Bounded rationality and heterogeneous expectations in macroeconomics

Massaro, D.

Chapter 1

Introduction

Individual expectations about future aggregate outcomes are the key feature that distinguishes social sciences and economics from the natural sciences. In every segment of macroeconomics expectations play a central role. Modern macroeconomic theory builds on a dynamic, stochastic, general equilibrium (DSGE) framework based on optimizing behavior. Structural relations derived within this framework and used for policy analysis are obtained from first-order conditions (Euler equations) that characterize optimal private-sector behavior. These conditions explicitly involve private-sector expectations about the future evolution of endogenous variables. Individual economic decisions today thus depend upon expectations about the future state of the global economy. Through these decisions, expectations feed back into the actual realizations of economic variables. Markets are therefore expectations feedback systems and any dynamic economic model depends crucially on its underlying expectations hypothesis.

Since the seminal works of Muth (1961), Lucas (1972) and Sargent (1973), *rational expectations* (RE) have become the leading paradigm on modeling expectations in economics. The idea of rational expectations has two components: first, that the behavior of each individual can be described as the outcome of maximizing an objective function subject to perceived constraints; and second, that the constraints perceived by everybody in the economy are mutually consistent. In such a framework, all agents are the same, expectations are model consistent and coincide on average with realizations, without systematic forecasting errors. Early arguments in

favor of the representative rational agent framework have been provided by leading scholars and date back to Alchian (1950) and Friedman (1953). The general underlying idea is that natural selection driven by realized profits will eventually eliminate non-rational behavior and lead to a market outcome dominated by rational profit maximizing agents. Therefore the economy can be described "*as if*" all agents were perfectly rational.

The RE approach has the important advantage of imposing strong discipline on individual forecasting behavior, minimizing the number of free parameters to explain data. However, RE models rest on the unrealistic assumption of perfect knowledge of the economy. Sargent (1993), for example, notes that "*when implemented numerically or econometrically, rational expectations models impute much more knowledge to the agents within the model (who use the equilibrium probability distributions in evaluating their Euler equations) than is possessed by an econometrician, who faces estimation and inference problems that the agents in the model have somehow solved*". Sargent's argument underscores that, not only the agents have to be endowed with a substantial amount of information in order to form RE, but even if perfect knowledge of the market were available, RE requires extremely strong computing abilities of the agents to solve the model, i.e., to make decisions such that all predictions and beliefs are consistent with the outcome of all agents' choices.

Several tests of RE models, conducted in the last thirty years, have shown that their predictions are often at odds with empirical observations. Much evidence has been collected against the practice of describing human behavior as rational. Conlisk (1996) classifies this evidence as either "direct", concerning rationality tests on individuals cognitive abilities relevant to economic decisions, or "confounded", concerning tests in which the rationality hypothesis is tested jointly with other hypotheses in economic settings. Direct evidence against rationality consists, for example, in showing that individual responses to simple economic decisions typically present systematic errors and psychological biases (see, e.g., Tversky and Kahneman (1974), Grether and Plott (1979), Tversky and Thaler (1990), and Kahneman, Knetsch, and Tversky (1991)). Indirect evidence against rationality has been collected from empirical tests of the pre-

dictions of economic models built under the assumptions of RE. Few examples include households consumption data, which are often at odds with standard life cycle theory (see, e.g., Thaler (1990), Flavin (1993), Carroll (1994), and Shea (1995)), survey data on expectations of inflation and other variables, which commonly reject the unbiasedness and efficiency predictions of RE (see, e.g., Frankel and Froot (1987), De Bondt and Thaler (1990), Ito (1990), and Capistran and Timmermann (2009)), and the anomalous behavior of asset prices (see, e.g., Mehra and Prescott (1985), Fama and French (1988), Lee, Shleifer, and Thaler (1991), and Schiller (2000)).

Bounded Rationality

In recent years, a significant part of economics witnessed a paradigm shift to an alternative, behavioral view, where agents are *boundedly rational* (see, e.g., Conlisk (1996) for a survey). Generally speaking, a boundedly rational agent is modeled as being able to choose what he perceives as the best alternative in a decision making process, but he does not know the exact structure of the economic environment. In standard optimizing theory, agents act as if they perform exhaustive searches over all possible decisions and then pick the best. However, Simon (1955, 1957) emphasizes that individuals are limited in their knowledge and in their computing abilities, and moreover that they face search costs to obtain sophisticated information in order to pursue optimal decision rules. Simon argues that, because of these limitations, bounded rationality with agents using simple satisficing rules of thumb for their decisions under uncertainty, is a more accurate and realistic description of human behavior than perfect rationality with fully optimal decision rules. Kahneman and Tversky (1973) and Tversky and Kahneman (1974) provide evidence from psychology laboratory experiments that, in simple decision problems under uncertainty, individual behavior can be described by simple heuristics which may lead to significant biases. An interesting discussion on the use of simple heuristics as opposed to rational behavior is contained in the Nobel Memorial Lectures in Simon (1979), while a more recent overview on bounded rationality can be found in Kahneman (2003). When predicting future variables, bounded rationality implies that individual agents do not know the true equilibrium

distribution of aggregate variables, therefore ex-ante predictions and ex-post outcomes need not to coincide on average. Hence, a boundedly rational agent is not able to solve for the equilibrium of the expectational feedback system. In contrast, he uses simple heuristics and keeps on updating his strategies as he learns about the economic environment through feedbacks about his past decisions. Researchers adopting RE models often argue that rationality is a useful assumption to describe the equilibrium outcome of this trial-and-error processes. According to this view, the repeated interaction of boundedly rational agents whose beliefs co-evolve with the economic environment in a dynamic feedback system, leads to the same outcome "*as if*" agents were perfectly rational. Convergence to rational behavior has been the topic of investigation of many theoretical papers on bounded rationality. In macroeconomics, much work has been done on *adaptive learning*, see e.g. Sargent (1993) and Evans and Honkapohja (2001) for detailed overviews. Boundedly rational agents do not know the true law of motion of the economy, but instead use time series observations to form expectations based on their own perceived law of motion, trying to learn the model parameters as new observations become available. Much of this literature focused on the stability of RE equilibria and addressed the possibility of agents learning to form rational expectations. In fact, adaptive learning may enforce convergence to RE equilibria but it may also lead to non-RE equilibria, such as the learning equilibria in Bullard (1994). As synthesized by Bullard (1996), "*some rational expectations equilibria are learnable while others are not. Furthermore, convergence will in general depend on all the economic parameters of the system, including policy parameters*". Learnability of RE equilibria depends on the structure of the feedback system between individual expectations and the economic environment, as the signals received from the market might be deceptive to agents trying to acquire rational expectations through learning.

Heterogeneous Expectations

Although adaptive learning has become increasingly popular as an alternative paradigm to model private-sector expectations, most models still assume a representative agent who is learn-

ing about the economy. Kirman (1992, 2006) and Hommes (2006) summarize some of the arguments in support of heterogeneous expectations. One commonly referred to reason is the "no trade" argument, which states that in a world where all agents are rational and it is common knowledge that everyone is rational, there will be no trade. Several no trade theorem have been obtained in the literature (see, e.g., Milgrom and Stockey (1982) and Fudenberg and Tirole (1991)). However, no trade theorems are in sharp contrast with the high daily trading observed in real markets, and this reinforces the idea of heterogeneous expectations. Moreover, heterogeneity in individual expectations has been abundantly documented empirically. For example, Frankel and Froot (1987, 1990), Allen and Taylor (1990) and Taylor and Allen (1992) find that financial experts use different forecasting strategies to predict exchange rates. More recently, Carroll (2003), Mankiw, Reis, and Wolfers (2003), Branch (2004) and Pfajfar and Santoro (2010) provided supporting evidence for heterogeneous beliefs using survey data on inflation expectations, while Hommes, Sonnemans, Tuinstra, and van de Velden (2005), Adam (2007), Pfajfar and Zakelj (2010) and Hommes (2011) find evidence for heterogeneity in learning to forecast laboratory experiments.

In this thesis we depart from standard models by relaxing, as suggested by Sargent (1993), the second component of RE, i.e., mutual consistency of perceptions. In fact, in the models considered in the thesis, we fully maintain the assumption that individual choices are made optimally, given subjective expectations. To be more precise, we assume that agents only have knowledge of their objectives and of the constraints that they face, but they do not have a complete economic model of determination of aggregate variables. Individual decisions are then taken optimally on the basis of subjective expectations of future evolution of endogenous variables. In the light of the theoretical arguments and empirical evidence mentioned above, we assume that individual expectations are heterogeneous. Therefore different agents will, in general, take different decisions when facing the same economic problem as a function of their prediction rules. Prediction rules can differ in terms of sophistication, where the most sophisticated rule corresponds to the RE predictor. Starting from the idea of costly information pro-

cessing, shared with the literature on "rational inattention" (see e.g. Sims (2003)), we assume that the higher the sophistication of a rule, the higher the deliberation cost an agent pays to use it. Instead of considering the fractions of agents employing each rule as fixed and exogenously given, we let them evolve over time as a function of their *fitness*. We thus employ an evolutionary approach where a Darwinian "survival of the fittest" mechanism is at work. A strategy that has performed better according to some measure, e.g. related to its forecasting error, is more likely to be adopted by a higher fraction of agents.[1] In order to model the updating process for fractions we use a discrete choice mechanism (see Manski and McFadden (1981)), introduced in the learning literature by Brock and Hommes (1997) to describe the endogenous selection among heterogeneous expectation rules.

Thesis Outline

This thesis presents applications of a bounded rationality and heterogeneous expectations framework to (New Keynesian) macroeconomic models of inflation and output dynamics. The results of the analysis are articulated in four complementary chapters. Chapters 2 and 3 focus on theoretical monetary policy considerations in the presence of heterogeneous beliefs. Chapter 2 presents a simple macro model with heterogeneous expectations while Chapter 3 provides a micro-foundation of heterogeneous expectations in a New Keynesian setting. Chapters 4 and 5 are devoted to the empirical validation of the heterogeneous expectations framework, using time series data (Chapter 4) as well as experimental data (Chapter 5). A working paper has been extracted from each chapter: Anufriev, Assenza, Hommes, and Massaro (2008) is based on Chapter 2, Massaro (2011) is based on Chapter 3, Assenza, Heemeijer, Hommes, and Massaro (2011) is based on Chapter 4, and Cornea, Hommes, and Massaro (2011) is based on Chapter 5.

Chapter 2 studies inflation dynamics under heterogeneous expectations and investigates the

[1] Empirical evidence that proportion of heterogeneous forecasters evolve over time as a reaction to forecast errors has been provided, among others, by Frankel and Froot (1991), Bloomfield and Hales (2002), Branch (2004), and Hommes (2011), using survey data as well as experimental data.

robustness of monetary policies when agents use simple heuristics to predict future inflation, and they update their beliefs based on past forecasting performance. In this chapter we use a simple frictionless DSGE framework, which allows us to obtain analytical results about global dynamics in the model. The results of the analysis show how macroeconomic stability depends on the composition of the set of forecasting strategies and on the policy reaction coefficient of a Taylor-type interest rate rule. A result of particular interest is that the Taylor principle is no longer sufficient to guarantee uniqueness and global stability of the RE equilibrium, as multiple equilibria and non-rational beliefs may survive evolutionary competition.

Chapter 3 extends the analysis performed in Chapter 2 in two important respects. First, it derives a New-Keynesian framework consistent with heterogeneous expectations starting from the *micro-foundations* of the model. We then use this framework characterized by monopolistic competition, nominal rigidities and heterogeneous beliefs for the analysis of monetary policy and macroeconomic stability. Second, it introduces the sophisticated RE predictor in the set of forecasting strategies available to the agents in the economy. Due to the increased level of complexity, this chapter uses computational methods to derive policy implications for a monetary authority aiming at stabilizing the dynamic feedback system where macroeconomic variables and heterogeneous expectations co-evolve over time. A salient result is that policy attempts to achieve determinacy under RE may destabilize the economy even when only a small fraction of boundedly rational agents are present in the system.

Chapter 4 focuses on the empirical validation of the heterogeneous expectation framework. In this chapter we estimate a behavioral model of inflation dynamics with monopolistic competition, staggered price setting, and heterogeneous firms using U.S. time series data. In our stylized framework there are two types of price-setters, fundamentalists and naive. The estimation results show statistically significant behavioral heterogeneity and substantial time variation in the weights of different forecasting rules.

Chapter 5 investigates the individual expectation formation process in laboratory experiments with human subjects within a New Keynesian framework. Our data show that individuals

use simple heuristics to forecast aggregate macroeconomic variables within the experimental economies, and that individual learning takes the form of switching towards the best performing heuristics. We then use a simple model of individual learning with a performance-based evolutionary selection among forecasting strategies to explain experimental outcomes under different monetary policy regimes.

Each chapter is self-contained, with its own introduction, conclusion, notes and appendices as needed. For this reason, each chapter can be read independently from the others. A common bibliography is collected at the end of the thesis.