



UvA-DARE (Digital Academic Repository)

Learning to forecast: Genetic algorithms and experiments

Makarewicz, T.A.

[Link to publication](#)

Citation for published version (APA):

Makarewicz, T. A. (2014). Learning to forecast: Genetic algorithms and experiments Amsterdam: Tinbergen Institute

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <http://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Chapter 5

Summary

Price expectations remain a matter of controversy in the economic literature. Recent market developments, including the 2007 financial meltdown, challenge the traditional view of perfect rationality and market efficiency. Learning-to-Forecast experiments (LtF) indicate that, when facing feedback between expectations and realized prices, subjects remain heterogeneous and follow simple behavioral rules, with their exact choice depending on the nature of the feedback. The most important difference depends upon whether the market has a negative or positive feedback structure. Under negative feedback, for example in a producers economy of a perishable good, the realized price is negatively correlated with expectations. This drives laboratory subjects to adaptive behavior and consequently they coordinate on the rational fundamental solution.

In contrast, financial markets feature a self-fulfilling positive feedback structure. For example, optimistic agents will increase demand for a financial asset, what indeed drives up the price of this asset. LtF experiments demonstrate that human subjects in such a positive feedback environment learn to extrapolate price trends, which results in erratic and, most likely, non-converging price oscillations.

The aim of this thesis is to study learning to forecast behavior. We focus on a Genetic Algorithms based learning model, its fit to LtF experimental data and application to various market settings. We also run an experiment to study the robustness of the LtF design. The thesis consists of three related but independent chapters.

In Chapter 2, we develop a model in which agents use Genetic Algorithms to optimize a general forecasting heuristic to form their price expectations (Hommes and Lux, 2013). We apply the model to four experimental settings:

- linear positive and negative feedback (Heemeijer et al., 2009);
- linear positive and negative feedback with large and unanticipated shocks to the fundamental price (Bao et al., 2012);

- non-linear cobweb economy (Hommes et al., 2007; van de Velden, 2001);
- non-linear two-period ahead asset pricing market (Hommes et al., 2005).

Our model outperforms Rational Expectations, a number of other homogeneous expectation rules (naive, adaptive, trend extrapolation and contrarian), and a simple Heuristic Switching Model with heterogeneous expectations. Specifically, the model has an excellent short-horizon out-of-sample predictive power. In addition, it is the only model able to explain the *individual* behavior, and successfully predict the experimental dynamics in the *long run* (fifty periods ahead). It also justifies using the Heuristic Switching Model as a stylized description of experimental dynamics. This is an important contribution to the literature, since in recent years policy makers are looking for economic models with robust behavioral micro-foundations.

In Chapter 3, we use the Genetic Algorithms Learning-to-Forecast model to study information networks in financial markets. The goal is to evaluate the impact of information flows on market stability and individual coordination. We focus on a non-linear two-period ahead asset pricing model, in which the agents can learn whether to extrapolate the observed price trend, but also whether to trust the average mood of the friends that the agents observe through a network. We show that without a network, agents can coordinate either on the fundamental value, or on trend-following behavior. The latter makes the market switch between the fundamental solution and erratic price oscillations. Adding a network of any architecture or size destabilizes the market, which almost never stays in the fundamental solution. This follows from two related observations. First, agents learn to coordinate on stronger price trend following rules. Secondly, agents learn contrarian behavior: they are more optimistic if they observe their friends to be selling in the past and *vice versa*. The reason for the contrarian behavior is that the agents try to ‘out-smart’ their friends. For example, in the moment when the bubble starts to collapse, an agent realizes the negative trend, as well as the fact that until the peak of the bubble, her friends were buying the asset. She therefore has an incentive to trade against her friends.

The model predicts that, despite the contrarian attitude, the agents learn a similar type of behavior and therefore remain well coordinated. This is an important insight into the literature’s conflicting view on herding. Experiments demonstrate that people use contrarian (anti-herding) strategies, which indeed occurs in our simulations. On the other hand, indirect measures used for the market data suggest some degree of herding. In the model, similar learning of *contrarian* attitude drives agents to *similar behavior*, which may be mistaken for herding.

Chapter 4 reports an experiment based on a simple linear asset pricing market,

in which subjects were asked to forecast the prices, trade the asset or do both. The importance of this experiment is that it establishes a link between Learning-to-Forecast and Learning-to-Optimize experiments. It can therefore serve as a useful benchmark for future theoretical and experimental investigations on learning in asset pricing markets.

We find that regardless of the treatment, subjects coordinate on non-fundamental outcomes: stable price far from the fundamental value or persistent price oscillations. The Learning-to-Optimize and mixed treatments are more unstable, with the highest bubble (around 3.5 times the fundamental price at the peak) in one of the groups from the mixed treatment. These results show that the learning to optimize is even more difficult for the subjects than the learning to forecast. Furthermore, the results of the LtF experiments are robust, in the sense that price oscillations are not just an artifact of the forecasting treatment. We also confirm statistically significant heterogeneity of subjects' behavioral rules.

The most surprising result comes from the mixed treatment, in which both the trades and the corresponding price forecasts are observable. This allowed us to explicitly test their consistency. We found only a quarter of our subjects to trade optimally conditional on the implied expectations of the asset return. This stands as a striking warning against the idea that economic agents are perfect optimizers — an idea popular even among the proponents of boundedly rational price expectations. Further studies shall focus on theoretical models of learning to optimize, and its link to learning to forecast.