



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 1

Introduction

1.1 Learning, rationality and markets

Learning is one of the defining capacities of man, a fact that has been recognized by most of social sciences and philosophy. The 20th century witnessed advances in modern logic, specifically inquiries about what can be represented and proven within boundaries of a formal language. Since such formal languages stand as a necessary backbone of rationality, these logical investigations gave important insights into the matter of learning itself, with the most important applications in the field of artificial intelligence and computer sciences (Negnevitsky, 2005). On the other hand, psychologists, neuroscientists and social scientists in general have been trying to understand how learning is realized within our minds (Carey et al., 2014). The human brain continues to fascinate with the ability to solve complex problems in a matter of seconds, which seemingly contradicts its affinity for systematic biases and quirks (Kahneman, 2011).

Economic sciences stand as an interesting exception. The bulk of economic literature focuses on the paradigm of perfect rationality, in which economic agents see through the underlying market structure, understand the strategic nature of their interaction and possess mutual knowledge of each other's (and common) rationality. The resulting rational equilibrium is 'perfect' in the sense of, broadly speaking, self-consistent and optimal response to the market constraints and decisions of other agents.

Learning is often mentioned as a justification for perfect rationality, nevertheless, it appears seldom as an explicit feature of typical mainstream economic models, especially in the context of finance or macroeconomics (see Evans and Honkapohja, 2001; Sargent, 1993, for a discussion). In contrast to what is implied by modern logic (Binmore, 1987) economists typically take the learning as automatically *perfect* and thus having no important dynamics of its own. This practice follows the classical arguments of

Friedman (1953); Lucas Jr. (1972); Lucas Jr (1986); Muth (1961). It is difficult to find a contemporary explication of the rational paradigm's virtues (or validity), since its proponents take it as a methodological standard, a modeling tool that requires no further justification or explanation. An exception can be found in the concluding remarks of Blundell and Stoker (2005). How does the rational approach fare in modeling practice?

One of the key instances of learning in the economic context, and the focus of this thesis, are price expectations, namely *learning to forecast* prices. In a plethora of economic problems, agents need to decide on a specific action in the present, whereas the profitability of this action depends on future prices. For example, firms may face a production lag: they have to set up the produced quantity today, but they will sell it tomorrow, with tomorrow's demand and price. Another important case are investors who buy financial assets, like stocks, in the hope that these assets later gain more value. In order to make their decisions in a reasonable fashion, both the producers and investors need somehow to forecast future prices.

The framework of perfect rationality emphasizes Rational Expectations (RE), understood as model-consistent price forecasts. In the example of the firms, every producer will optimize production based on a price forecast, that in turn follows the *accurate* belief of common mutually optimal responses to the underlying market clearing mechanism and fundamentals. The latter include distribution of individual market power, demand structure and production technologies. This leads to market clearing (aggregate production is equal to the demand), where price forecasts are self-fulfilling and no firm could improve its profit. The example of financial markets has a similar RE equilibrium: investors optimize their trades conditional on what turns out to be the realized price given these individual trades.

One can show that a RE equilibrium exists under mild assumptions about the market structure. However, the claim of the perfectly rational framework is that the RE equilibrium is the actual state of empirical markets. As mentioned above, this strong belief relies on a particular explication of individual *learning*. The argument goes that if agents would make poor price forecasts, they would lose profit in comparison with other agents. In a rational response, they would adjust their expectations. Therefore, agents will not make systematic mistakes and RE appears as a fixed point of the aggregate learning process (Muth, 1961).¹ However, proponents of RE typically do not take this reasoning as an explicit model feature. Instead, they would think that

¹In addition, one can use a similar evolutionary argument based on so called social learning: agents that make mistakes are replaced by smarter agents. Therefore RE is selected, or 'learned', in aggregate terms (Friedman, 1953).

learning pushes the agents to the RE equilibrium, which can therefore be used as a valid approximation of market dynamics. In other words, learning to forecast is *perfect*: it comes as a crucial foundation of the model, yet it does not have to be investigated on its own.

This approach has two conceptual problems, however. First of all, it assumes what remains an empirical issue. We know that human learning is bounded (Kahneman, 2011), also by the very nature of limitations of logical languages (Binmore, 1987). It is therefore unclear *a priori* whether learning to forecast can end up with a ‘perfect’ outcome. This issue becomes even more severe if the economic environment is complex and evolves over time.

The second conceptual issue is that the RE equilibrium, as a fixed point of the individual learning and decisions, is inherently *static*. In the example of the producers economy, the structure of the RE equilibrium remains intact if the production problem becomes repeated: every period has the same realized (market clearing) price. Furthermore, every agent knows that from the beginning. Instead of trading repeatedly on a period-to-period basis, agents in the first period could simply sign a ‘social contract’ that would cover all their actions in all future periods (regardless of the time horizon). And *nothing* would change in the realized market dynamics.²

Under RE, the market price can change only if there is a shock to the underlying fundamentals, such as production technology, preferences or asset dividend. Such shocks are purely exogenous: they are not anticipated and do not follow from anything in the model itself, but rather appear *deus ex machina*. Furthermore, they cannot push the economy to any meaningful dynamics. For example, an exogenous productivity shock would simply make the firms re-optimize their production portfolios, jump to the new market-clearing equilibrium and stay there forever (or until a new shock hits the economy). Similarly, DSGE models, which stand as a backbone of modern macroeconomics, explain the business cycle purely by such exogenous shocks (see An and Schorfheide, 2007, for a theoretical example together with an empirical application). These models do impose additional frictions (*e.g.* on price adjustment) in order to match the empirical persistence of macro dynamics; but their spirit remains intact.

The static nature of RE equilibria stands in contrast to what one intuitively thinks

²Assuming complete markets, the same holds for a (dynamic model) with uncertainty which leads to the Arrow–Debreu equilibrium. Here we leave aside the issue of multiplicity of RE equilibria. The agents in such a case could coordinate on jumping between the possible equilibria, even on a basis of period-to-period random signal. Nevertheless the reasoning stays intact, as again in the very first period the agents could sign an appropriate contract that would specify the period-to-period coordination mechanism conditional on the realized state of the economy. See Benhabib and Nishimura (2012) for an example of these ‘sunspot equilibria’.

of learning. Namely, agents learn by a process of trial-and-error, experiment with different strategies or wait until they have enough data for appropriate statistical inference, which then they update every time they access new relevant observations. And even if the agents are able to eventually learn the RE equilibrium, a shock to the economy would force them to re-evaluate their knowledge (such as their beliefs about productivity or fundamentals). Learning is therefore a ‘sticky’ and largely unpredictable process, which would impose additional inertia in economic variables that the RE approach cannot acknowledge, or does not want to acknowledge (Sims, 1980). Therefore, learning implies *bounded rationality* (Simon, 1972).

1.2 Evidence from markets and experiments

In the context of price expectations, how important is the fact that learning is a dynamic and possibly imperfect process? There is no *a priori* theoretical answer to this fundamentally empirical question. Instead, we can judge the relevance of learning to forecast through studies of market data and experiments.

Learning is not easily identifiable in market surveys, as these rarely accord with the assumption of *ceteris paribus* and hence can serve only as an indirect proof. Over the last three decades, however, economists gathered substantial evidence that the empirical price expectations are much more complicated than the rational framework implies. A natural alternative is thus to interpret the data as a signal of bounded rationality and learning dynamics. Specific examples include:

- Consumer inflation expectations: studies of micro-data surveys clearly indicate that households retain heterogeneous price expectations and most likely learn only from their individual experience, disregarding the full history of macro data (Malmendier and Nagel, 2009). They can be subject to systematic biases like ignoring the business cycle (Thomas Jr., 1999);
- Financial markets: asset prices exhibit patterns that are ‘paradoxical’ from the RE perspective (see De Long et al., 1990, for a discussion), including the following three most famous examples:
 1. difficulties in explaining asset prices with bare fundamentals (Israel and Moskowitz, 2013);
 2. investor overreaction to market signals (Bondt and Thaler, 2012);
 3. excess volatility of stock prices (Shiller, 1981, 2003).

Furthermore, in the recent two decades we experienced several financial meltdowns (Kindleberger and Aliber, 2011; Reinhart and Rogoff, 2008), which undermined the conviction of financial markets' efficiency and rationality. Among the most recent crises were:

1. 'Black Wednesday' event, which caused the UK to leave the European Exchange Rate Mechanism (Söderlin, 2000),
 2. East Asian crisis of the late 1990ties (Best, 2010; Mendoza, 2010);
 3. 'dot-com', or internet technologies bubble (Griffin et al., 2011; Morris and Alam, 2012);
 4. financial meltdown of 2007 (Erkens et al., 2012; Goodhart, 2008; Reinhart and Rogoff, 2009; Shiller, 2008), with an interesting case study of Iceland (Aliber and Zoega, 2011);
 5. bubbles in the markets of the so called 'cryptocurrencies', like Litecoin, Dogecoin or Bitcoin (Yermack, 2013).
- Housing market: house prices in the USA experienced a bubble due to systematic over-evaluation of the fundamentals before 2007 (Case and Shiller, 2003), and the signs of this can be traced back even to 1990ties (Goodman Jr. and Ittner, 1992). Similar patterns can be observed in other countries, see Ambrose et al. (2013) for a three and a half century data set of the housing market in Amsterdam;
 - Producers inflation expectations: survey inflation forecasts fail to comport rational expectations (Mavroeidis et al., 2014), while at the same time explaining an important part of firms' decision making (Nunes, 2010a).

Laboratory experiments serve as an alternative to market surveys in evaluating human behavior (Smith, 2010). They offer a controlled setting in which researchers can directly setup the structure, fundamentals, information feedback and incentive schemes of the investigated market.³

The most popular economic experiments are built over Game Theory settings, with simple benchmark games such as 'p-beauty contest' (Duffy and Nagel, 1997; Ho et al., 1998), ultimatum game (Güth et al., 1982), centipede game (McKelvey and Palfrey, 1992) and prisoner's dilemma (Andreoni and Miller, 1993; Fehr and Gächter, 2000).⁴

³For the sake of fairness, it must be emphasized that the use of experiments for testing financial and macroeconomic models does raise some methodological issues of the so called external validity, see Guala and Mittone (2005) for a discussion.

⁴Apart from the question of rationality, the first experiments on Game Theory sparked a discussion about the so-called other-regarding or social preferences, see Fischbacher and Gächter (2010) for literature overview and an example.

A more sophisticated applications, often based on oligopoly games, became popular in Industrial Organization (Huck et al., 1999; Offerman et al., 2002). The rational solutions to these games, understood as (refined) Nash Equilibria, are not necessarily the unique experimental outcome, with costly punishment in the public good game being the most famous example of a robust non-rational finding (Fehr and Gächter, 2000). And even if subjects converge to a Nash Equilibrium, they may need time to do so (Smith, 2010). For the example of the Industrial Organization setting, Offerman et al. (2002) investigate a simple oligopoly game. The authors show that subjects learn different equilibria (perfectly competitive, Cournot-Nash and collusive) depending on the information framing, despite the same underlying production and demand functions.⁵ These findings constitute strong evidence in support for the importance of individual learning.

An interesting example is the above mentioned ‘p-beauty contest’, which models a positive feedback between predictions and realization. This captures the spirit of financial markets, with self-fulfilling investor sentiments (Sonnemans and Tuinstra, 2010). In addition, the subjects are given perfect information about the underlying strategic structure of the game and therefore have all means to compute the self-consistent equilibrium already in the very first period of the game. However, experiments show that subjects require repeated interaction before they converge to the Nash Equilibrium. An even more striking finding is that the subjects remain largely heterogeneous and seem to use strategies of diversified levels of sophistication (compare with the survey on consumer’s inflation expectations by Malmendier and Nagel, 2009). This suggests that they undergo a process of learning, which remains ‘imperfect’ from the classical perspective.

Another class of financial market experiments comes with the seminal work by Smith et al. (1988). The authors investigated a market, in which subjects could trade an asset with a declining fundamental price. The result was consistent miss-pricing: subjects would overprice the asset for a substantial number of periods, and the resulting bubbles eventually crashed when the asset fundamental value approached zero. Follow-up studies focused on testing a number of suggested explanations of the experimental bubbles: uncertainty about rationality of other market participants, or ‘rational’ attempts to outsmart other subjects and play out the bubble (Lei et al., 2001); problems with understanding the declining fundamental (Huber and Kirchler, 2012; Kirchler

⁵The outcome of Offerman et al. (2002) shows another issue with the rational framework. In the example of the producers game, equilibrium price and welfare crucially depend on whether the firms are price-takers or play a Nash Game. Since the rational framework contains no explicit learning, it cannot explain why in some markets agents realize that they have significant market power, while in other they consider themselves price-takers.

et al., 2012; Noussair et al., 2001); lack of trading experience (Dufwenberg et al., 2005; Smith et al., 1988); even ‘psychological irrationalities’ like trading to avoid boredom (Active Participation Hypothesis; Lei et al., 2001). It seems that these factors at best only partially explain the experimental bubbles (Noussair and Tucker, 2013).

The trading experiments suggest that the subjects indeed learn (for instance the role of experience is clear), but they learn behavior of a very different nature than what the rational framework postulates. Instead of the perfect, model-consistent price expectations, subjects are rather trying to come up with simple behavioral ‘rules of thumb’ that are *good enough so far*. As a result, their behavior is much more backward-looking and ‘imperfect’ than expected.

Since price forecasts are not directly observed in trading decisions, the natural step for experimentalists was to set up laboratory studies in order to control the forecasting itself (Marimon et al., 1993).⁶ In these *Learning-to-Forecast experiments* (see Assenza et al., 2014a; Hommes, 2011, for a comprehensive literature overview), subjects play the role of forecasting advisers to computer agents, such as financial investors or producers. Subjects give a price forecast, which the computer agents then use to optimize their decisions. This leads, through a specific market clearing condition, to the realized price and the subjects are paid based *only* on their forecasting performance.

As a result, we obtain an economically founded feedback between prices and price predictions that can be used to study individual learning conditional on the specific structure of the feedback. The two most important cases turned out to be negative and positive expectations feedback. In the earlier example of the producers, if they expect a high price, they overproduce and the realized price will be low. This is a case of negative feedback. In contrast, financial markets are characterized by self-fulfilling moods: if the investors are optimistic, they will buy more assets, which drives up the asset prices, yielding a positive feedback between expectations and realized market prices.

Heemeijer et al. (2009) show the significant differences of the two types of feedback. The authors report that in their simple linear setup, under the negative feedback treatment subjects quickly converge to the fundamental price, whereas the positive feedback treatment can induce significant price oscillations (see Sonnemans and Tuinstra, 2010,

⁶This experimental design has been criticized as simplistic, because in real markets agents are directly asked for trading decisions instead of price expectations. Among critics of experimental economics prevails a misconception that laboratory experiments should be as close to real economic environments as possible. From this perspective, any simplification of the experimental economy is perceived as a design flaw. This is a misunderstanding of the role of experiments. Their goal (and what natural scientists realized already in 16th century) is to *isolate* one specific factor in a setting *simple enough* that we can study it *directly*, without a need of dissecting it from other phenomena, unlike in real markets (see Smith, 2010, for some discussion).

for a discussion of positive feedback). A follow-up study by Bao et al. (2012) confirms that these dynamics are robust against large and unanticipated changes to the fundamental price. Next to the sign, the complexity of the feedback plays an important role. Hommes et al. (2007) investigate a negative feedback system based on a nonlinear cobweb economy and find ‘excess volatility’: prices fluctuate irregularly around the fundamental price. The more the economy is unstable under the assumption of naive expectations, the more volatile the market becomes. As an example of a complex positive feedback, Hommes et al. (2005) study a two-period ahead non-linear asset pricing market. Their subjects coordinate on price oscillations of varied amplitude and period, which can furthermore change their pattern in a single 50-period session.

In line with findings of the other experiments, Learning-to-Forecast experiments indicate a high level of heterogeneity between the subjects, even at the level of their behavioral rules (Anufriev and Hommes, 2012; Heemeijer et al., 2009; Hommes, 2011). Again, this can be interpreted as a sign of independent learning, which is furthermore sensitive to the specific market environment. The RE solution clearly does not fit these learning dynamics. A pattern that seems to emerge instead is that the subjects learn, in rough terms, adaptive type of expectations under negative and trend following behavior under positive feedback.

The most important limitation of Learning-to-Forecast experiments is that they are based on markets in which the computerized agents behave optimally conditional on the subjects’ price forecasts.⁷ In reality this assumption does not necessarily hold. Economists would typically take it for granted, however, and so there are only few studies that directly challenge it. A notable exceptions is a recent experiment by Bao et al. (2013), in which the authors ask their subjects both to (1) forecast prices and (2) setup the production (*Learning-to-Optimize* experimental design) in a negative feedback producers economy. They identify behavior that seems to contradict the assumption of optimal trading conditional on price forecasts. However, the design of the feedback eventually pushes the subjects to the fundamental solution. The non-optimal behavior may be more relevant in less stable markets with positive or non-linear negative feedback feature, but to our best knowledge this has not been systematically examined (see Assenza et al., 2014a, for a recent overview of the existing literature on the topic). We will investigate this in detail in Chapter 4.

⁷As explained, this assumption is also the cornerstone of the rational framework. See Smith (2010) for a discussion about other underlying assumptions of rational behavior that do not survive laboratory testing.

1.3 Models of learning

1.3.1 Rational learning

The bulk of the economic literature did not respond to the criticism of perfect rationality and retains this framework to this day. As a result, theoretical models in economics seldom contain explicit learning features. For instance, this is visible in the design of typical DSGE models, which serve as the workhorse of the theoretical policy-oriented macroeconomics. However, there are some cases of ‘rational learning’. The most popular (especially in Game Theory and finance) is Bayesian updating: agents have a prior belief that they update conditional on unveiling market signals. This can lead to interesting dynamics, such as information cascades (Anderson and Holt, 1997). A more recent development are models of rational inattention (Sims, 2003, 2010): agents cannot perfectly process market information and instead optimize a noisy reaction to the noisy market environment. The approach of rational inattention is somewhat uncommon due to the technical challenges involved in solving such models, but applications can be found in macroeconomics (Maćkowiak and Wiederholt, 2009) and market organization literature Willems (2012).⁸ Neither Bayesian updating nor rational inattention can exhaust the full meaning of learning, however, since learning ought to be understood as something more fundamental, active and insecure than just being subject to an information noise, which is processed in an optimal fashion.

Some proponents of the rational paradigm followed this notion of learning and tried to use it to defend perfect rationality. One important case is the work by Guesnerie (1992) on Eductive Learning. It is based on an idea that agents, following the belief that no agent will use dominated strategies, can iteratively reduce their strategy sets by eliminating those strategies that became dominated after previous iteration of this algorithm. Eductive Learning can be applied to varied economic models (including macro and finance models) and often supports rational solution as the limiting case of the belief of common rationality (Guesnerie, 2002).

Another attempt to use learning as a defense of the rational paradigm comes with so called adaptive learning (Evans and Honkapohja, 2001), which has important applications in macroeconomics (see Evans and Honkapohja, 2009, for extensive overview and examples). Empirical agents, regardless of their rationality, in practice face a task of econometric nature, such as estimating relevant price or production elasticities.

⁸Willems (2012) gives a very interesting example of a monopoly firm that over time has to learn its demand of a linear form. The author notes that the model becomes intractable once the learning involves a noisy ‘rationally inattentive’ response to both the slope and the constant of the demand. This striking result casts serious doubts about the empirical relevance of rational inattention models.

ties (Sargent, 1993). Under adaptive learning, agents perceive a law of motion of the economy and learn its parameters over time, *e.g.* through recursive least squares. A natural question is where such a feedback can lead, and the literature on adaptive learning takes prime interest in the conditions for stability and uniqueness of the RE solution (Evans and Honkapohja, 2001). Additionally, adaptive learning can be used as a selection criterion in case of multiplicity of RE equilibria (Evans and Honkapohja, 1999). Finally, in reality agents presumably have to learn not only the parameters of the law of motion of the economy, but its specific functional form as well. This can lead to non-fundamental dynamics, especially in the case of significant non-linearities, and sparked a literature on restrictive perception equilibria (Branch, 2004; Bullard, 1994; Hommes and Sorger, 1998; Hommes and Zhu, 2014).

The example of Eductive Learning and adaptive learning are similar in spirit. They serve as the natural and unquestionably involved explication of the original, somewhat informal argument of Muth (1961) for the RE. However, due to the intentions of their authors, they cannot escape the criticism of the perfect rationality paradigm that they were supposed to counter in the first place. Both approaches, without any regard to the empirical evidence, simply assume a high level of individual rationality as a necessary characteristic of learning.⁹ As a result, Eductive Learning is unable to explain why in the experiments the subjects require repeated interaction to converge to the rational solution (Duffy and Nagel, 1997; Sutan and Willinger, 2004).¹⁰ On the other hand, adaptive learning is at odds with the heterogeneity and ‘irrationality’ of the household and experimental price forecasts (Malmendier and Nagel, 2009).

1.3.2 Learning and experiments: EWA and HSM

The issues with rational learning convey that we must not construct learning models in order to satisfy or explicate some *a priori* given theoretical model of prediction formation (such as RE), since we cannot have one in a reliable fashion. To the contrary, the learning to forecast remains an empirical phenomenon and hence we have to use empirical data to construct and judge learning models. As discussed earlier, experiments with their controlled settings can serve here as an ideal starting point. The best known examples of such a methodological approach come from Game Theory. First, models from evolutionary game theory can be understood as a form of social learning. Gale et al. (1995) applied it to the experimental findings on the Ultimatum Game. Bowles

⁹Guesnerie (2002) openly acknowledges that the Eductive Learning requires ‘two extreme rationality assumptions’, namely Bayesian updating and common agreement on individual rationality.

¹⁰Eductive Learning is interpreted an instantaneous mental process, which proceeds the actual game that unfolds according to the rational solution.

and Gintis (2011) provide an extensive review of the example of evolution of human cooperation, whereas Vriend (2000) discusses social versus individual learning in a market context. Next, experiments on specific games with ‘unintuitive’ Nash Equilibria (such as the above mentioned p-beauty contest and the ultimatum game) sparked a large literature on how the agents can learn selecting strategies (see Rand et al., 2013, for a study of the ultimatum game). Important cases include attraction-based learning: reinforcement learning (Roth and Erev, 1995) and fictitious learning (Fudenberg, 1998), together with their generalization in the form of Experience-Weighted Attraction model (EWA; Camerer and Ho, 1999; Ho et al., 2008).

EWA models found successful applications to various experimental data, which stands as an important lesson of how economic experiments can be used for selecting and tuning theory (see Kocher and Sutter, 2005, for an example and literature overview). In the context of this thesis, the main problem of these models is that they were designed specifically to benchmark games, typically of static nature with a small strategy space. In comparison, firms or financial investors face dynamic problems with continuous action space (like a production choice based on price expectations). Even the most general EWA model cannot be directly adapted to such settings, and therefore these models play a limited role in behavioral finance or macro.

Once we reinterpret ‘strategy’ as a prediction rule instead of a point prediction, however, reinforcement learning can be applied to price-expectations feedback. As mentioned above, the type of market feedback seems to explain the subjects’ choice of forecasting heuristics. This leads to a class of Heuristic Switching Models (HSM; Brock and Hommes, 1997), in which agents switch between simple prediction rules (like adaptive, anchor and adjustment or trend extrapolation expectations) depending on their historical ability to forecast prices. A virtue of HSM is that it explicitly models behavioral heterogeneity, and thus can be easily calibrated to experimental data (Hommes, 2013). Among recent examples, Anufriev et al. (2013) use a two-heuristic HSM to evaluate the differences between the positive and negative feedback treatments of Heemeijer et al. (2009), while Anufriev and Hommes (2012) apply a four-heuristic HSM to the nonlinear asset pricing experiment of Hommes et al. (2005). In both cases HSM outperformed benchmark homogenous models, including RE, in replicating heterogeneous, diversified dynamics of these experiments. Furthermore, the HSM approach proved a successful approximation of dynamics of empirical foreign exchange rates (Dieci and Westerhoff, 2010), house prices (Bolt et al., 2011), macroeconomics (De Grauwe, 2011; Massaro, 2012) and asset prices (Boswijk et al., 2007; Westerhoff and Reitz, 2003).

Nevertheless, HSM is limited in one important aspect. It is natural to expect people

to use and switch between different behavioral rules, which remains the cornerstone of HSM. However, there is an infinite space of forecasting heuristics, while HSM is typically based on an *a priori* selection of a handful of prediction rules. Furthermore, this selection in practice is different in different settings. For instance, Westerhoff and Reitz (2003) consider interaction only between fundamentalists and chartists, whereas Anufriev and Hommes (2012) study a model with four heuristics that do not contain the pure fundamental rule. As a result, HSM cannot fully explain the empirical heterogeneity of price forecasting, and constrains the underlying forecasting heuristics in an unsatisfactory fashion.

1.3.3 Agent-based models of learning

Another approach comes with the so called agent-based models (ABM). For the sake of analytical tractability, economists typically use a notion of representative agents to describe the vast number of real market participants, as if the economy was populated by a handful of characteristic individuals (Hartley, 2002). This does not exclude some degree of an underlying heterogeneity, both with the assumption of RE (Heathcote et al., 2009) or without it (with an example of HSM). Nevertheless, such a representation requires restrictive assumptions on the agents' nature (for example homothetic utility functions; see Kirman, 1992, for a discussion), as well as on the agents' interaction and market structure (Colander et al., 2008). ABMs are based on alternative, bottom-up approach (Delli Gatti et al., 2011): the model keeps track of *every* individual, which is left as an explicitly independent decision and learning routine. Hence, the market aggregate outcomes are computed directly from the local interactions (Delli Gatti et al., 2010; Tesfatsion and Judd, 2006).¹¹

ABM approach became popular in 1990ties (Farmer and Foley, 2009), with seminal works of Arifovic (1995) and Kirman and Vriend (2001). By their complex nature, ABMs offer the perfect testing and modeling ground for explicit asymmetries, non-trivial local interaction, bounded rationality and the emergent aggregate properties of these (LeBaron and Tesfatsion, 2008; Lengnick, 2013). In addition, ABMs allow for a direct evaluation of dynamical properties of complex systems, instead of relying on the more traditional equilibrium approach (Arthur, 2006). ABMs found successful

¹¹Sometimes the term ABM is used to refer to any model in which agents of the same type retain heterogeneous beliefs and decisions, even if it is possible (*e.g.* through the methods of statistical physics) to aggregate the individuals into a set of representative agents (for example, see Hommes, 2013; Lux and Marchesi, 1999). Hence, some authors prefer the name of Agent-based Computational Economics (ACE) to signify a model that does not use representative agents at all (Kirman, 1992). This thesis will *not* follow this practice.

applications in finance (Arifovic, 1995; Egenter et al., 1999; LeBaron, 2001; Lux and Marchesi, 2000), macroeconomics (Assenza et al., 2014b; Dawid, 2006; Deissenberg et al., 2008; Delli Gatti et al., 2011, 2010, 2009) and market organization (Chen, 2012; Kirman and Vriend, 2001; Vriend, 2006). Chapter 3 of this thesis focuses on an ABM application to information networks in financial markets.

There is no free lunch, however. A potential downside of the complexity of the ABM approach is that such models can only be investigated through numerical simulations, *e.g.* Monte Carlo experiments (Metropolis et al., 1953; Metropolis and Ulam, 1949). Therefore, a typical ABM cannot yield an explicit formulation of the emergent equilibrium conditions (like an equilibrium market clearing solution) or the emergent individual behavior (such as the Euler equation in an infinite horizon optimization problem).¹² A further implication is that ABMs often generate dynamics and properties that are *ex ante* unpredictable and can be difficult to interpret *ex post*.¹³ This led many economists to believe that ABMs are unintelligible ‘black-boxes’, a methodological attitude that originates from the problem of ‘wilderness of bounded rationality’ (Conlisk, 1996; Sims, 1980).

The specific challenge for ABMs in this context is that these models remain sensitive to the assumed individual behavior (LeBaron, 2006). A common practice is to select the behavioral skeleton of an ABM in such a way that its aggregate outcomes match the relevant stylized facts (like a distribution of prices or production decisions). This approach, however, may fail the Lucas critique: it is possible that a model has aggregate dynamics similar to a particular data set only ‘by chance’ and hence is useless for evaluating different economies, or the effects of larger fundamental shocks or policy changes. This issue has been widely recognized (LeBaron, 2001) and led to a rising popularity of calibrating ABMs with experimental data (Duffy, 2006) in the hope of providing these models with sound *empirical* micro-foundations (see also Macy and Willer, 2002, for a sociology sciences perspective). Chapter 2 of this thesis is inspired by this research agenda.

Most of the ABM studies are focused directly on the formation of individual actions

¹²For a typical ABM, the relevant rational (Walrasian, RE or Nash) equilibrium is likely to be a part of the set of the model’s fixed points. However, it is impossible to assess uniqueness or derive stability conditions of this set. See Chapter 3 for an exemplary financial market in which introducing a network unexpectedly distorts the stability of the RE solution.

¹³Many economist take as natural that the learned equilibria should have ‘smart’, rational properties (Sims, 1980), whereas ABMs often lead to ‘irrational’ outcomes. For instance, the model of Assenza et al. (2014b) yields an endogenous business cycle without a need of additional fundamental shocks. This may fit better the stylized behavior of the empirical GDP, but also means that the agents within the model do not converge to a ‘perfect’ behavior of any type. As discussed above, there is strong evidence against relevance of the ‘perfect equilibrium’, and we leave this part of ABM criticism for other discussions.

(like trading decisions), thus ABM literature on explicit expectation learning is scarce. An interesting exception comes with the works of Arifovic (1995, 1996), who uses Genetic Algorithms to model social learning of price expectations (see Haupt and Haupt (2004) for technical introduction and Dawid (1996) for discussion and examples of economic applications of Genetic Algorithms). Hommes and Lux (2013) combine this approach with the insight of HSM, namely that economic agents operate on the level of heuristics instead of point predictions. The authors consider a model in which every agent learns independently by optimizing a general forecasting rule with Genetic Algorithms. The authors show that this model beats the RE in explaining aggregate price dynamics and individual behavior in the experiment by Hommes et al. (2007). Chapter 2 of this thesis follows applies this approach to a diversified set of Learning-to-Forecast experiments.

1.4 Thesis outline

The main question of this thesis is how do agents learn to forecast in diversified market environments? This includes the role of price-expectations feedback, heterogeneity, and the relation between forecasting and trading. The thesis consists of three related but independently written chapters, which are built around the themes discussed earlier in this introduction.

Chapter 2 adapts the ABM design by Hommes and Lux (2013) of the Genetic Algorithms (GA) based learning and apply it to various experimental settings.¹⁴ For the sake of robust microfoundations, we follow the insight of Heemeijer et al. (2009) and have our agents use a mixture of adaptive and trend following expectations heuristics. We focus on the model's *empirical fit* to four different Learning-to-Forecast experiments, which we investigate through novel Monte Carlo simulation studies of one-period and 50-period ahead out-of-sample prediction performance. The specific experiments include: linear positive and negative feedback (Heemeijer et al., 2009); linear positive and negative feedback with large unanticipated shocks to the fundamental (Bao et al., 2012); non-linear cobweb economy (Hommes et al., 2007; van de Velden, 2001); and two-period ahead non-linear asset pricing market (Hommes et al., 2005). We show that for all four experiments, our GA model outperforms rational expectations, simple homogenous expectation model and the simple HSM, both at the *aggregate* and the *individual* level. In addition, we show that HSM is a good stylized approximation of the experimental dynamics.

¹⁴This chapter is based on joint work with Mikhail Anufriev and Cars Hommes.

The goal of Chapter 3 is to study the effect of information networks on price stability and individual forecasting coordination in financial markets. We take the GA model in the setting of the two-period ahead non-linear asset pricing market, and add information networks. Next to extrapolating the trend, agents can also learn to trust (or not) the recent mood of their friends. We find that the information networks destabilize the market. Agents remain well coordinated in terms of their price expectations. Nevertheless, they also learn contrarian strategies, trying to outsmart their friends, whose past decisions lag behind the market price cycle. This result confirms and explains the experimental studies on herding/contrarian strategies (Cipriani and Guarino, 2009; Drehmann et al., 2005). Interestingly, we show that the specific network architecture or size does not seem to play a significant role in the emergent market dynamics.

Chapter 4 discusses an experimental study of a simple linear asset pricing market, in which subjects are asked to predict price, directly trade the asset, or do both.¹⁵ Our investigation offers a link between Learning-to-Forecast and Learning-to-Optimize experimental designs for a positive feedback type of economy. Non-fundamental dynamics occur in all of three treatments, with the forecasting treatment being the most stable. Subjects use price trend or asset return extrapolating type of rules and can coordinate on large price oscillations, but can also converge to a non-fundamental price. Regardless of the treatment, large behavioral heterogeneity persists. This confirms that the subject behavior, which was observed in previous Learning-to-Forecast experiments, is robust against task specification. In fact, the Learning-to-Forecast setting, surprisingly, can result in the most stable learning dynamics, whereas Learning-to-Optimize and Mixed treatments are more unstable and may lead to a cycle of repeated bubbles and crashes. Moreover, the mixed treatment indicates that trading consistently with the price expectations is difficult for the subjects, and only a quarter of them properly optimize their behavior. A general conclusion from this chapter is that the learning to optimize seems more difficult for the subjects than the learning to forecast.

¹⁵This chapter is based on joint work with Te Bao and Cars Hommes, available online as CeNDEF Working Paper 14-01.