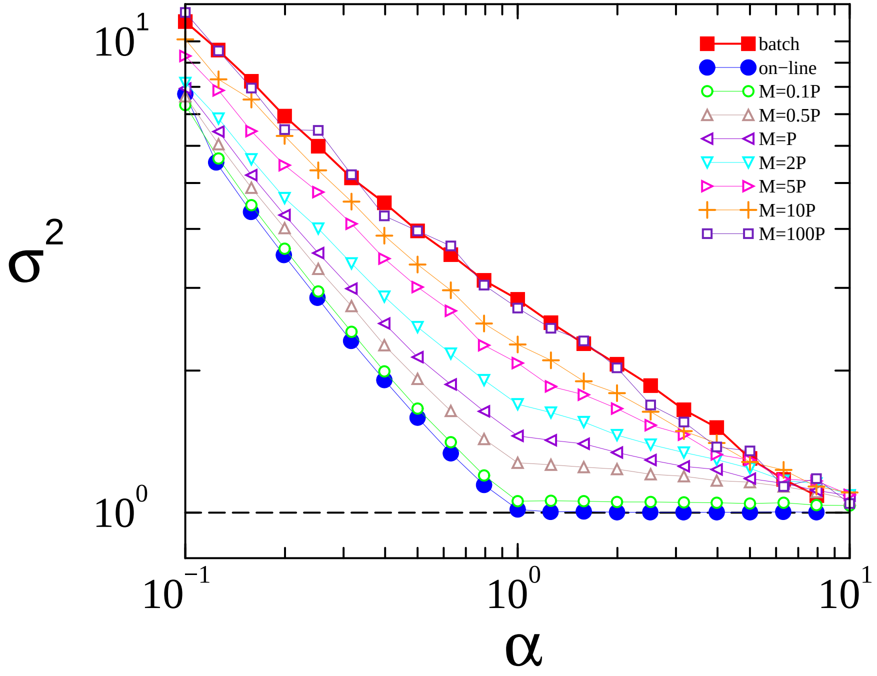

The phase transition

$$\chi = \sum_{\tau} G(\tau) = \begin{cases} \text{finite ?} & \text{-> system ergodic} \\ \text{infinite ?} & \text{-> system non-ergodic} \end{cases}$$

$$\alpha < \alpha_c, \chi = \infty, H = 0$$

non-ergodic, perturbations persists, memory

$$\alpha > \alpha_c, \chi < \infty, H > 0$$

ergodic, perturbations decay, no memory

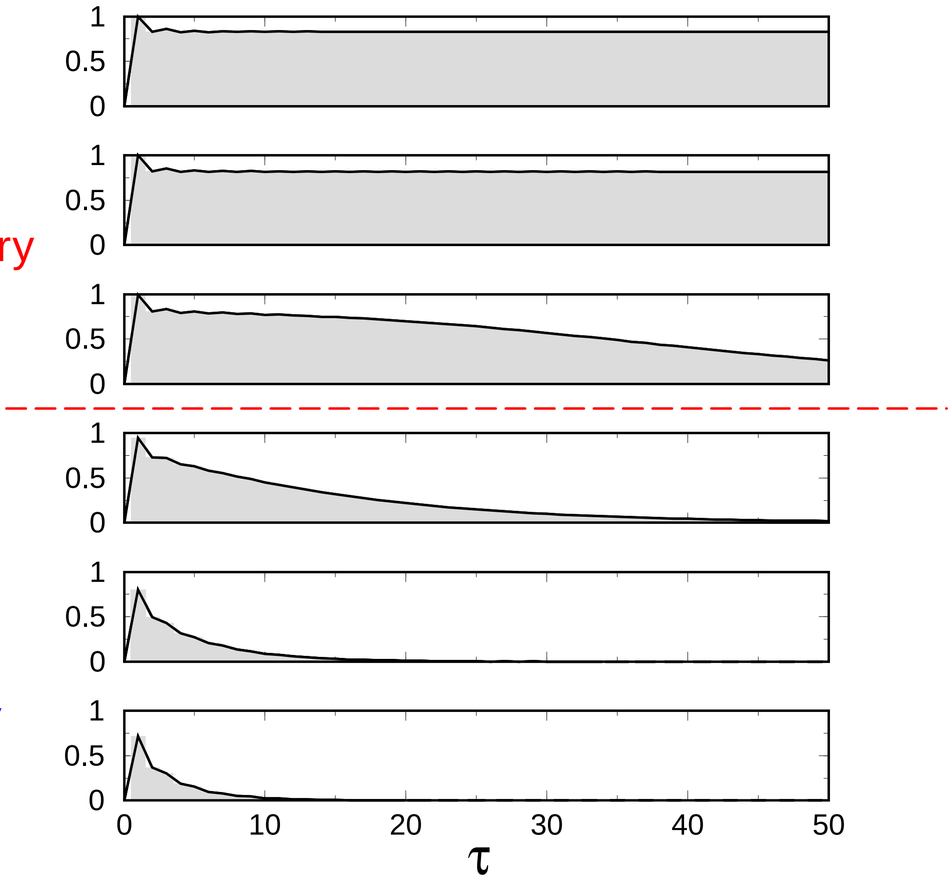

Picture in phase space

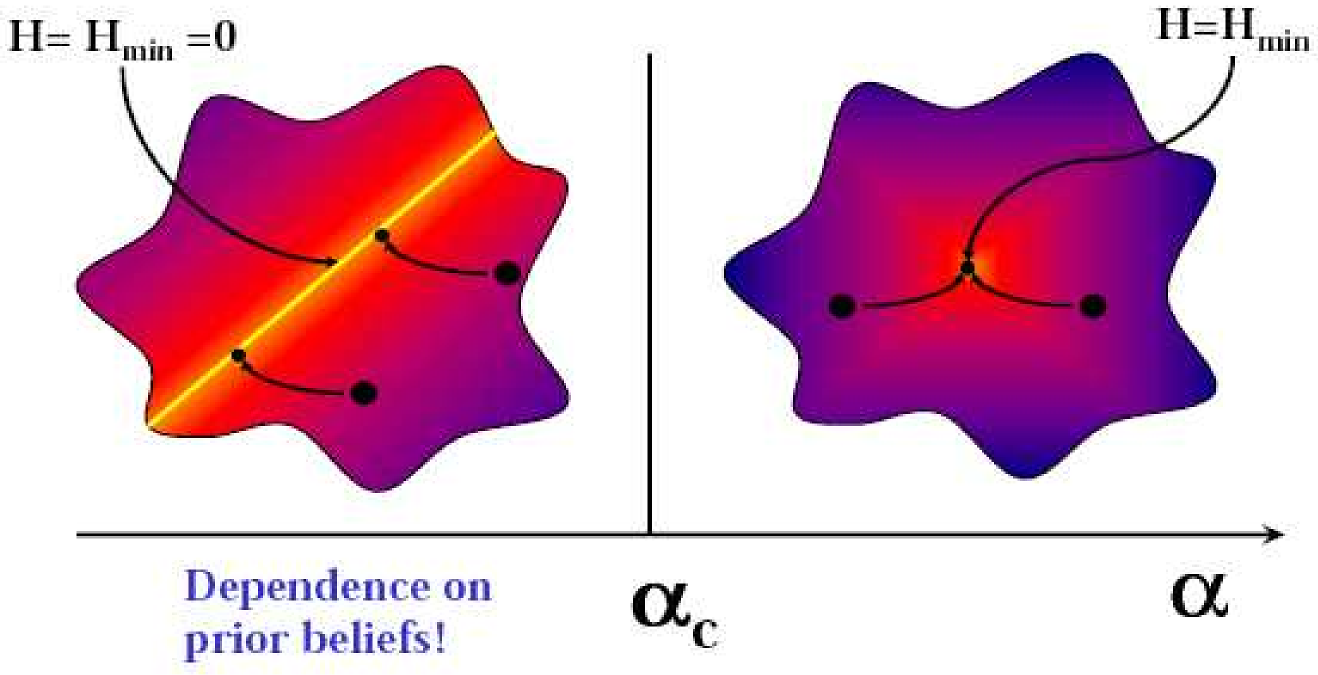

No replica symmetry breaking in standard MG.

[Marsili]

RSB in modified MGs

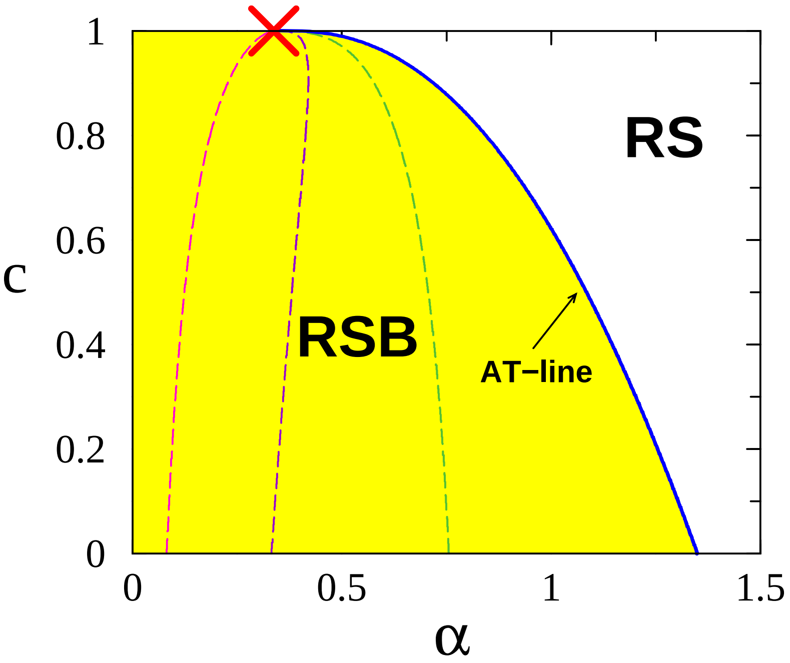

dilute MG

[Galla JSTAT 2005]

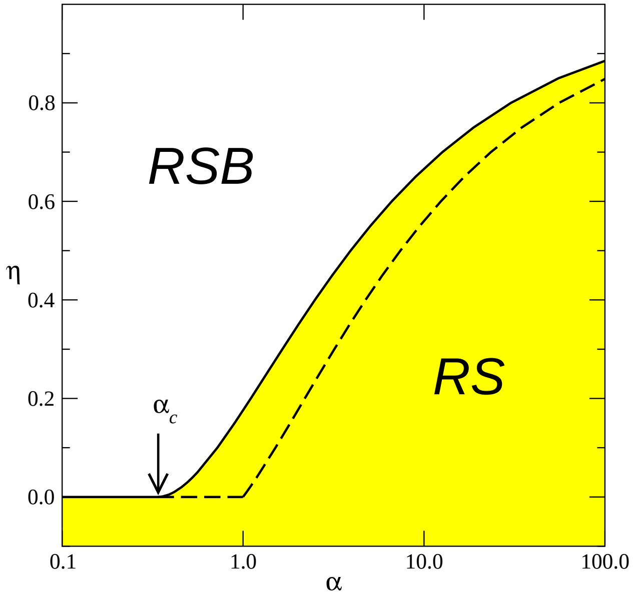

MG with impact correction

[Heimel, De Martino, J Phys A 2001]

[De Martino, Marsili J Phys A 2001]

also in El-Farol with heterogeneous resource level [De Sanctis, Galla, in preparation]

The Physicists view

- This is all very nice ...

The Physicists view

- This is all very nice
- ... but does one see anything like feature of real-market data in this model ?

The Physicists view

- This is all very nice ...
- ... but does one see anything like feature of real-market data in this model ?

Actually ... No

MG for economists

Basic stylised facts

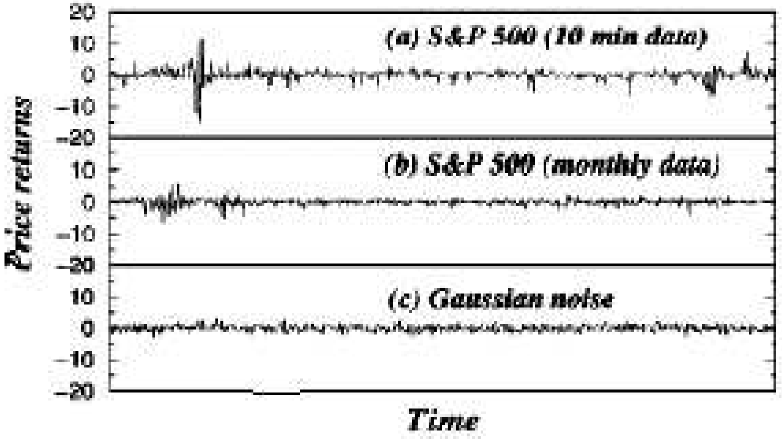

FIG. 2. Sequence of (a) 10-min returns, from database (i), and (b) one-month returns, from database (iii), for the S&P 500, normalized to unit variance. (c) Sequence of i.i.d. Gaussian random variables with unit variance, which was proposed by Bachelier as a model for stock returns [1]. For all three panels, there are 850 events —i.e., in panel (a) 850 min and in panel (b) 850 months. Note that, in contrast to (a) and (b), there are no "extreme" events in (c).

H.E. Stanley et al. / Physica A 269 (1999) 156–169

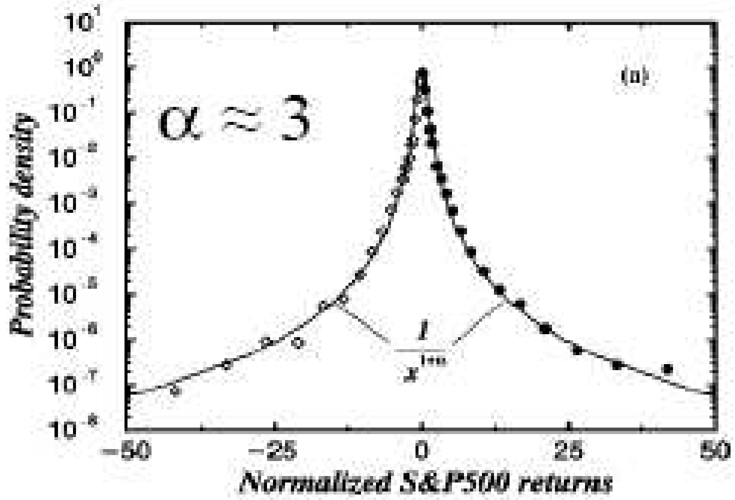

Basic MG

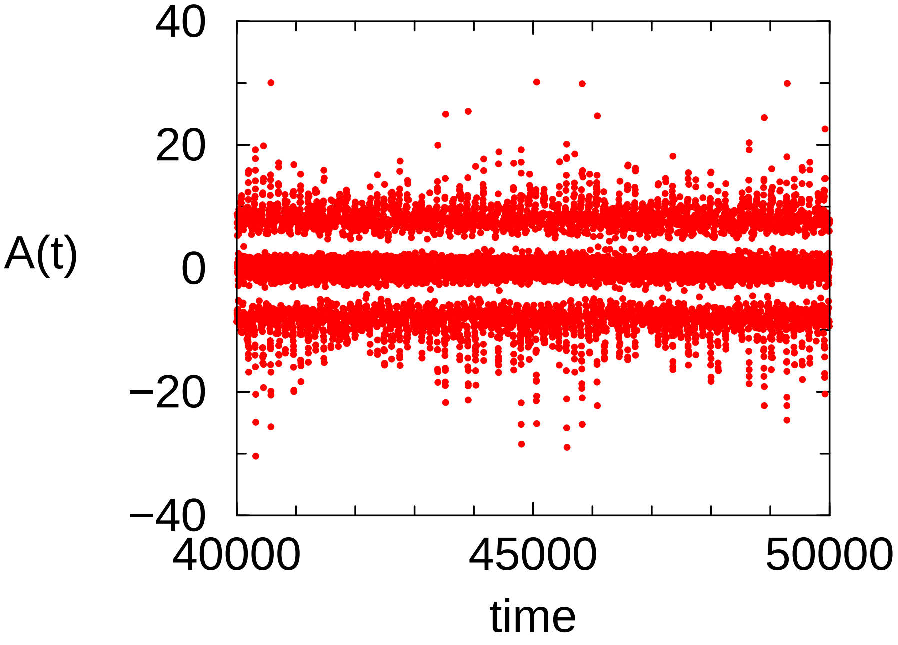
A(t)
40
20
0
−20
−40
40000
45000
50000
time

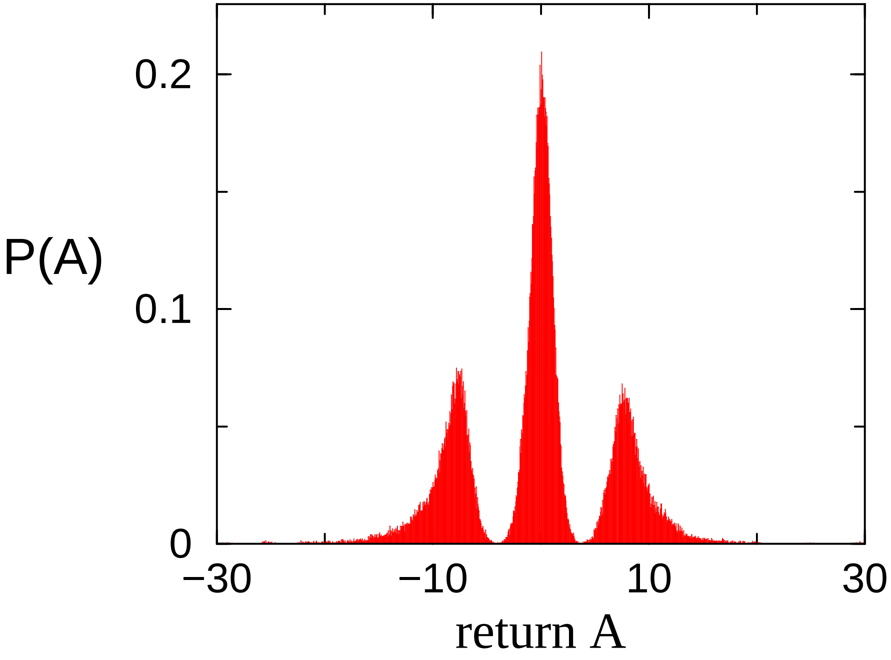
P(A)
0.2
0.1
0
−30
−10
10
30
return A

Way out

What are the minimal additions one has to make to make it more realistic ?

- give the agents the choice not to play -> grand-canonical MGs
- give them dynamically evolving capitals

Both things have similar effects: the trading volume is no longer constant ($= N$ up to now), but can evolve in time.

General idea

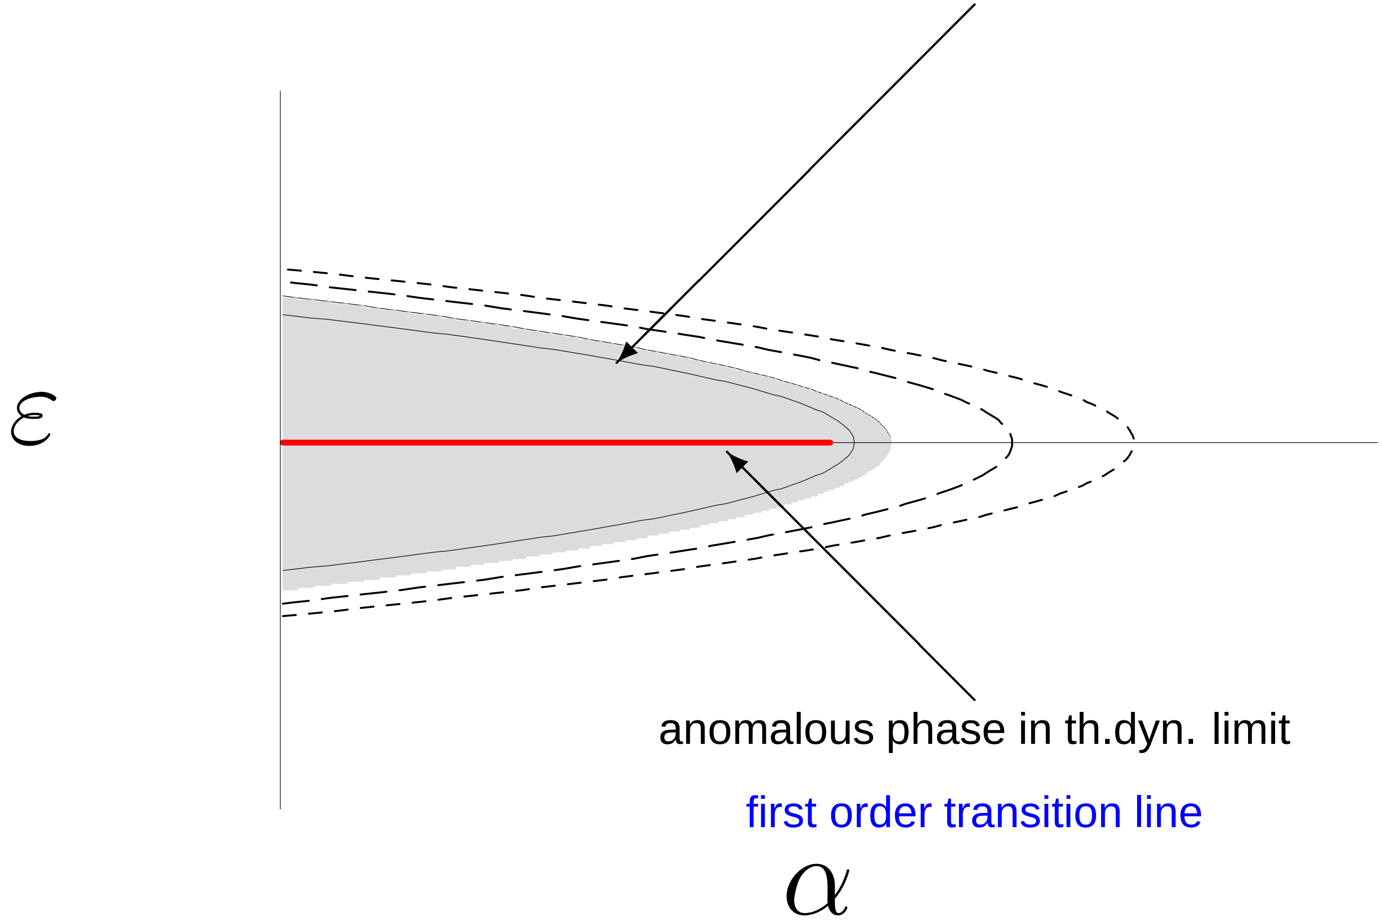

critical region at finite N:

stylised facts+interesting dynamical features

anomalous phase in th.dyn. limit

first order transition line

General idea

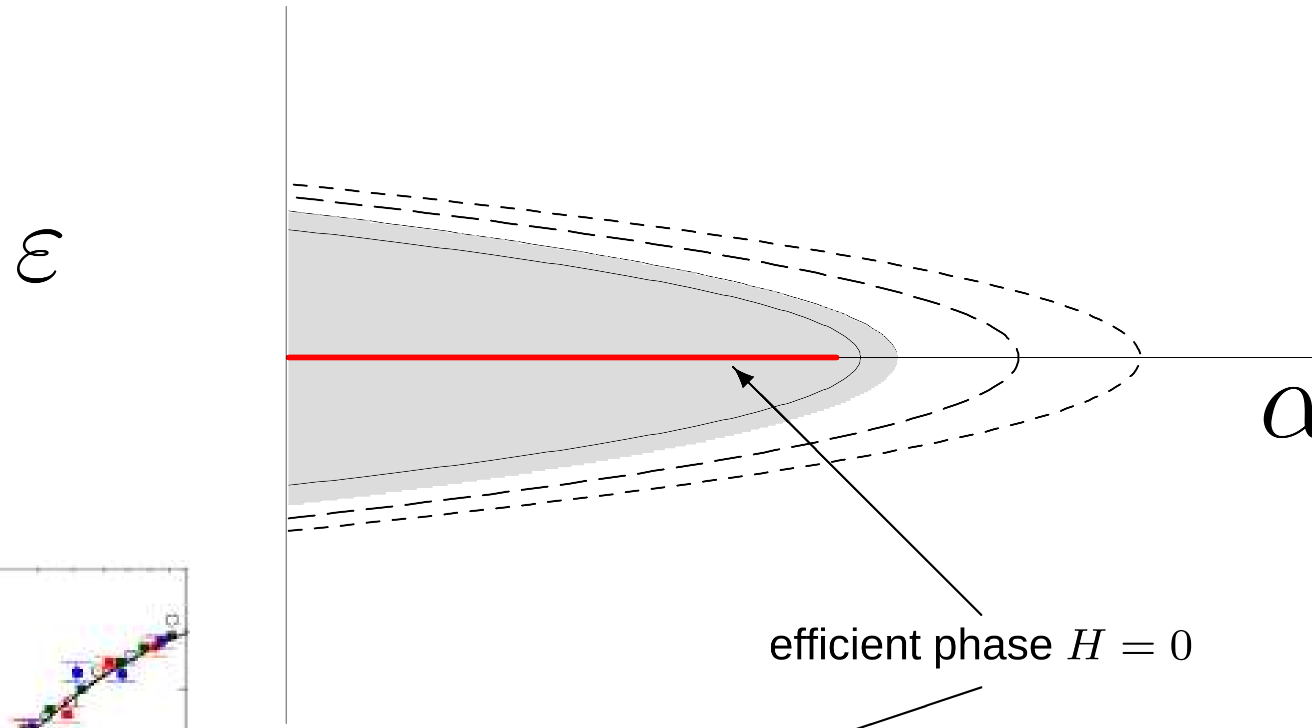

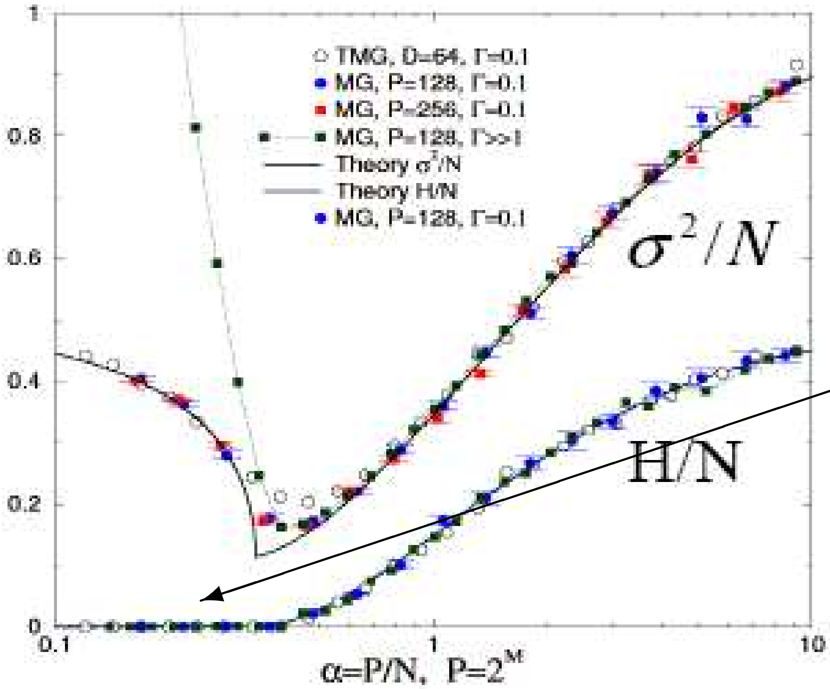

efficient phase $H = 0$

Stylised facts

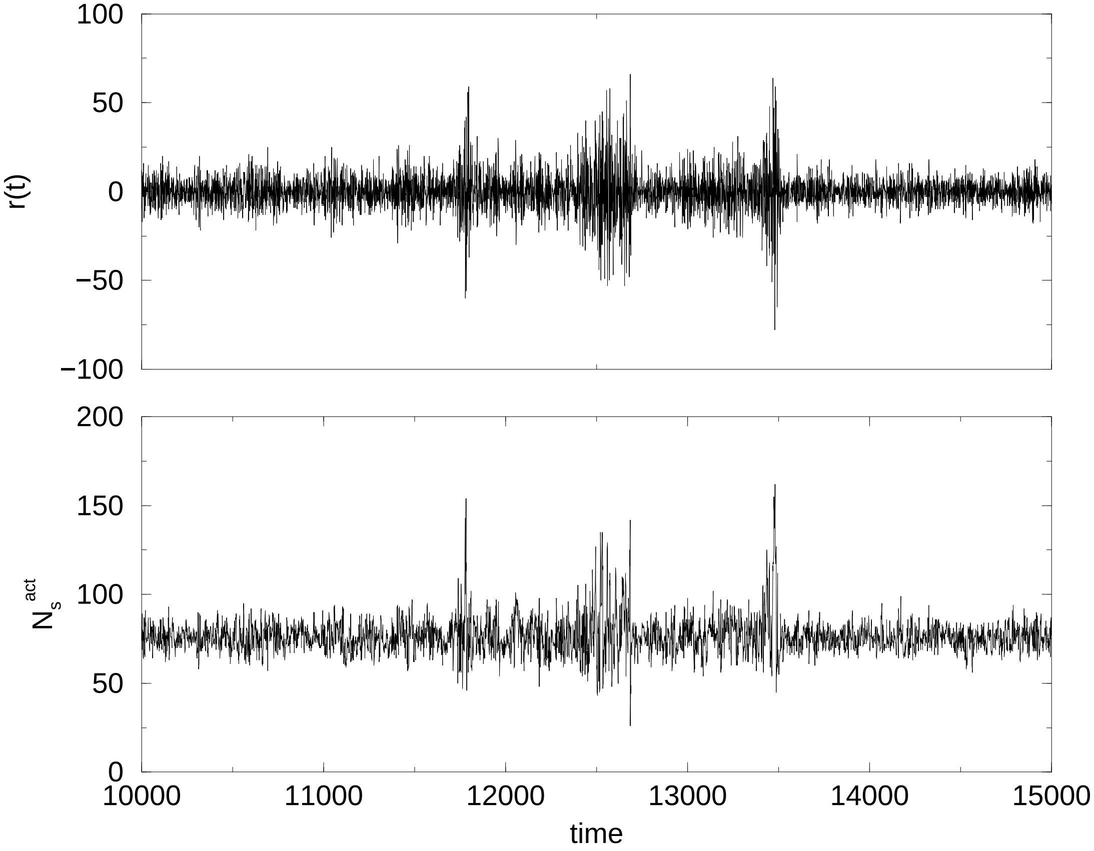

[Challet, Marsili, Zhang 2001]

MG with dynamical capitals

MG with 2 strategies per player and dynamical capitals:

[Challet, Chessa, Marsili, Zhang (2001)]

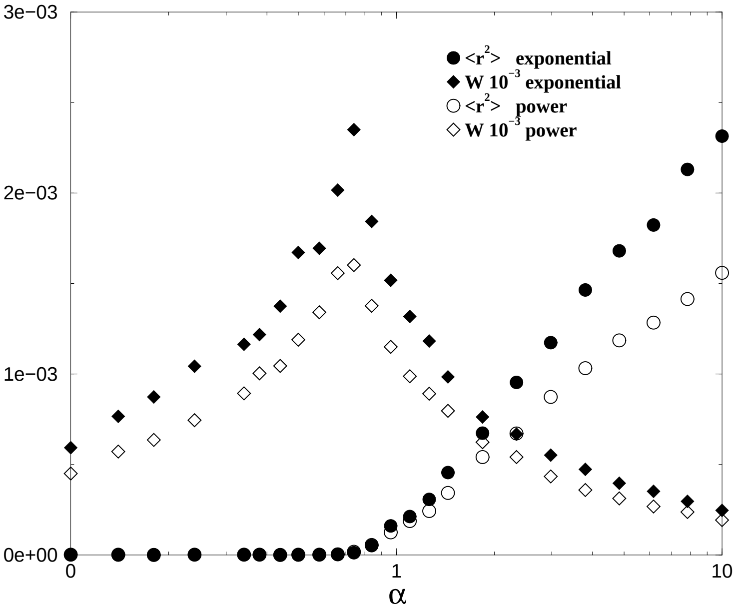

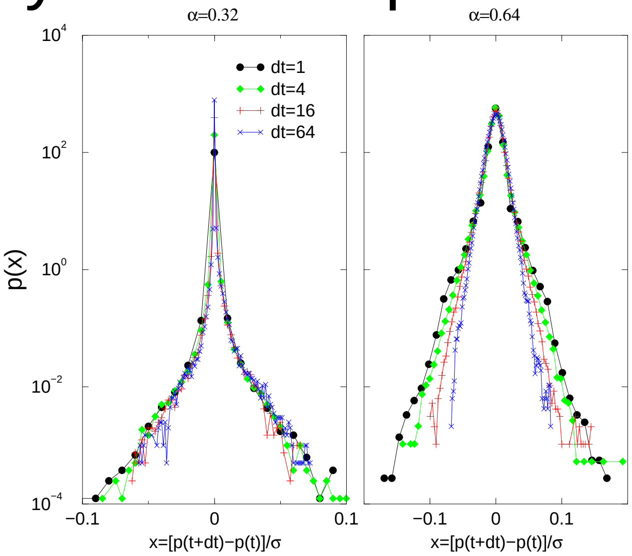

stylised facts,

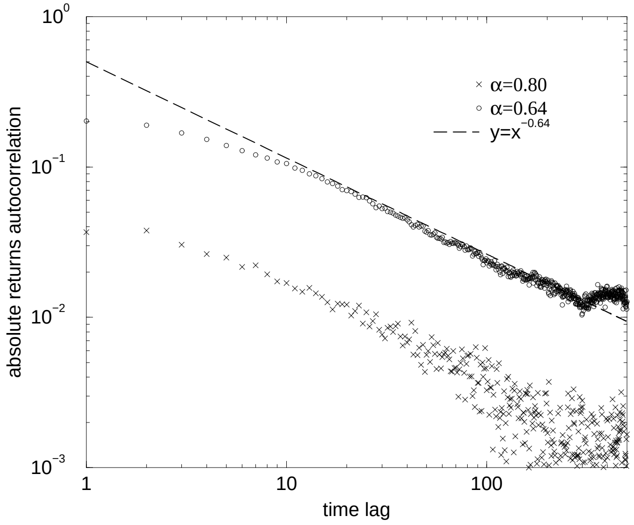

but only close to/below the phase transition

MG with dynamical capitals

MG with 2 strategies per player and dynamical capitals:

But no analytical theory.

complicated/tedious:

one has fast-evolving variables (the decisions of the agents) and slow ones (the capitals)

Simple MG with dynamical wealth

Simple MG with dynamical capitals:

$$c_i(t+1) = c_i(t) - \underbrace{\varepsilon c_i(t)}_{\text{investment}} \; \underbrace{a_i^{\mu(t)} \frac{A(t)}{V(t)}}_{\text{MG-type payoff}}$$

Similar to a replicator system with random couplings.

[T. Galla, ‘Random replicators with Hebbian interactions’, JSTAT 2005]

Simple MG with dynamical wealth

One strategy only per player - exact analytical solution:

Transition persists, and wealth → ∞ at transition in the infinite system

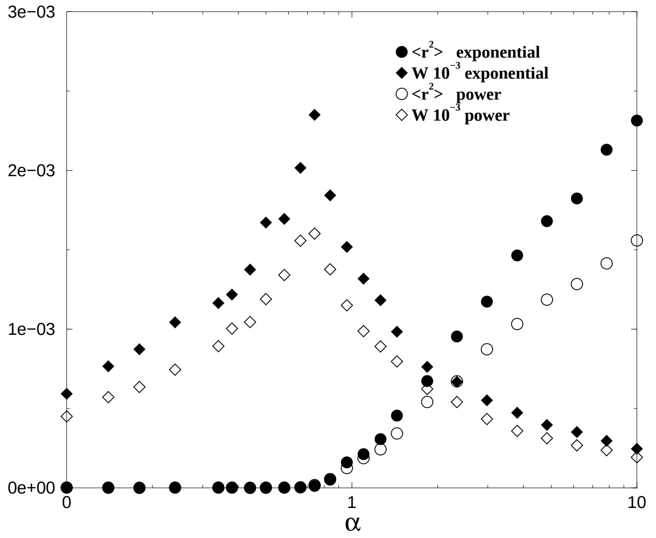

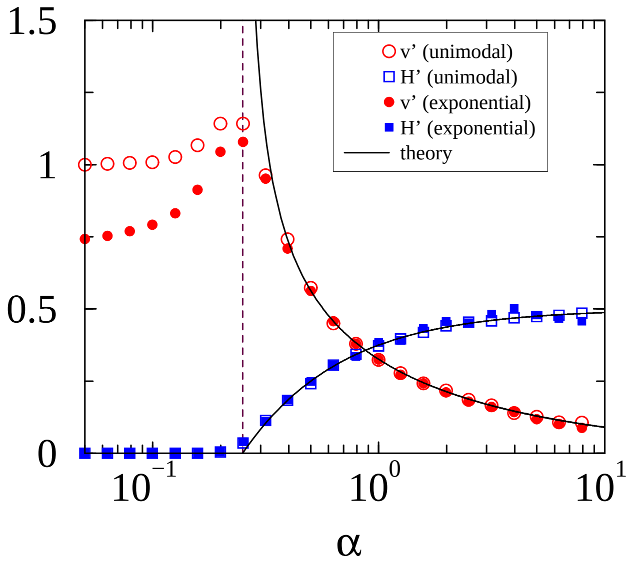

Simple MG with dynamical wealth

Distribution of returns (re-scaled to unit variance):

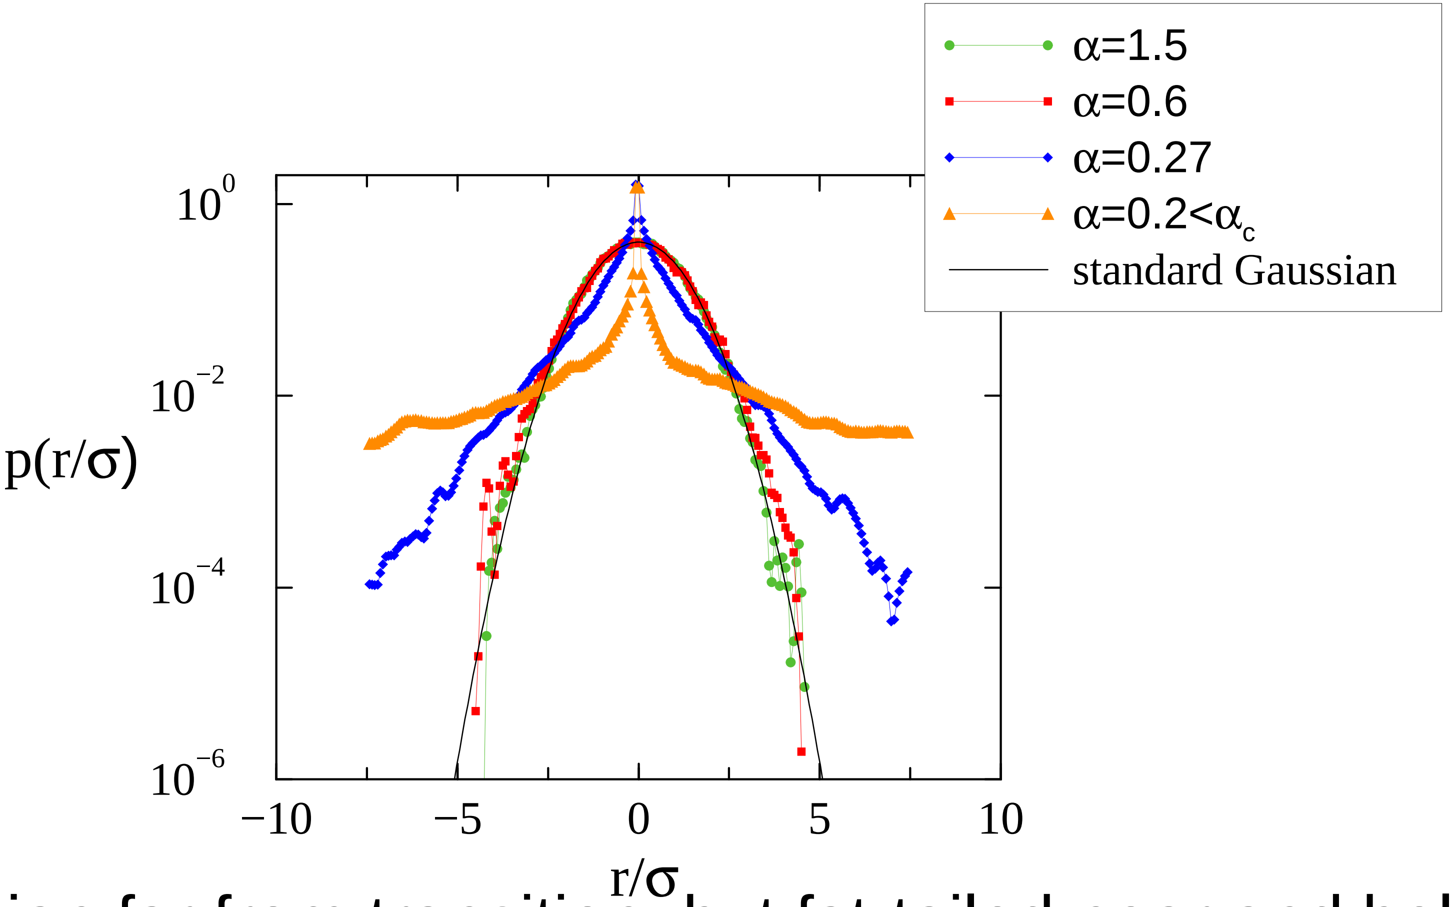

Gaussian far from transition, but fat-tailed near and below.

Simple MG with dynamical wealth

Distribution of wealth (re-scaled to unit variance):

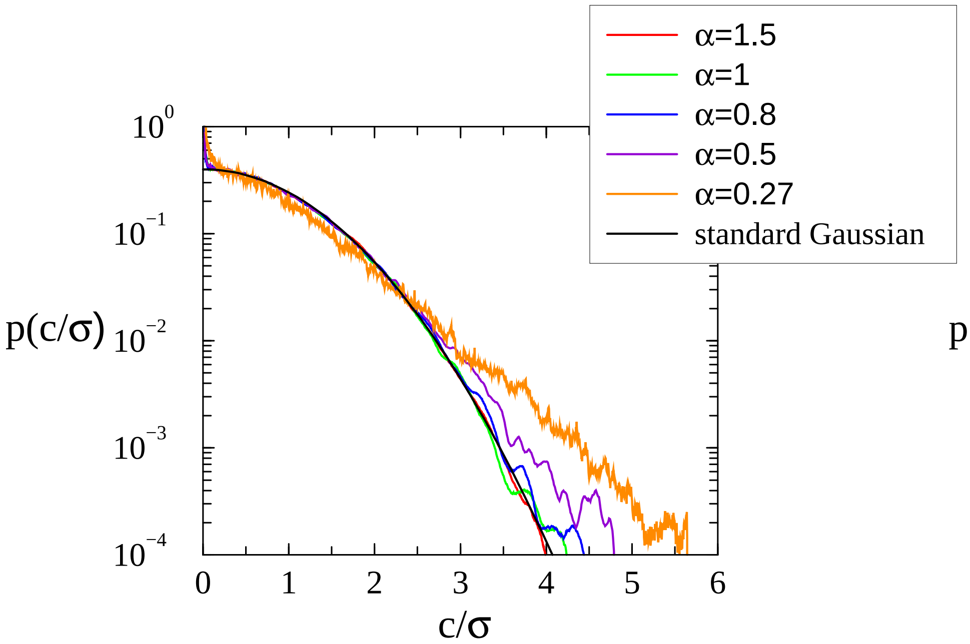

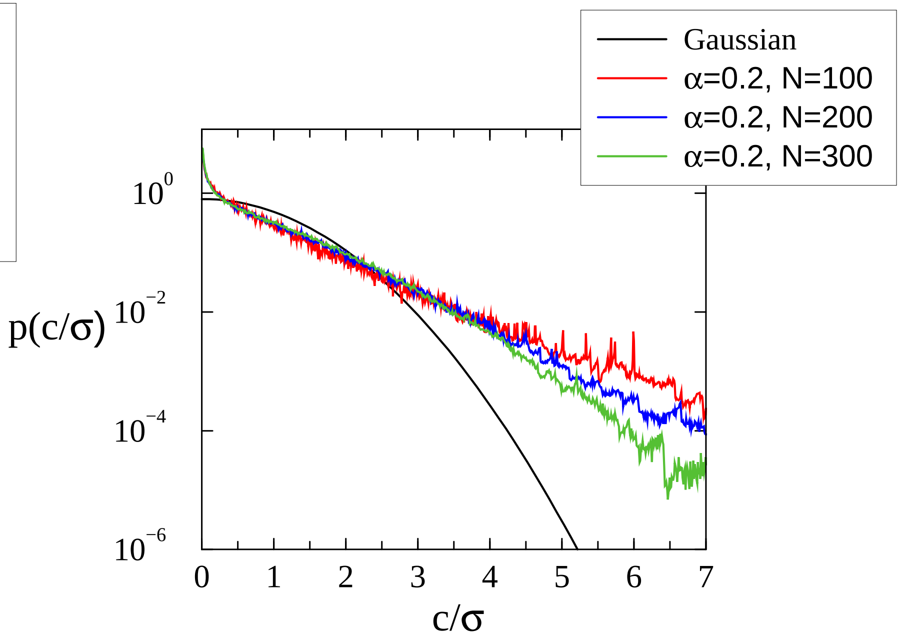

Fat tailed non-Gaussian distribution not a finite-size effect below transition ?

Tobin Tax in MGs

Tax revenue as function of trading fee

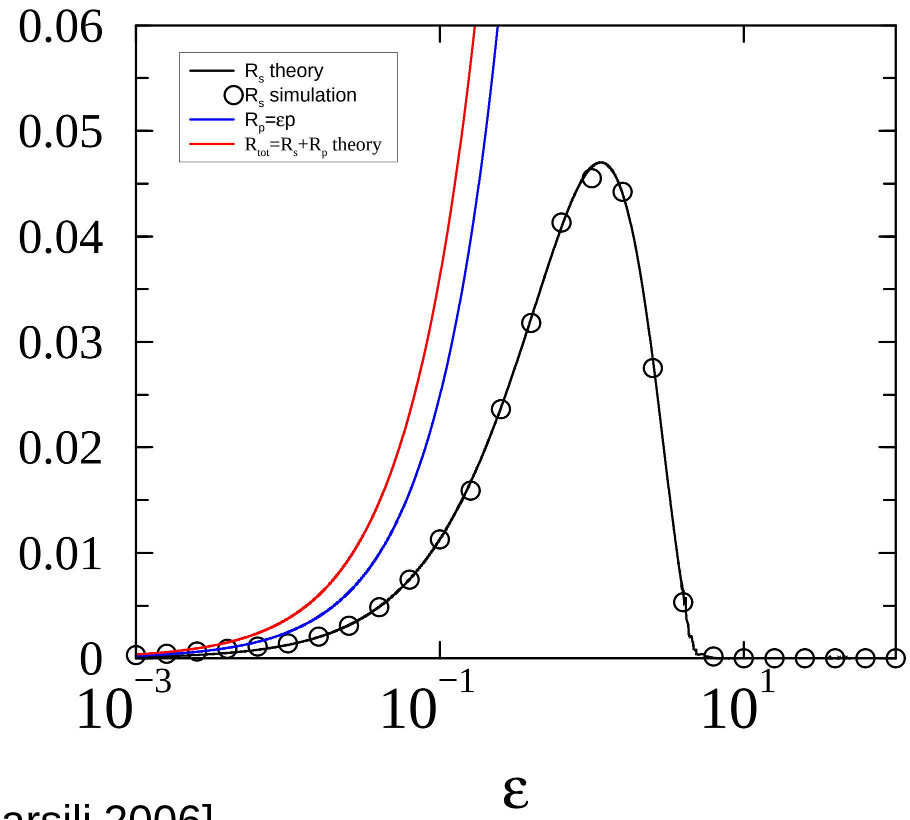

[Bianconi, Galla, Marsili 2006]

[Galla, Zhang in progress]

Conclusions

- MG has attracted attention from physics, mathematics and economics
- physics: spin glass problem with off-equilibrium dynamics
- open questions:
  - solution in non-ergodic region
  - critical exponents, RG ...
  - relation to spin-glass models and Hopfield model
- also to do: find more realistic extensions which are still analytically tractable