Simple Forecasting Heuristics that Make us Smart: Evidence from Different Market Experiments

Mikhail Anufriev Cars Hommes Tomasz Makarewicz

Rebuilding Macroeconomics Conference

HM Treasury, London, UK, 1-2 October, 2018

Outline

Behavioral Macroeconomics

Many deviations from rationality

- **rational inattention** (Sims, etc.)
- **sparse dynamic optimization** (Gabaix, etc.)
- **survey data** (Coibion, Gorodnichenko, Schneider, etc.)
- **learning from experience** (Malmendier, etc.)
- **expectations bias and overreaction** (Shleifer, etc.)
- **adaptive learning** (Eusepi and Preston, etc.)
- ...

- **New Behavioral Elements in this Talk:**
  1. laboratory experiments with human subjects;
  2. anchor and adjustment forecasting heuristics in a complex macro environment

Example Simple Heuristics: the Dog and the Frisbee

Simon, 1955; Kahnemann & Tversky, 1974; rational heuristic Haldane, 2012; simple heuristics that make us smart, Gigerenzer and Todd, 1999

simple heuristic: "run at a speed so that the angle of gaze to the frisbee remains roughly constant"

Which heuristic?: GA-learning model selects **optimal** heuristic within a **given plausible class** in a complex environment

Main contributions of the paper

- **Agent-based GA-model explaining different experimental data sets**
- Agents use Genetic Algorithms to **optimize a linear, anchor and adjustment** forecasting heuristic
- **Contribution:**
    1. **parsimonious** empirical micro-foundation of individual forecasting behavior in a macro system
    2. explanation of **individual** (micro) and **aggregate** (macro) behavior in three different macro experiments;
    3. improved **one-period** and **50-period** ahead forecasting of experiments (compared to 11 alternative learning models);
    4. characterization of the **mean/median/variance** of empirical distribution of forecasting heuristics in different environments

Why Experimental Macroeconomics?

- If a macro theory does **not** work in a (simple) controlled laboratory environment with human subjects, why would it work in **reality**?
- a macro experiment is a **group experiment**, where we study **individual** (micro) as well as the **emerging aggregate group** (macro) behavior in a **controlled** laboratory environment
- **empirical foundation** for individual decision rules for **behavioral agent-based models** (ABMs) to **discipline** wilderness of bounded rationality
- laboratory test for **policy analysis** to test policies in more realistic controlled lab macro environment

Outline

Learning to Forecasts Experiments (LtF)

Marimon, Spear and Sunder, (1993, 1994, 1995)
Repeated Keynesian Beauty Contest Game (Nagel, 1995)

- individuals **only** have to forecast prices, **ceteris paribus**, with all other behavior computerized by theory
- price depends on **average forecast**: $p_t = f(\bar{p}_t^e)$ or $p_t = f(\bar{p}_{t+1}^e)$
  - **one-period** versus **two periods** ahead
  - **positive feedback**: f is increasing
  - **negative feedback**: f is decreasing

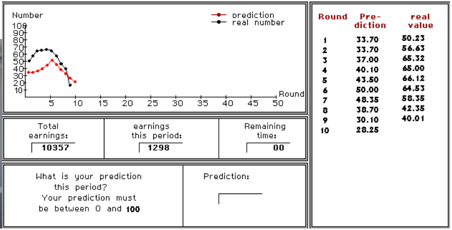

Positive versus Negative Feedback Experiments

Heemeijer et al. (JEDC 2009); Bao et al. (JEDC 2012)

- **negative feedback** (strategic substitutes environment)

$$p_t = 60 - \frac{20}{21}[\sum_{h=1}^{6} \frac{1}{6} p^e_{ht}] - 60] + \epsilon_t$$

- **positive feedback** (strategic complements environment)

$$p_t = 60 + \frac{20}{21}[\sum_{h=1}^{6} \frac{1}{6} p^e_{ht} - 60] + \epsilon_t$$

- **common feature**: same RE equilibrium 60
- **only difference**: sign in the slope of linear map $+0.95$ vs -0.95
- ϵ_t small **IID shocks** or large permanent **fundamental shocks**

Feedback Mappings in LtFE

negative feedback

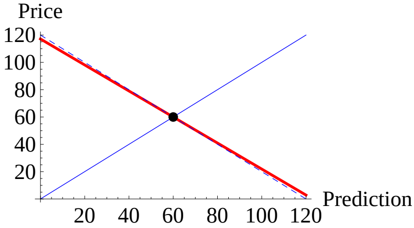

$$p_t = 60 - \frac{20}{21}\left(\overline{p_t^e} - 60\right) + \varepsilon_t$$

positive feedback

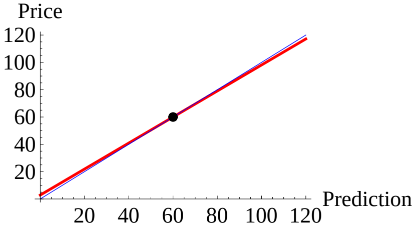

$$p_t = 60 + \frac{20}{21}\left(\overline{p_t^e} - 60\right) + \varepsilon_t$$

(strongly) **positive feedback systems** exhibit many **almost self-fulfilling equilibria**

Lab experiment (top) and 50-period simulation (bottom)

experimental data Heemeijer et al. (2009); small IID shocks

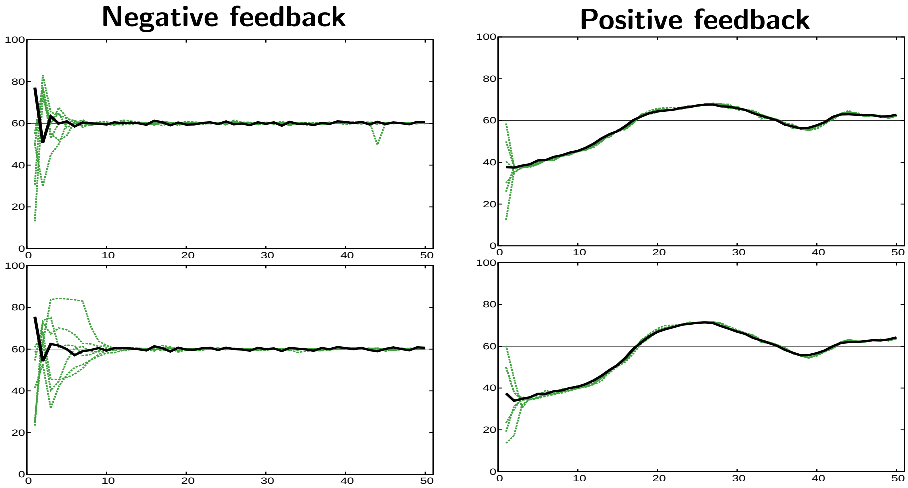

Lab experiment (top) and 65-period simulations (bottom)

experimental data Bao et al. (2012); large permanent fundamental shocks

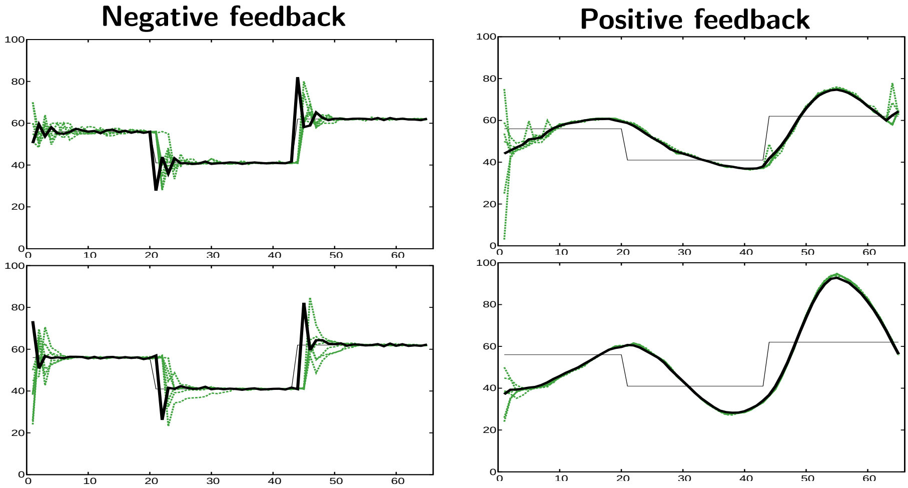

Asset Pricing Experiment

Hommes, Sonnemans, Tuinstra, v.d. Velden, Rev. Fin. Stud. 2005

(unknown) law of motion

$$p_t = \frac{1}{1+r}\left((1-n_t)\frac{p^e_{t+1,1} + \cdots + p^e_{t+1,6}}{6} + n_t\, p^f + \bar{y} + \varepsilon_t\right)$$

- **asset pricing experiment** (with/without robot trader)
  - two-period ahead
  - positive feedback
  - mean dividend $\bar{y} = 3$ and interest rate $r = 0.05$ are **known**
    rational fundamental price $p^f = \bar{y}/r = 60$ **not known**
    (but can be computed)
  - n_t is fraction of **fundamental robot traders**;
    increases as price moves away from fundamental

Asset Pricing Experiment

Hommes, Sonnemans, Tuinstra, v.d. Velden, Rev. Fin. Stud. 2005

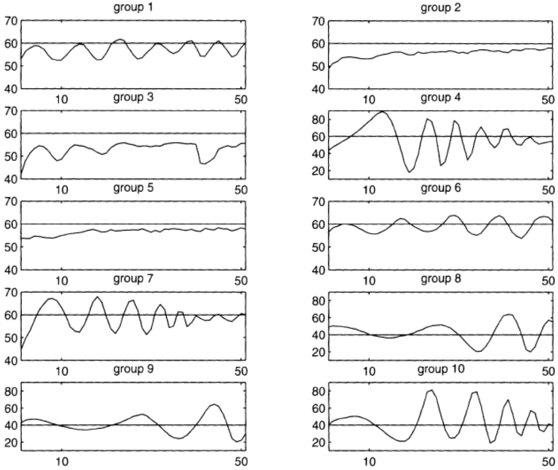

Hommes et al. (2005) experiment: *diversified* patterns between groups.

Lab experiment (top) and 50-period simulations (bottom)

experimental data Hommes et al. (2005)

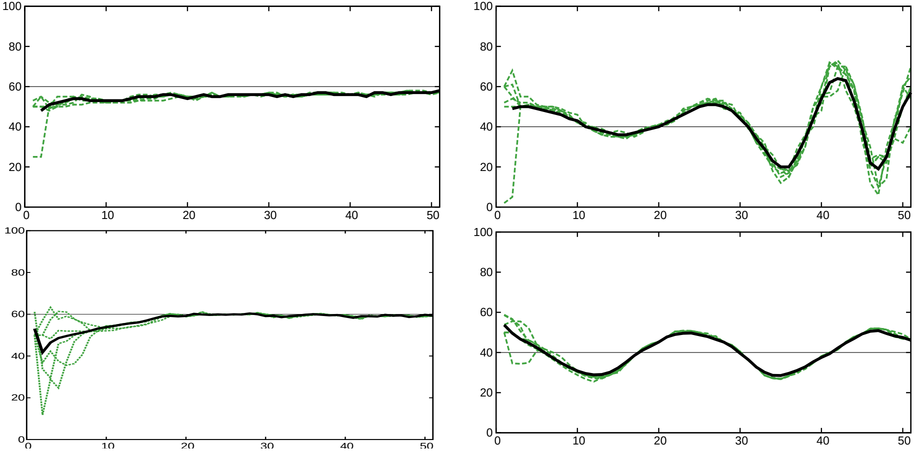

Outline

Behavioral Heuristics Switching Model

Brock and Hommes, ECMA 1997; Anufriev and Hommes, AEJ:Micro 2012

- agents choose from a number of simple **forecasting heuristics**
- **performance based reinforcement learning:**
  agents evaluate the **performances** of all heuristics, and gradually **switch** to more successful rules;

fractions of belief types are gradually updated in each period:
(discrete choice model with asynchronous updating)

$$n_{ht} = \delta n_{h,t-1} + (1-\delta)\frac{e^{\beta U_{h,t-1}}}{Z_{t-1}}$$

where Z_{t-1} is normalization factor.

- U_{ht} **fitness measure** (e.g. utility, forecasting performance, etc.)
- β is **intensity of choice**.
- δ **asynchronous** updating

Heuristic Switching Model: four forecasting heuristics

Anufriev and Hommes, AEJ:Micro 2012

- **adaptive expectations** rule, $[w = 0.65]$

$$\text{ADA} \qquad p^e_{1,t+1} = 0.65\, p_{t-1} + 0.35\, p^e_{1,t}$$

- **weak trend-following** rule, $[\ \gamma = 0.4]$

$$\text{WTR} \qquad p^e_{2,t+1} = p_{t-1} + 0.4\,(p_{t-1} - p_{t-2})$$

- **strong trend-following** rule, $[\ \gamma = 1.3]$

$$\text{STR} \qquad p^e_{3,t+1} = p_{t-1} + 1.3\,(p_{t-1} - p_{t-2})$$

- **anchoring and adjustment heuristic** with learnable anchor

$$\text{LAA} \qquad p^e_{4,t+1} = \tfrac{1}{2}\left(p^{av}_{t-1} + p_{t-1}\right) + \left(p_{t-1} - p_{t-2}\right)$$

Problem: HSM fits experimental data well, but **where do these heuristics come from**?

Outline

Agents use First Order Forecasting Heuristic

- Agent i uses a **first order forecasting heuristic** h to predict p_t:

  anchor and adjustment rule

$$p^e_{i,h,t} = \alpha_{i,h,t} p_{t-1} + (1 - \alpha_{i,h,t}) p^e_{i,t-1} + \beta_{i,h,t}(p_{t-1} - p_{t-2}).$$

- The rule h requires **two** parameters: an **anchor** $\alpha_{i,h,t}$ and a **trend** $\beta_{i,h,t}$
- General constraint: $\alpha \in [0, 1]$, $\beta \in [-1.1, 1.1]$.
- The rule generalizes popular HSM heuristics: naive, adaptive expectations and trend extrapolation.
- **RE**: $\alpha = 0$, $\beta = 0$, $p^e_{i,t-1} = p^f$.

Problem: how do agents learn the **anchor** and **trend parameters** of the heuristics?

GA's learning optimal forecasting heuristics

Simple heuristics that make us smart (Anufriev et al., 2018)

- Every agent i has a list of $H = 20$ different heuristics $(\alpha_{i,h.t}, \beta_{i,h.t})$.
- When agent i observes the last realized price p_{t-1}, she tries to re-optimize the rules with GA evolutionary operators:
  1. **reproduction**: **survival of the fittest**:
     sample (with replacement) 20 new heuristics from the old depending on their *forecasting performance*
     probability of selection $\Pi_{i,h,t} = \frac{\exp(-\mathrm{SE}_{i,h,t})}{\sum_{k=1}^{H}\exp(-\mathrm{SE}_{i,k,t})}$.
  2. **mutation** and **cross-over**: with some small probability "mutate" and "combine" them (modify (α, β) of each heuristic);
  3. **election:** compare the new and the old heuristics in terms of their hypothetical forecasting performance – pick the better ones.

Natural selection: worse forecasting heuristics are likely to be *replaced* by better heuristics; *inefficient experimentation* screened out.

GA model details

GA optimizes anchor and trend parameters α and β

Parameter	Notation	Value
# of agents	I	6
# of heuristics per agent	H	20
# of learning parameters	N	2
Allowed α price weight	$[\alpha_L, \alpha_H]$	$[0, 1]$
Allowed β trend extrapolation	$[\beta_L, \beta_H]$	$[-1.1, 1.1]$
# of bits per parameter	$\{L_1, L_2\}$	$\{20, 20\}$
Mutation rate	δ_m	0.01
Crossover rate	δ_c	0.9
Fitness measure	$V(\cdot)$	$\exp(-SE(\cdot))$

GA simulation analysis

1. Monte Carlo simulations — run the GA model 1000 times
2. 50-**period ahead** simulations — take experimental initial periods and run the GA learning model 1000 times: **how far does the model diverge from the data in the long-run**?
3. **One-period ahead** simulations — take experimental data until period τ and predict period $\tau + 1$: **how well does the model predict the data in the short-run**? (Sequential Monte Carlo analysis, 1000 runs.)

Outcome: GA model beats RE, simple heuristics, adaptive learning, HSM and other GA's.

50-period ahead Monte Carlo simulations (1000)

experimental data Heemeijer et al. (2009); small IID shocks

Negative feedback

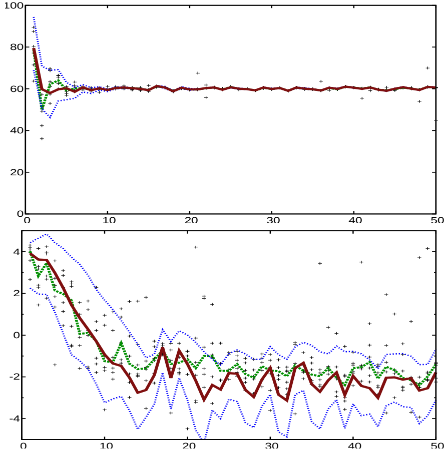

SD individual predictions

Positive feedback

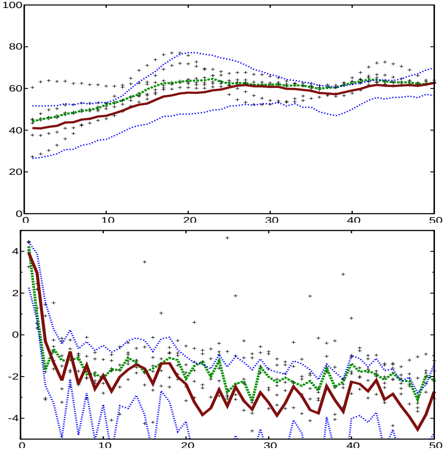

SD individual predictions

65-period ahead Monte Carlo simulations (1000)

experimental data Bao et al. (2012); large permanent fundamental shocks

Negative feedback

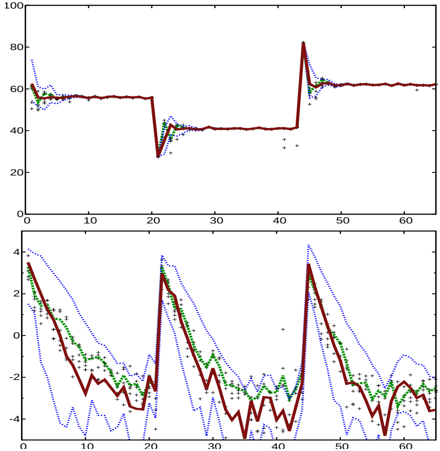

SD individual predictions

Positive feedback

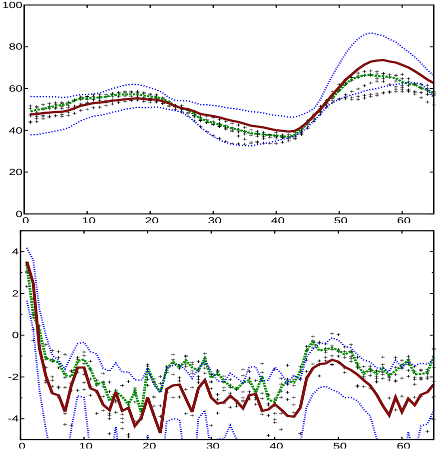

SD individual predictions

Mean anchor α_t and trend β_t parameters (1000 MC runs)

experimental data Heemeijer et al. (2009); small IID shocks

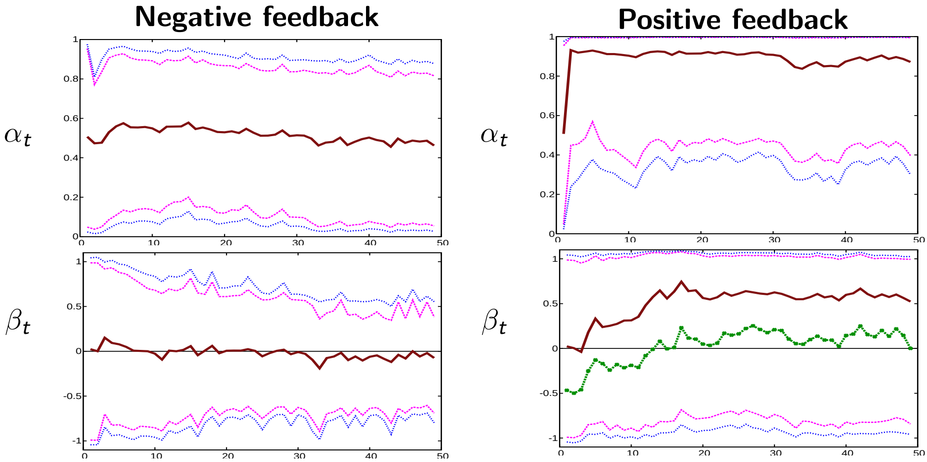

Mean anchor α_t and trend β_t parameters (1000 MC runs)

experimental data Bao et al. (2012); large permanent fundamental shocks

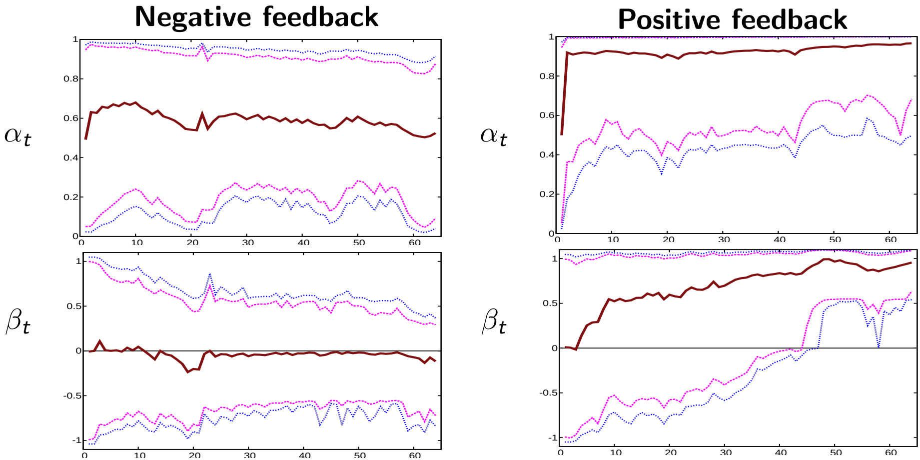

Mean/Median Heuristic

Under **negative feedback** agents learn to use **adaptive expectations**:

$$p^e_{i,t} \approx 0.5 p_{t-1} + 0.5 p^e_{i,t-1}$$

Under **positive feedback** agents learn to become **trend-follower**:

$$p^e_{i,t} \approx 0.95 p_{t-1} + 0.05 p^e_{i,t-1} + 0.9(p_{t-1} - p_{t-2})$$

Main insights

- GA model explains **convergence under negative feedback**, as well as persistent **oscillations under positive feedback**
- GA agents learn to give **more weight to the recently observed price** under positive feedback
- GA agents **learn to extrapolate a trend** under positive feedback.

Measuring predicting error: 65-period ahead MSE

experimental data Bao et al. (2012); large permanent fundamental shocks

	Negative feedback		Positive feedback	
MSE	**Prices**	**Forecasts**	**Prices**	**Forecasts**
Trend extrapolation ($\beta = 1$)	2736	1289	101.3	113.3
Adaptive expectations ($\alpha = -0.75$)	**3.629**	**10.75**	55	62.14
Contrarian ($\beta = -0.5$)	6.984	*14.45*	58.46	65.95
Naive	94.44	110.9	46.62	52.9
RE	13.871	20.923	55.133	60.859
LS learning	262.1	230.7	228.8	235.8
HSM - 2 types	73.57	87.86	90.8	101.8
HSM - 4 types	236.08	267.59	**32.18**	**37.01**
GA: $\beta \in [\mathbf{-1.1}, \mathbf{1.1}]$	8.01	21.97	*43.49*	**49.44**
GA: $\beta \in [\mathbf{0}, \mathbf{1.1}]$	*6.333*	17.39	*43.49*	*49.64*
GA: Action-based	29.73	90.7	179.8	200.2
GA: AR1	21.28	59.71	88.02	98.71

Measuring predicting error: 1-period ahead MSE

experimental data Bao et al. (2012); large permanent fundamental shocks

	Negative feedback		Positive feedback	
MSE	**Prices**	**Forecasts**	**Prices**	**Forecasts**
Trend extrapolation ($\beta = 1$)	114.061	121.329	1.183	2.165
Adaptive expectations ($\alpha = 0.75$)	**3.689**	**10.332**	3.776	4.618
Contrarian ($\beta = -0.5$)	5.92	*12.534*	4.737	5.559
Naive	9.979	16.81	2.411	3.286
RE	13.871	20.923	55.133	60.859
LS learning	21.109	30.023	3.283	5.156
HSM - 2 types	38.309	45.679	0.9996	**2.024**
HSM - 4 types	6.28	13.68	0.42	**2.01**
GA: $\beta \in [\mathbf{-1.1}, \mathbf{1.1}]$	10.247	21.464	*0.342*	2.059
GA: $\beta \in [\mathbf{0}, \mathbf{1.1}]$	*4.208*	15.267	**0.341**	*2.036*
GA: Action-based	34.865	67.753	40.333	66.451
GA: AR1	7.939	24.022	1.111	3.555

Asset Pricing Experiment; two-period ahead

Hommes, Sonnemans, Tuinstra, v.d. Velden, Rev. Fin. Stud. 2005

Long-run GA model simulation: two attractors

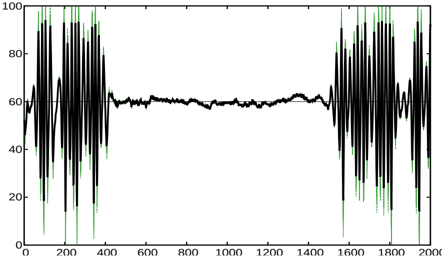

Prices and predictions

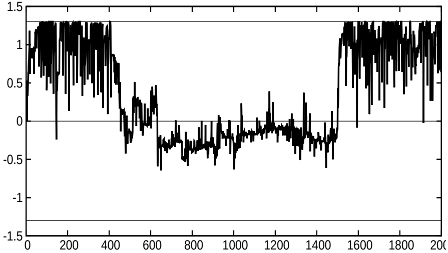

Average used trend extrapolation

GA model: two "attractors" coexist, and **agents can switch** between them $\rightarrow$ inherently **unstable learning**.

Positive feedback: complexity matters

Last period: distribution of trend extrapolation learned by agents.

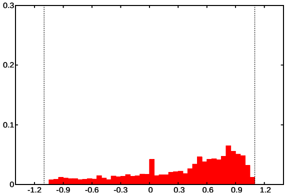

Linear pos. feedback
Heemeijer et al., 2009

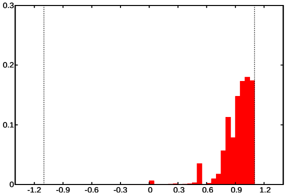

Linear pos. feedback
with shocks to the
fundamental
Bao et al., 2012

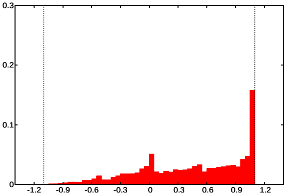

2-period ahead
non-linear feedback,
Hommes et al., 2005

Result: the more “complex” the feedback, the more likely more **extreme** trend extrapolation.

Outline

Conclusions

- **intuitive** and **parsimonious** way of selecting forecasting heuristics
- GA model of **individual learning**: agents learn and adapt forecast heuristics to their specific market environment.
- GA model explains **individual** (micro) as well as **aggregate** (macro) behaviour in different experiments
- **Positive feedback** environments are unstable, with fluctuations **amplified** by **coordination on trend-following heuristic**
- **coordination failures** in systems with unique equilibrium; coordination on almost self-fulfilling equilibria

Questions? Comments? Literature

Thank you for your attention!

1. Anufriev, M., Hommes, C.H. and Makarewicz, T. (2018), Simple heuristics that make us smart: evidence from different market experiments, *Journal European Economic Association*, forthcoming.
2. Hommes, C.H. (2018), Behavioral & Experimental Macroeconomics and Policy Analysis: a Complex Systems Approach, ECB working paper, forthcoming.
3. Hommes, C.H., Makarewicz, T., Massaro, D. and Smits, T. (2017), Genetic algorithm learning in a New Keynesian macroeconomic setup, *Journal of Evolutionary Economics* 27, 1133–1155.