Learning Processes of Dynamically Optimizing Heterogeneous Agents

An Examination of Least Squares Learning Approaches by the Example of a Basic Consumption Model

Inauguraldissertation zur Erlangung des akademischen Grades
eines Doktors der Wirtschaftswissenschaften durch die
Wirtschaftswissenschaftliche Fakultät der
Westfälischen Wilhelms-Universität Münster
vorgelegt von

Dipl.-Volkswirt Philipp Gerlach

aus Münster

Erstgutachter:	Prof. Dr. Gustav Dieckheuer
Zweitgutachter:	Prof. Dr. Wolfgang Ströbele
Dekan:	Prof. Dr. Stefan Klein
Tag der mündlichen Prüfung:	24. November 2009

Preface

Most of the dynamic general equilibrium models in current macroeconomic theory assume the existence of a representative agent whose expectations are rational. While these assumptions facilitate the models' technical tractability, it is clear that the behaviour they describe is to a great extent simplified and idealized. Currently in the scientific landscape considerable efforts are made to relax these assumptions. In particular the rational expectations assumptions has been the subject of critical review, which has led to the emergence of learning theory, among others. While learning is now being extensively applied in macroeconomics, usually in the form of adaptive learning, the models still often make use of a representative agent. However, it is a distinct possibility that this assumption is not a mere technicality but can have significant consequences to the outcome of the learning processes in a model.

This work examines how the assumption of a representative agent can lead to spurious results in the context of adaptive learning and how adaptive learning can be applied when heterogeneous agents are used by the example of a basic consumption model with learning agents. It has been written at the Westfälische Wilhelms-Universität Münster, in part during my employment as a scientific assistant at the Institut für Industriewirtschaftliche Forschung, and has been accepted as a doctoral thesis in 2009. I would like to thank my thesis advisor Prof. Dr. Gustav Dieckheuer for his critical support, his patience and the scientific liberty which allowed me to accomplish this work. Just as well I want to thank my second advisor Prof. Dr. Wolfgang Ströbele for valuable hints and suggestions.

A work like a doctoral thesis is virtually impossible without the understanding support of one's friends and family. Among these I can only name but a few. First of all I am highly indebted to Dr. Marcel Mlakar for many valuable discussions, critical comments and his steady encouragement. In the same way I would like to thank my former colleagues at the Institut für Industriewirtschaftliche Forschung for their support, the inspiring collaboration and the good times we shared. Many thanks go to Nadine Piolot for meticulously proofreading the manuscript in spite of having already very much on her plate, and to Dr. André Pohlkamp for our mutual motivation to compensate for the scientific work with sports and leisure once in a while. My mother and my sister Johanna have constantly encouraged me and I could always count on their moral support over the years, which I would like to thank them for. But most of all I want to thank Kathrin Meier-Kolthoff for her unfailing motivation, for enduring occasional difficult times of misery and bad temper on my side, and for the way I could always rely on her understanding, sympathy and affection.

Münster, April 2010 Philipp Gerlach

Contents

Abbreviations

ALM	Actual law of motion
AR(1)	Auto-regressive of order 1
EE learning	Euler equation learning
EEE learning	Extended Euler equation learning
IH learning	Infinite horizon learning
MSV	Minimal state variable
ODE	Ordinary differential equation
OLS	Ordinary least squares
PLM	Perceived law of motion
REE	Rational expectations equilibrium
RLS	Recursive least squares
RPE	Restricted perceptions equilibrium
SCE	Self confirming equilibrium
SRA	Stochastic recursive algorithm
VAR	Vector auto-regressive

Symbols

Chapter 2

$E_t x$	expectation of x in period t
$T(\cdot)$	mapping from PLM to ALM
N	index of the period of time in which beliefs are held (in early approaches to adaptive learning)
R_t	moment matrix
a, b, d	parameters of the households' PLM
c_t	aggregated private consumption
c_t^i	private consumption of household i
g_t	government expenditure
n	number of households
s, t	time index
u_t	disturbance within the AR(1) process for g
v_t	unrelated stochastic term ("sunspot")
y_t	aggregated income / general endogenous variable
z_t	general vector of a constant and of determining variables
ϕ	vector of estimated parameters of the PLM
Ψ	vector of actual parameters of the ALM
Ω_t	information set available in period t
α, δ	parameters of the households' private consumption function
β	parameter of the AR(1) process for g
γ	gain
η_t	general disturbance
τ	time variable for associated ODE
ω	general vector of determining variables

Chapter 3

X	upper case letters denote the absolute value of a variable
x	lower case letters denote the natural logarithm of the deviation from the non-stochastic steady state
$\bar{X}$	upper case letters with an upper bar denotes the non-stochastic steady state
x^i	a superscripted index *i* indicates a variable of an individual agent *i*
$\tilde{x}$	a tilde over a variable denotes the average of an individual variable of all agents
$\hat{E}_t^i$	expectation of agent *i* in period *t* (possibly non-rational)
E_t	rational expectation in period *t*
$T(\cdot)$	mapping from PLM to ALM
$U(\cdot)$	instantaneous utility function
B	wealth
C	consumption
M	parameter of income process
R	gross one-period rate of return
Y	income
e, k, l, m, p	parameters of the PLM (beliefs)
i, j	agent index
n	number of households
s, t	time index
β	temporal discount factor
ε	extraneous sunspot variable
μ	parameter of sunspot process
v	disturbance of income process
π	disturbance of sunspot variable
ρ	parameter of income process
σ	intertemporal elasticity of substitution at the steady state
τ	time variable for associated ODE
υ	error term

1. Introduction

Many macroeconomic models used today take the form of micro-founded dynamic general equilibrium models. Models of this class aim to derive relations between macroeconomic aggregate variables from the microeconomic behaviour of the model's agents. Employing this micro-founded approach, such models contrast with models in the tradition of the Keynes-Hicks IS-LM framework, which largely posit functional relationships between aggregate variables directly.[1] Among the plethora of macroeconomic models, purely macroeconomic IS-LM-style models on the one hand and micro-founded dynamic general equilibrium models on the other hand seem to constitute two important poles. But it is difficult to label those poles properly, since the taxonomy of macroeconomic models is far from being clear. For example, the attribute "general equilibrium" is often used to denote micro-founded models, while others argue that IS-LM-style models which cover all markets can also be of a "general equilibrium" kind in contrast to partial models.[2] Here, I will use the terms "micro-founded" versus "purely macroeconomic" to designate the two approaches.[3]

While the underlying concepts of the micro-founded general-equilibrium approach dates back to ideas which have been advanced prominently by Walras as well as Arrow and Debreu, models of the purely macroeconomic kind have dominated macroeconomics up to the end of the 1970s. Then, the micro-foundation approach became popular when it was employed in the real business cycle models of the 1980s,[4] and today it is the foundation of many workhorse models in macroeconomics, such as the New-Keynesian general equilibrium models.[5] It should be mentioned that the current predominance of

---

1 An overview of the IS-LM framework with various extensions can be found in Dieckheuer (2003).

2 See De Vroey (2006). In this paper, the author instead proposes the main methodological distinction to be drawn between a Marshallian approach and a Walrasian approach.

3 In theory, a micro-founded general equilibrium model does not have to be dynamic as well, but in practice it usually is. So when in the following the term "micro-founded" is used this may also imply that the model in question is a dynamic model.

4 See, for example, Plosser (1989).

5 For an overview of the history of macroeconomic models, see Blanchard (2000) and Woodford (1999).

micro-founded models is nevertheless still the subject of a controversial methodological discussion.[6]

The distinctive feature of micro-founded models is the fact that the interrelations between aggregate macroeconomic variables do not have to be posited but can be traced back to the individuals' behaviour which causes the interrelation in question. Still, this approach is not necessarily less dependent on assumptions than the purely macroeconomic models. Instead of assuming a functional form and parameters on a macroeconomic level it is now necessary to make assumptions about microeconomic functions and their parameters. Since virtually all features of micro-founded models are derived from these functions the model can be very sensitive to their specification. Moreover, the way how a derived macroeconomic interrelation depends on the specifications of the underlying functions can be much less obvious in such a model than in a purely macroeconomic model. For that reason it is important to exercise great care in defining the microeconomic foundation of a general equilibrium model. This includes not only the specification of the microeconomic functions but also the assumptions and restrictions which are made implicitly in the process of building and solving the model, for example by employing particular solving techniques. Among the decisions which must be taken when building a micro-founded dynamic general equilibrium model, two are of special interest in the context of learning: How to aggregate the individual agents' decisions to macroeconomic variables and the specification of expectations.

I will turn to expectations first. In a micro-founded model, the individual agent's behaviour is described by a set of given microeconomic functions, which are typically the result of the agent solving a problem of dynamic optimization subject to constraints. Because of the dynamic nature of the optimization problem the microeconomic behavioural functions take into account not only the observable current state variables but also expectations of future state variables. Thus, expectations constitute a vital element in such models with dynamically optimizing agents. It is therefore not surprising that the rise of micro-founded models coincided with the "rational

6 See for example Blinder (1987), section VIII, Hoover (2008), Kirman (1992) or Wren-Lewis (2006).

expectations revolution"[7], since the rational expectations hypothesis provides an elegant way to reduce and solve the agents' dynamic optimization problems. The rational expectations hypothesis has consequently become the standard assumption in most dynamic general equilibrium models.

Further, the very idea of using micro-foundations in a macroeconomic model implies the question of how to bridge the gap between the microeconomic and the macroeconomic level, or more precisely, how to derive the behaviour of aggregate variables from the decisions of individual agents. A straightforward way to deal with this problem is to assume the existence of a representative agent within each class of agents, for example a representative consumer or a representative firm, which embodies the "typical" or "average" behaviour of the whole class. The assumption here is that it is possible to specify an agent's microeconomic behavioural functions in such a way that his choices will always coincide with the aggregate choices of the whole class of agents. If such an agent can be constructed, then it is possible to scale up the behaviour of this representative agent to the aggregate.

While there are also models with heterogeneous agents, the great majority of micro-founded models use the representative agent assumption. Due to its easy application, the technique has become the standard way to solve the aggregation problem, despite the fact that from a methodological point of view it can be problematic. Indeed, this assumption can be considered the single most controversial ingredient of micro-founded models and it fuels much of the debate about their legitimacy.[8] Much of the debate focuses on the conditions under which a representative agent can be constructed at all. The details of this discussion are beyond the scope of this work and will be left aside here. But even if a representative agent could be constructed (or if it is simply assumed that this is the case), using him in the solution of a model can still be problematic. In particular, a result of a model may depend on the application of the representative agent assumption, or another element of the model or a solution method may be only applicable in conjunction with the assumption. This does not necessarily mean that the

7 For a historical overview, see section 2.1

8 An overview of theoretical objections to the representative agent hypothesis is given in Kirman (1992).

application of the representative agent assumption is always inadequate per se. It is possible that in some models the assumption is a valid simplification which does not change any outcome of the model and that all other elements of the model are independent from it, but that cannot be simply taken for granted. Instead, when using the representative agent assumption, it should be verified that the model would also work without it and that it does not change any results. The representative agent assumption is nothing more than a tool to simplify a complex heterogeneous world, so the results of a model should be as independent from its application as possible.

As a contribution to this field of research, this work will examine how the results of a basic dynamic optimization model change when rational expectations are replaced by learning and when the representative agent is replaced by individual agents which are heterogeneous to some degree. In order to keep the analysis from getting too complex, it is done here on a partial model which focuses on the consumers' optimization problem, treating the consumers' incomes as exogenous processes. Because of this restriction the model is rather a microeconomic one in contrast to a fully-fledged macroeconomic model, which would have to endogenize the income by including the production side of the economy. However, already the partial model used here highlights some effects which arise due to the introduction of learning and of heterogeneous agents. The question whether the results of this work hold in the same way for more complex macroeconomic models is left for future research.

I will argue that the introduction of learning makes most sense if individual agents are considered. When adaptive learning is included in a micro founded model with dynamically optimizing agents, there are two known ways which are usually called "Euler Equation Learning" versus "Infinite Horizon Learning".[9] The question of which of these methods is the adequate one is subject of an ongoing debate.[10] I will demonstrate that the infinite horizon learning is more robust with respect to the heterogeneity of the agents. Moreover, I will argue that some of the confusions about the equivalence of the two methods are due to the fact that some solving techniques, which depend on the representative agent assumption and rational expectations, are no

9 These terms will be defined below in section 3.2 and 3.3.

10 See Homkapohja, Mitra and Evans (2003) and Preston (2005), p. 94 ff.

longer suitable when learning is included.[11] Yet these techniques have become so familiar that their application is rarely subjected to critical review. Finally, I suggest an extension of the Euler equation learning method and demonstrate that this extended method allows individual non-representative agents to learn the rational expectations solution under straightforward assumptions.

In the next chapter, first of all the concept of rational expectations and its properties are shortly reviewed. Then learning is introduced as a dynamic form of expectation formation, with a special focus on recursive least squares learning as the most generic learning method. Constant gain learning will be shortly discussed as an interesting alternative. In chapter 3 a basic consumption model will be examined. It serves as a workhorse example by which first the usual solution technique will be demonstrated, which employs the representative agent assumption and the rational expectations assumption. This solution will subsequently be compared to the situation with heterogeneous agents. In section 3.2 Euler equation learning will be demonstrated and critically reviewed, with a special focus on the impact of the representative agent hypothesis. In this context it will be examined how the solution technique, which is made possible by the representative agent assumption, can be misleading when learning is introduced. Infinite horizon learning will be demonstrated in section 3.3 as a more robust alternative which does not rely on the representative agent assumption. In section 3.4 I suggest an extension to Euler equation learning as a less demanding alternative to infinite horizon learning. Chapter 4 summarizes the discussion and considers some possible areas for future research.

---

[11] This discussion is inspired by the work of Bruce Preston, who has proposed infinite horizon learning as an alternative method. While he puts forward some arguments in favour of infinite horizon learning, the methodological relationship to the representative agent assumption was not pointed out explicitly yet. See for example Preston (2005) or Preston (2005a).

2. Rational Expectations, Bounded Rationality and Learning

2.1. Historical Overview: The Way to Rational Expectations

In many cases an agent needs to consider his future situation and the future economic environment when making economic decisions. Expectations have therefore played a significant role in economic thinking from the very beginning.[12] Yet, in static or comparative static models there is no need to form expectations, so instead of variables for expected values their actual outcomes are usually used. It was therefore not before the advent of dynamic economic models that expectations received more attention and were formally included in models.

The easiest way to handle expectations is to consider them as static. In this case the value of a variable expected for the next period is the same as the actual value:

(1) $$E_t x_{t+1} = x_t \,.$$

Static expectations can either be modelled explicitly as static using a formula like (1), or they are not mentioned at all, thereby often implicitly assumed as static. The dynamic version of the standard Keynes-Hicks IS-LM model, in which planned consumption is a function of the last period's income or production depends on last period's demand, is an example of the latter.

An often used alternative to static expectations are adaptive expectations. These generally take the form of

(2) $$E_t x_{t+1} = E_{t-1} x_t + \lambda\left(x_t - E_{t-1} x_t\right) \qquad \text{with } 0 < \lambda < 1 \,.$$

Here, agents update their expectations with the expectation error from the last forecast they made, but only to a certain degree given by λ. The concept of adaptive expectations was first mentioned in Fisher (1930), but has become popular since Phillip

[12] Evans and Honkappohja (2001) cite an anecdote from Aristotle dating back to about 600 B.C. See Evans and Hokapohja (2001), p. 6.

Cagan included it in his standard model of hyperinflation in 1956.[13] With $\lambda < 1$ this approach smoothes expectations whenever actual values change, thereby contributing to the stability of a model. But adaptive expectations do not necessarily, as it is sometimes stated,[14] represent a higher degree of learning by the economic agents than in the case of static expectations. In fact, the higher λ, the more an agent revises his former expectation in the light of new information. So one could argue that the highest degree of learning would occur with $\lambda \to 1$, which is again the case of static expectations.[15]

Although in both approaches presented above agents do learn from past expectational errors up to a certain degree, this learning algorithm is rather crude and can result in systematically wrong expectations in many situations. This is at odds with the notion of rational behaviour: A rational agent would recognize that he repeatedly formed wrong expectations and would try to adapt his algorithm accordingly. Moreover, when doing so he would try to use all information available, including knowledge about economic theory and statistics. This is the basic idea behind rational expectations, which is attributed to Muth (1961).[16] However, the notion of agents using the relevant economic theory when making forecasts of uncertain variables has been mentioned before, one example is an article by Tinbergen from 1932.[17] Nevertheless Tinbergen's approach did not receive much attention at the time.

While Muth's paper was already widely admired at the time, the so called 'rational expectations revolution' arrived only in 1972, when Robert Lucas started applying Muth's concept (which in its original form was a microeconomic one and restricted to a

---

13 Fisher used a distributed lag model to find a statistically significant relationship between past inflation and nominal interest rates, whereas he considered the latter closely related to expected inflation. Although he used linear decreasing instead of geometrically decreasing weights for past values, this approach closely resembles adaptive expectations. See Dimand (1999), p. 747, Fisher (1930), part IV, chapter XIX, §6 and Cagan (1956), p. 37.

14 For example, see Felderer and Homburg (1993), p. 260.

15 However, the question of which learning algorithm would yield more accurate expectations depends on the model considered and cannot be decided generally.

16 See Muth (1961).

17 See Tinbergen (1932), p. 1972. In his article Tinbergen explores the possibility of expectations being 'reasonable' ('vernünftig' in the original German version of the text) and suggests replacing them by 'deductions according to economic theory' ('wirtschaftstheoretische Deduktion'), which captures the basic idea of rational expectations. For a detailed comparison between the model of his article and the one used by Muth (1961) see Keuzenkamp (1991). The term 'rational expectations' was also used by Hurwicz in 1946. See Sargent (1993), p. 6, note 1.

single market) to macroeconomics and the analysis of economic policy.[18] With this he initiated a process of profound rethinking about the way macroeconomic models should be built. Today the great majority of macroeconomic models incorporate the rational expectations hypothesis or one of its descendants.

2.2. Basic Concept of Rational Expectations

The concept of rational expectations is based on the assumption that in the long run rational agents will not form expectations which are systematically wrong. “Systematically wrong” here means that expectation errors would be related in a systematic way to observable influences, such as exogenous variables, lagged endogenous variables or lagged error terms. The existence of rational expectations can be justified by the assumption that agents compete with each other. If expectations exhibit errors which could be explained by observable influences, some agents could use this information to improve their expectations and so achieve an advantage over their competitors. In the long run this process of selection would lead to the elimination of all inferior expectations, leaving only expectations as rational as possible.[19]

In his seminal article, Muth describes his idea of rational expectations (by the example of the expectations of firms) as follows:

> “… expectations of firms (or, more generally, the subjective probability distribution of outcomes) tend to be distributed, for the same information set, about the prediction of the theory (or the “objective” probability distribution of outcomes).”[20]

Muth stresses that this does neither necessarily mean that expectations must be identical for all firms, nor that the firms use the “true” or objective model when forming their expectations. However, for simplification, rational expectations are often modelled as homogeneous and based on the objective model. This has become standard practice up

[18] For a short summary on Muth’s contribution see Brannon (2006). An overview of Lucas’ impact on macroeconomic research gives Sargent (1996).

[19] In that way competition can ensure that information is optimally used, as Muth demands: “Information is scarce, and the economic system generally does not waste it.” Muth (1961), p. 316.

[20] Muth (1961), p. 316.

to the point that today the term “Rational Expectations” usually means indeed that the agents’ expectations are based on the objective model.

An alternative way of characterizing rational expectations is given by Sargent (1993).[21] In this view rational expectations are an equilibrium concept which is characterized by (a) a rational choice (i.e. maximizing an objective function) by agents subject to their perceived constraints, and (b) mutual consistency of all perceived constraints in the system. In other words, this characterization can be described as follows: (a) Every agent forms beliefs about the world he lives in, and then makes a rational decision based on these beliefs. (b) A rational expectations equilibrium is reached if all beliefs are mutually consistent, including beliefs about beliefs of others (higher order beliefs), so that the observable information generated by the model does not contradict the beliefs of the agents.

I will now have a closer look at the role of rational expectations in a dynamic model in order to define a rational expectations equilibrium. Expectations in a dynamic model have an impact on the state of the modelled economy. If one defines expectations in a model as rational in the definition of Muth, one demands that expectations do not differ systematically from the resulting state of the economy, or more precisely: One demands that the agents do not perceive any contradiction between their expectations and the resulting states. So, the outcome under the impact of the agents’ expectations must “match” these very expectations. In this way expectations are part of a self-referential process.[22]

One possible and straightforward way to satisfy the condition that agents perceive the resulting states of the economy as indistinguishable from their expectations is that agents form their expectations as conditional expectations based on the functions and parameters of the true model (including the functions governing expectation formation). This approach leads to an equilibrium in the above mentioned self-referential process which is called *Rational Expectations Equilibrium* (REE). Again, this approach is not the only way to satisfy the condition. It is also possible that agents employ a

21 See Sargent (1993), p. 6 ff.

22 See Branch (2006), p. 135 f.

misspecified model to form their expectations, but nevertheless within their model they cannot detect any inconsistencies between the data generated by the true model and their misspecified model. This possibility will be examined in more detail in section 2.5.2.

The concept of an REE described above will be demonstrated by means of an example: Consider an economy with *n* households and income *y*, consumption *c* and government expenditure *g* given by

$$(3)\quad y_t = c_t + g_t$$

$$(4)\quad c_t = \sum_{i=1}^{n} c_t^i$$

$$(5)\quad c_t^i = \frac{1}{n}(\alpha y_{t-1} + \delta E_{t-1} y_t), \quad \delta \neq 1$$

$$(6)\quad g_t = \beta g_{t-1} + u_t, \quad |\beta| < 1, \quad E_{t-1}(u_t) = 0$$

where u is an i.i.d. disturbance with zero means. In (5) households smooth consumption over the last period's income and the expected income for the current period. Government expenditure follows an AR(1) process. β is bounded in order to rule out explosive time paths for *g*. For simplicity it is assumed that the parameter α, δ in (5) are the same for all households. Then the reduced form of this model is

$$(7)\quad y_t = \alpha y_{t-1} + \beta g_{t-1} + \delta E_{t-1} y_t + u_t .$$

Rational expectations are often denoted in a form like

$$(8)\quad E_t y_{t+1} = E(y_{t+1} | \Omega_t)$$

with Ω_t being the information set available at time *t* and $E(\cdot)$ being the mathematical expectation conditional on Ω_t. This notation is rather general and does not yet explain how the expectation is actually formed. For this it must be combined with the agent's view about how y_t is determined in the model.

Suppose the households believe that y can be described by the linear equation

$$(9) \quad y_t = a y_{t-1} + b g_{t-1} + u_t .$$

Suppose again that this belief is representative for all households. Such an equation, which characterizes the agents' perception of the economy, is called the *perceived law of motion* (PLM).[23] Then rational expectations imply

$$(10) \quad E_{t-1} y_t = E(y_t | \Omega_{t-1}) = a y_{t-1} + b g_{t-1} .$$

Inserting (10) into (7) yields the *actual law of motion* (ALM), which describes how the economy evolves if forecasts are made according to the PLM:

$$(11) \quad y_t = (\alpha + \delta a) y_{t-1} + (\beta + \delta b) g_{t-1} + u_t .$$

The PLM (9) gives rational expectations only if its coefficients *a*, *b* match the respective coefficients of (11), which implies that in a REE

$$(12) \quad a = \frac{\alpha}{1-\delta}, \quad b = \frac{\beta}{1-\delta} .$$

This approach for solving a rational expectations model is called *method of undetermined coefficients*. It consists of assuming a solution, inserting it into the reduced form of the model and then verifying it by matching the appropriate coefficients.[24] For a slightly different, but related way of solving for the REE look again

[23] See for example Sargent (1993) or Evans and Hokapohja (2001), passim.

[24] See for example Lucas (1975) p. 1118. Although it is often done for clarity, in general it is not necessary for this approach to specify the PLM and ALM explicitly. A very general exposition of the procedure which can be applied to most linear models is given by McCallum (1998). Aoki and Canzoneri mention that it can be difficult to guess the right solution beforehand and propose an alternative procedure. See Aoki and Canzoneri (1979), p. 64. Another well known algorithm for a common subclass of linear expectations models is given by Blanchard and Kahn (1980).

at the self-referential properties of the rational expectations model. The PLM with its coefficients a and b (the agents' beliefs) determines the ALM of y (the outcome). Thus we get a mapping from the coefficients of the PLM to the coefficients of the ALM. This mapping is usually called a *T-map* which plays a crucial role in learning algorithms, as will be discussed in section 2.4.3. In the example it takes the form

$$(13)\quad T\begin{pmatrix} a \\ b \end{pmatrix} = \begin{pmatrix} \alpha + \delta a \\ \beta + \delta b \end{pmatrix}.$$

The REE is then given by a fixed point of T, which again leads to the values of (12). This approach underlines what Sargent calls "consistency of beliefs" in an REE: If every one of the n households believes that the income evolves according to the PLM with coefficients given by (12), then their aggregated consumption will be such that the resulting ALM of y confirms their beliefs.

2.3. Multiple Equilibria in Rational Expectations Models

Models with rational expectations can have multiple rational expectations equilibria. For linear rational expectation models this is possible if current values of the model are determined by expectations of future ones.[25] For a demonstration I will change the model of the previous section by setting government expenditure as a constant plus a disturbance and letting consumption depend on the expected income of the current period and the period after:[26]

$$(14)\quad c_t = \alpha E_{t-1} y_t + \delta E_{t-1} y_{t+1}, \quad \delta \neq 0, \quad \delta + \alpha \neq 1$$

$$(15)\quad g_t = \overline{g} + u_t, \quad E_{t-1}(u_t) = 0\,.$$

[25] See Evans and Honkapohja (2001), p. 173. This dependence on expected future variables is called *forward-looking*.

[26] Consumption is here immediately given as aggregated. Setting g as constant simplifies the treatment. For a discussion of systems with an exogenous process like the one for g in the model before, see Evans and Honkapohja (2001), pp. 178 ff.

The reduced form of the model is[27]

$$(16)\quad y_t = \overline{g} + \alpha E_{t-1} y_t + \delta E_{t-1} y_{t+1} + u_t .$$

Under rational expectations this system has multiple solutions. Consider first the stochastic steady-state solution. Using the method of undetermined coefficients again, assume a solution of the form

$$(17)\quad y_t = a + u_t .$$

Forming rational expectations and inserting these into (16) gives the appropriate term for *a*, so the REE is

$$(18)\quad y_t = \frac{\overline{g}}{1-\alpha-\delta} + u_t .$$

However, other solutions are possible. The agents' PLM could include extraneous variables, namely the disturbance *u* in $t-1$, or an additional unrelated stochastic term v_t with $E_{t-1}(v_t) = 0$, or both. Then the set of solutions is given by

$$(19)\quad y_t = \frac{1}{\delta}\overline{g} + \frac{1-\alpha}{\delta} y_{t-1} + b u_{t-1} + d v_{t-1} + u_t$$

for any *b* and *d*.[28]

The existence of multiple solutions to some rational expectations models (or even an infinity of solutions like in the model above) has been considered by some early authors as a weakness of the rational expectations hypothesis.[29] But it can be shown that the

---

[27] This system is similar to the one discussed in Evans and Honkapohja (2001), pp. 174 ff. For details about obtaining the different solutions and a proof of completeness, see there or Evans (1985), pp. 1218 f.

[28] The process (19) is an ARMA (1,1) process which describes the full set of solutions, including the stochastic steady state for appropriate values of b, c and the starting condition y_{-1}. See Evans and Honkapohja (2001), pp. 176 f.

[29] See Burmeister, Flood and Garber (1983) as well as McCallum (1983), pp. 139 f.

existence of multiple solutions is a common feature of dynamic models and not due to the application of rational expectations.[30] In the face of multiple solutions the question arises which solution(s) economists should focus on, and if maybe some solutions can be ruled out on the grounds of further considerations. I will first look at the typology of possible solutions to rational expectations and then discuss different procedures to classify solutions.

The standard classification distinguishes between a *market fundamentals* solution and a set of *market fundamentals plus bubble* solutions. It can be shown that other proposed types of solution are in fact included in these two basic types.[31] The market fundamentals solution is defined as a solution which is based only on variables considered as market fundamentals. The set of market fundamentals is specific for each model, the intuition is that it consists of those variables which have a direct[32] effect on the state of the model according to its reduced form. In the model above equation (18) represents the market fundamentals solution. In contrast, a bubble is an additional component to the market fundamentals solution which is part of the solution only because it is expected to be so. Thus bubbles represent a kind of self-fulfilling prophecy. These solutions are given above by (19). Note that bubble components do not need to be explosive.

A subclass of bubble solutions are bubbles driven by variables which are otherwise completely unrelated to the rest of the system. In the example above these are given by (19) for $d \neq 0$. Solutions of this type are called *sunspot solutions*. [33] The central idea is that although expectations in a model may be perfectly rational, it is still possible that

---

[30] See McCallum (1983), p. 141. Moreover, he argues that multiple equilibria also exist in equilibrium models with maximizing agents (microfoundation). Under some conditions, explosive solutions can be ruled out in such models, but nevertheless multiple equilibria remain possible.

[31] Other types of solutions (which evolved partly due to specific solution algorithms) are the *forward-backward* solution by Blanchard and the *spurious indicator* solution by Taylor (1977). See Burmeister, Flood and Garber (1983).

[32] As opposed to an indirect effect only via expectations.

[33] These possibilities were investigated initially by Cass and Shell. They used the term „sunspots“ as a synonym for events far away and without actual influence. The word was chosen as an allusion to the theories of the 19th century economist William Stanley Jevons, who claimed to have found a connection between business cycles and the occurrence of sunspots. See Cass and Shell (1983). For a survey of the existence and properties of sunspot equilibria in various types of models, see Chiappori and Guesnerie (1991).

some extraneous variable influences the dynamics of the model. If some agents in a model include this variable into their expectations and so let their decisions be influenced by this variable, then the variable has in fact an impact on the behaviour of the model. It can therefore be rational for each agent to include it in his expectations. In this way the extraneous variable serves as a common signal which causes and "coordinates" self-fulfilling expectations in the economy.

Several criteria have been proposed to select economically relevant solutions from the set of possible solutions.[34] The non-bubble-solution in particular is often of special interest. Taylor (1977) suggested focussing on the solution in which the variable of interest exhibits the lowest unconditional variance (*minimum-variance criterion*).[35] But this criterion does not always point to a unique solution. It is for example possible that there are several variables of interest, each having their minimal variance in different solutions.[36] Another criterion is the widely used *stability criterion* of Blanchard and Kahn (1980). Originally, the procedure described there was apparently not intended as a selection criterion among various solutions. Instead it is a solution algorithm for linear rational expectations models which also yields conditions for existence and uniqueness of the solution. In the construction of the algorithm it is required that expectations do not follow an explosive path, so one has to rule out explosive solutions before applying the algorithm to a model. However, in an economic analysis it is often the aim to determine under which conditions solutions become explosive. As a selection criterion the procedure is therefore not suitable. In addition, there are models with multiple solutions which are each stable, and in this case the criterion is not capable of selecting one.

As an alternative McCallum proposes the *minimal state variable* (MSV) criterion.[37] This criterion closely mimics the intuition behind the market fundamentals solution and

[34] The following is a summary of the discussion in McCallum (1999), pp. 625 ff.

[35] See Taylor (1977), pp. 1383 f. However, it is admitted in this article that this criterion might not be reasonable.

[36] McCallum adds that it is also not clear why an agent should choose that kind of solution rather than another. See McCallum (1999), p. 625. However, this argument is valid as well for Blanchard and Kahn's stability criterion and for McCallum's suggested MSV-criterion.

[37] The construction of the MSV-solution and the procedure to select it was first proposed in McCallum (1983) and further developed in McCallum (1999).

gives a formal procedure to identify it. It consists of several conditions. First, a MSV solution for a linear model is a linear function of no more than a minimal set of state variables, i.e. exogenous or predetermined variables which (directly) determine endogenous variables. “Minimal” here means that

> “... it is impossible to delete … any single variable, or a group of variables, while continuing to obtain a solution valid for all admissionable parameter values.”[38]

There can still be more than one solution which satisfies this condition. To select among these, a second condition[39] is that the MSV solution is still based on a minimal set of state variables for all special cases of the model’s parameter values. That means in particular that if a certain parameter value (e.g. zero) in the model has the effect that its associated state variable does not have an influence on the endogenous variables any more, then this state variable must also vanish from the MSV solution for this special parameter value. Thus, the second condition leaves a single solution, which by construction is the market fundamentals solution. This solution will be referred to as the MSV solution in the following.

Because of its bubble-free nature the MSV solution is of special interest in an economic analysis. However, selecting this solution is just a deliberate choice made by the economist analysing a model. It is still not sure whether this is the solution which agents would choose in a dynamic model. Agents will agree on a particular solution if a process of coordination between agents exists with this solution as its limit point. Within this process one no longer regards the agents’ expectation formation algorithm itself as static. One rather allows expectation formation to change dynamically and one allows agents to adapt their expectation formation to their experience. In this way expectations evolve with the progression of the model. This dynamic view of expectations is the subject of learning theory.

---

38 McCallum (1983), p. 145. There may be several possibilities to define a minimal set of state variables, which nevertheless point to the same MSV solution. See McCallum (1999), p. 626.

39 This condition is sometimes referred to as „subsidiary principle“. See Evans (1986), p. 147.

2.4. Learning

2.4.1. Learning as a Dynamic View of Expectations Formation

One important point of criticism concerning the assumption of rational expectations is that it implies that agents use the true model's structure and parameters to form their expectations. It follows that the agents know a lot more about the economy they live in than the economist who constructed the model does. This does not seem to be a realistic description of how expectations are formed in the real world, and has consequently been criticised by many economists.[40] In order to bring the agents in a model back on an equal footing with the modelling economist the idea of *bounded rationality* was introduced.[41] This concept requires to add behavioural foundations of expectations in a model. This can be done by assuming that agents act like econometricians: They might not know the true model, but they can infer information about it from past experience. The uncertainty about the true model can be of different degree.[42] Agents can be modelled as 'classical econometricians', who have a clear idea of the model's structure in mind and try to estimate its parameters. Alternatively they can be modelled as 'Bayesian econometricians', who are not sure about the model and the parameters, but can specify their uncertainty. Finally, they could even be unsure about whether they regard themselves as classical or Bayesian econometricians. Here I will limit the analysis to agents of the first kind. Note that the fact that agents are certain about the model's structure does not guarantee that their perceived model structure is the same as that of the true model they live in: For example, in their perceived model they could employ functions which are based on a different set of variables, or they could base their model on wrong functional forms.

If agents infer information about the model from past experiences, they will form expectations no longer based on the true model, but on their perceived model. These

---

[40] An example is given by Paul de Grauwe, who stressed the need for a middle way between naïve or adaptive expectations on the one hand and rational expectations on the other hand. In a recent article he wrote: "There is a need to move away from extreme assumptions when modelling human behaviour. Human beings are neither the dumb automatons of the old models nor the supreme creatures of knowledge and understanding of the new models." De Grauwe (2006).

[41] This idea goes back to the work of Simon (1957). It was applied to dynamic economic systems for example in Sargent (1993).

[42] The following designations are taken from Sargent (1993), p. 22.

expectations are only rational if the agents succeed to estimate the model's parameters correctly. Otherwise expectations are no longer rational, but they can still be in accordance with the agents' rationality in the sense that it may be impossible to form better expectations on the grounds of available information. A second feature of such agents is that they revise their inferences about the model when new information arrives, so they are in effect *learning*. It is important to note that this leads to an additional[43] recursion in a dynamic model: The agents form expectations based on their perceived model, their expectations influence the outcomes of the model, which in turn are observed by the agents and are thus used to revise their perceptions of the model. In this way learning provides a way to model explicitly the above mentioned self-referential process of which rational expectations are an equilibrium. With learning, expectation formation is no longer static but evolves according to a dynamic process which will only come to rest if expectations are confirmed by the outcomes, that is in an expectations equilibrium.

Looking at it from a single agent's point of view, it becomes clear that learning can also describe a way of how agents can coordinate their expectations. Every agent forms his expectations, all agents' expectations are aggregated by the model and yield commonly observed outcomes, which again are used by every agent to update their views of the model. Consistency of individual beliefs, as required by Sargent in his description of rational expectations, is a possible limit point of such a process. The process also provides further insight into the question of which of the solutions to a rational expectations model are dynamically attainable, as outlined at the end of the previous section.

The idea of updating agents' expectations repeatedly with the advent of new information has of course already been employed in a very crude and basic way in adaptive expectations. But learning provides a much more general approach to adaptation, of which adaptive expectations are only one special case.[44] Whereas the expectations of a variable themselves are updated according to new information when

[43] Additional means here in addition to recursive functions which may be present in the model anyway.

[44] Adaptive expectations are equivalent to expecting a constant and estimating the constant with a constant gain learning rule. See appendix A.1 as well as Evans and Honkapohja (2001), p. 49.

adaptive expectations are used, in the case of learning the updating concerns the agents' beliefs and in consequence their rules of expectations formation. Thus, in learning the adaptation takes place on a meta-level as compared to adaptive expectations.

Learning agents can be modelled in a multitude of ways, including various methods of artificial intelligence such as neural networks, genetic algorithms, classifier systems and fuzzy logic.[45] Here I will follow an established line of work which models learning as estimating unknown parameters by means of statistical methods. A recurrent question will be under which conditions expectations with learning will converge to an equilibrium, and whether this equilibrium can be considered as a rational expectations equilibrium. If convergence to an equilibrium occurs, this equilibrium is called *learnable*.

In the following section so called iterative or eductive learning will be presented as first examples of learning algorithms. In 2.4.3 I will then turn to adaptive learning and describe the often used RLS algorithm as the main exponent of learning algorithms with decreasing gains. Convergence and stability issues will be treated in 2.4.4. In section 2.5 constant gain learning will be discussed as an important alternative to decreasing gain learning. Section 2.6 summarizes the discussion.

2.4.2. Early Examples of Learning Algorithms and Eductive Learning

An early example of the learning approach as described in the previous section is stated in DeCanio (1979). In the article a cobweb model with an observable exogenous influence and an unobservable disturbance is used to investigate whether the rational expectations equilibrium of the model is learnable. The agents in the model try to forecast the current price level. Their information set consists of past values of the price level and produced quantities as well as of past and current values of the exogenous variable. The expectation is made by means of a linear function which calculates the expected price level from the elements of the agent's information set and a set of weight parameters. This function can describe rational expectations for one specific set of parameters.

---

[45] See Sargent (1993), Chapter 4. Some examples of recent economic applications are Kooths (1998), Ringhut (2002), and Mlakar (2006).

In this model, the agents start their expectation formation with an arbitrary set of parameters and use the associated expectation function to form their expectations for a number of periods. In the next step the outcomes from these periods in connection with the observable variables are used to calculate a new, improved estimation of the parameters. These are again used for a number of periods, and so on. Each set of parameters generates a new series of outcomes, which in turn is used to generate a new set of parameters. Note that it is precisely their expectations changed by learning which in every round invalidates the agents' last estimations. Only if the expectations form an expectations equilibrium their expectations formation is correct in the sense that it is not contradicted by the agents' observation.

The process can be described as a sequence of parameter sets evolving according to a difference equation which links each new parameter set to its predecessor. It can be shown that this sequence of parameter sets converges to the parameter set describing rational expectations under certain conditions.[46] A related model is investigated in Bray (1982).[47] Note that in this process estimations are made only over distinct periods of time, each governed by constant expectations functions. All agents *simultaneously* update their expectations and then start collecting data again. More generally described, in this learning process the agents use a PLM with unknown parameters to calculate their expectations. That induces an ALM, whose parameters are estimated and used in the agents' PLM in the next round.

The following will give a first illustration of learning as it is modelled in the example above. For simplicity, I will again use the example of section 2.3, of which the main results are repeated here for convenience.[48] The reduced form of the system is

$$(20)\quad y_t = \bar{g} + \alpha E_{t-1} y_t + \delta E_{t-1} y_{t+1} + u_t \,.$$

Let us focus on learning of the MSV solution

---

46 See DeCanio (1979), p. 53 f.

47 In this article two different learning procedures are discussed. The procedure in Proposition 4 is the one which is similar to DeCanio (1979). See Bray (1982), p. 325.

48 This is also equivalent to a simplified version of the system in Evans (1985).

(21) $y_t = a + u_t$

which will be considered as the agents' PLM.[49] Let *N* indicate the index of the period of time during which the agents use a_N as a parameter for their expectations. Agents start with a random value of a_0. During the period of time *N*, the agents' expectations are

(22) $E_{t-1}y_t = a_N, \quad E_{t-1}y_{t+1} = a_N$

so the corresponding ALM is

(23) $y_t = \overline{g} + (\alpha + \delta)a_N + u_t$.

At the end of period *N* the agents re-estimate *a* for their expectations in the next period of time $N+1$. Here, this amounts simply to the mean value of y_t during *N*, which results in the difference equation[50]

(24) $a_{N+1} = \overline{g} + (\alpha + \delta)a_N$.

This describes the path how *a* is updated over time. If

(25) $|\alpha + \delta| < 1$

then the difference equation (24) converges to

(26) $\lim_{N \to \infty} a_N = \dfrac{\overline{g}}{1 - \alpha - \delta}$

---

49 Learning a constant is the simplest possibility and is used here only for expository purpose, but the approach works similar for other PLMs. It is interesting to note that in the case of a constant, learning algorithms and adaptive expectations work roughly similar, while of course this would not be true any more if agents tried to learn the parameters of more complicated PLMs.

50 Note that (24) is equivalent to an iteration of a T-map from the PLM to the ALM as in (13) in section 2.2. However, in less simplistic models the estimation of the parameters from the ALM is more involved, so that this equivalence does not hold any more. Nevertheless, the case at hand illustrates the motivation for the concept of E-stability as discussed below.

which yields the MSV solution. So if the model's parameter satisfy condition (25), the learning algorithm will ascertain that expectations converge to rational expectations in the limit. Moreover, this also shows that in this model the MSV solution is learnable for a suitable PLM and a certain set of parameter values.

One weakness of this approach is that the learning behaviour of the agents is still somewhat artificial. Though it may be conceivable that agents do not update their expectations in every period *t*, for example if such updates are costly, it is not natural to assume that all agents update their expectations simultaneously. The answer to this criticism is to allow for an updating of expectation rules in each period. But that means for nearly all models[51] that it is no longer possible to perform the estimation of parameters over a set of data which are the result of constant expectation rules like in the articles mentioned above. Instead the estimations are based on a set of data which stem from different expectation functions for each period. That requires a modified approach to learning, which will be discussed in the next section.

A different interpretation of the learning scheme described above is the concept of *eductive learning*.[52] In the models above learning takes place in real time: It takes the form of repeated updates each of them being performed when new data arrives. In contrast to that, eductive learning consists of a process of mental reasoning by the agents. Similar to related concepts in game theory,[53] an agent repeatedly forms expectations with respect to the expectations of others. He starts by forming an expectation, then forms his expectations of the expectations of others and modifies his expectation to achieve a "best response" to the others' expectations. But since that behaviour would be rational for all agents, he has again to consider similar modifications of expectations by the other agents and has to update his expectations again. This evolution of expectations (in "notional" time) can be described by difference equations like in the examples above. Although eductive learning constitutes an elegant

---

[51] That is, for all models which require more than one data point for the estimation of the parameters.

[52] The idea behind eductive learning is mentioned for example in Evans (1985), p. 1222. A brief example is given in Evans and Hokapohja (2001), pp. 372 ff.

[53] The relations between stability criteria for equilibria of eductive learning (iterative E-Stability) and the corresponding concepts of game theory (rationalizable solutions) is examined in Evans and Guesnerie (1992).

way to model learning, it relies on the assumption that an agent knows enough about the model to calculate his "best response" in each round of reasoning. So in these models agents do know a lot more than the agents described above. The focus of the learning algorithm here is more on determining the stability conditions for a particular expectational equilibrium in a model with strategic agents than on the consequences of the fact that agents do not know the model and have to estimate its parameters.[54] In the following this line of research will therefore be left aside.

2.4.3. Adaptive Learning and Recursive Least Squares

In contrast to the approach in section 2.4.2, in the models below learning will take place in each period. This requires that the learning process must be specified differently. An intuitive way of modelling learning, which corresponds well to the idea of modelling agents as econometricians, would be to assume that agents estimate the parameters of their expectations by running ordinary least squares (OLS) regressions of the history of the variables determined in their PLM on the history of the determining variables (and possibly an intercept, if present in the PLM).

In this section the approach will be illustrated at the example of univariate linear models. Consider a general model of the following structure: Let y be a scalar and ω be a vector of determining variables with constant second moments, and define the vector

$$(27)\quad z_t = \begin{pmatrix} 1 \\ \omega_t \end{pmatrix}.$$

The vector ω_t can include any variable which is part of the agents' information set at time t, so it can include exogenous variables as well as lagged endogenous variables. Let y_t be a linear function of z_{t-1} with unknown parameters plus a disturbance η_t.[55]

---

[54] It turns out that eductively stable equilibria are generally also stable under adaptive learning (described in the following section), but not vice versa. For a comparison between the two types of learning see Evans (2001).

[55] An OLS regression will yield unbiased estimates if the disturbance and the vector z of determining variables are uncorrelated, which we assume here. Otherwise it is possible to modify the approach by using instrumental variables. See for example Ljung and Söderström (1987), pp. 21 ff.

The agents believe that this function describes the actual dependence of y_t from z_{t-1} and try to estimate its parameters. The estimated parameters may be stacked in the vector ϕ. So the agents' PLM can be written as

(28) $y_t = \phi' z_{t-1} + \eta_t$.

The agents form their expectations on the basis of this PLM using the most recent values of the parameter ϕ and of determining variables ω that are available to them.

However, like in the examples before, the *actual* value of y_t may depend not only on z_{t-1} but also linearly on the agents' expectations of $y_t, y_{t+1}, \ldots$ at time $t-1$ and thus on the parameter values ϕ they use to form them. The resulting actual relationship between y_t and z_{t-1} (ALM) may therefore be different from the one the agents believe. Let the ALM be described by

(29) $y_t = \Psi' z_{t-1} + \eta_t$

with Ψ being a vector of parameters (similar to ϕ above). Then these parameter values include the indirect influence of z_{t-1} on y_t via the agents' expectations. Thus Ψ depends on the parameter values ϕ the agents used to form their expectations. This dependence can be expressed as a mapping from the parameters of the PLM to the parameters of the ALM which is usually called a *T-map*:[56]

(30) $\Psi = T(\phi)$.

Thus the *T-map* represents the way how the agents' parameter estimates change the relationship between y_t and z_{t-1}. The ALM can then be described alternatively by

(31) $y_t = T(\phi)' z_{t-1} + \eta_t$.

[56] Such a T-map was already used in the example on p. 12, equation (13).

To estimate their parameters ϕ, the agents run an OLS regression of y_t on z_{t-1}. Generally, given only the PLM, such a regression could be written as

$$(32)\quad \phi_s = \left(s^{-1} \sum_{t=1}^{s} z_{t-1} z'_{t-1} \right)^{-1} \left(s^{-1} \sum_{t=1}^{s} z_{t-1} y_t \right).$$

This yields the vector of estimated parameters in time *s* as an OLS estimation on data from 1 through *s*.[57] But to capture the idea of updating parameters in each period, an OLS regression can alternatively be written in a recursive form (recursive least squares, RLS): [58]

$$(33)\quad \begin{aligned} \phi_t &= \phi_{t-1} + t^{-1} R_t^{-1} z_{t-1} \left(y_t - z'_{t-1} \phi_{t-1} \right) \\ R_t &= R_{t-1} + t^{-1} \left(z_{t-1} z'_{t-1} - R_{t-1} \right) \end{aligned}$$

where R_t is the moment matrix of z_t based on data through $t-1$. With suitable starting values for R_t and ϕ_t this is equivalent to (32). Since estimation with this version of the algorithm is performed whenever new data arrive, RLS is sometimes called an *on-line estimation* as opposed to OLS as an *off-line estimation*.

While I will not enter into a detailed analysis of this algorithm here, it is nevertheless useful to give an intuitive interpretation of it. In the upper equation of (33), the vector of estimations is updated in each period. The amount of updating which takes place is determined by four elements. Starting from right to left, the first element is the *prediction error*, that is the difference between actual outcome and prediction $(y_t - z'_{t-1}\phi_{t-1})$, which is due to the fact that the parameter estimation may be not correct as well as to the disturbance which influences y_t. The prediction error is multiplied by

57 See for example Sargent (1993), p. 37.

58 See Sargent (1993), p. 41 or Evans and Honkapohja (2001), p. 33. For a derivation of the RLS from the OLS see Ljung and Söderström (1987), pp. 17 ff. Because in eq. (33) a matrix inversion has to be performed in each period, it is not well suited for computation. Therefore an equivalent form of the RLS exists which circumvents this disadvantage by setting $P_t = t^{-1} R_t^{-1}$ and updating it directly via the matrix inversion lemma. For details see Ljung and Söderström (1987), p. 19 f. or Harvey (1995), p. 120 f. However, the form in eq. (33) allows for a better interpretation.

the vector of determining variables z_{t-1}. These two elements make sure that the estimations are updated in the right direction and that updating is stronger for large prediction errors. The latter characteristic is somewhat refined by the next element, a *weighting matrix*.[59] Roughly speaking, since large prediction errors are not necessarily due to bad estimations but can also be the result of a high variance of z, the step of including the weighting matrix R_t^{-1} (i.e., the inverse of the estimated second moments matrix of z) can compensate for this effect: The higher the (estimated) variance of a regressor is, the less information it transmits, and its influence is consequently given less weight.[60]

Finally, each update of the estimations for each new observation is weighted with a *gain* factor[61] before it is added to the previous estimation values. Over the course of time these factors form a *sequence of gains* which determines the behaviour of the learning algorithm in time with respect to new observations. A decreasing gain like above makes sure that the estimates stabilize and "settle down" in time. Here, the specific function t^{-1} guarantees that every observation has the same weight in the estimation, exactly like in an OLS regression. In contrast to that, there are also learning algorithms with a constant gain,[62] in which recent observations have more weight than old ones. The lower equation of (33) can be seen as a simple updating of R_t.

2.4.4. Stability under Learning and E-Stability

In this section I will examine under which conditions the adaptive learning approach can converge to an REE solution of the model in question, that is, whether an REE solution is learnable. After defining the general problem at the example of

59 If this weighting matrix would be left out, the resulting learning algorithm would be "stochastic gradient learning" instead of RLS learning. For details, see Evans, Hokapohja and Williams (2005).

60 See Williams (2004), p. 6.

61 The gain γ is usually required to be a positive, nonstochastic and nonincreasing sequence. To ensure convergence to the equilibrium it is usually also demanded that it satisfies $\sum_{t=1}^{\infty}\gamma_t = \infty$ and $\sum_{t=1}^{\infty}\gamma_t^2 < \infty$. See Evans and Honkapohja (2001), pp. 36 and 123 f.

62 The relevance of the sequence of gains and constant-gain learning will be examined in more detail in section 2.5.1.

RLS learning, I will characterize the dynamics which result from the application of RLS learning to the agents' expectations. I will then define the criterion of *E-stability* as a simple way to determine convergence of the learning algorithm and discuss under which conditions this criterion can be applied. The relations between these elements of analysis will be summarized in *Figure 2.1*.

As mentioned in the previous sections, running a regression over several periods of a self-referential economic model and updating the parameter estimates in each period means that such a regression is inherently faulty. By using a least squares regression it is assumed that there exists a constant set of parameters which relates the determined variable to the determining ones. However, the sample of data on which the regression is run is not the result of a constant relationship between the variables: For each data point the agents had a different parameter estimate, so for each data point a different ALM was in effect. As a straightforward remedy one could imagine to include forecasts made on the basis of the estimated parameters for each period into the set of regressors. But that is not possible because these forecasts are themselves dependent on z, which leads to difficulties with collinearity in the regression.[63]

This feedback of the time-varying estimates on the system can be exemplified by replacing y_t by the ALM (31) in the RLS estimation (33):

$$(34)\quad \begin{aligned} \phi_t &= \phi_{t-1} + t^{-1} R_t^{-1} z_{t-1} \left[z_{t-1}' \left(T(\phi_{t-1}) - \phi_{t-1} \right) + \eta_t \right] \\ R_t &= R_{t-1} + t^{-1} \left(z_{t-1} z_{t-1}' - R_{t-1} \right). \end{aligned}$$

This is a stochastic recursive algorithm (SRA).

Although a linear regression is thus no longer applicable to the system, it is nevertheless possible that such a learning algorithm converges to an equilibrium. Whether this is the case, and how the equilibrium thus reached can be characterized, can be determined by

63 See for example Bray and Savin (1986), p. 1133. In this article they also point out that agents in such a model are not fully rational and do not behave optimally, because they behave as if the ALM of the variables would be constant, while actually it is constantly changing due to their learning. It would only be constant in the REE when learning dynamics vanish. But since agents cannot know the REE, following such a learning rule is the most rational behaviour available to them.

use of stochastic approximation techniques applied to SRAs. The foundations of this approach are taken from Ljung and have been applied to economic learning mechanisms by Marcet and Sargent.[64] Using this approach, an ordinary differential equation (ODE) is associated with the SRA in question. It can be shown that under some regularity conditions[65] the convergence properties of the recursive stochastic algorithm can be determined by studying the stability properties of the ODE. Roughly speaking, the associated ODE is constructed by determining the average updating direction for a fixed vector of estimations (average in the sense of taking the expected value over the distribution of z for the fixed vector of estimations).[66] The ODE for the example above would be[67]

$$\begin{aligned} &\frac{d}{d\tau}\phi_D(\tau) = R_D(\tau)^{-1} E(zz')\left(T\left(\phi_D(\tau)\right) - \phi_D(\tau)\right) \\ &\frac{d}{d\tau}R_D(\tau) = E(zz') - R_D(\tau) \end{aligned} \tag{35}$$

In Marcet and Sargent (1989) it was demonstrated that rather than analyzing the whole ODE, local convergence (and thus stability) of such a system can be analyzed by only considering the difference between the vector of estimations and its T-map, $T\left(\phi_D(\tau)\right) - \phi_D(\tau)$.[68] This result can be illustrated by (35): Since $E(zz')$ is constant, the second equation is stable and R_D converges to this value in the limit. Then in the first

64 See for example Marcet and Sargent (1989) and the references cited therein. Various recursive stochastic algorithms are treated in Ljung and Söderström (1987).

65 These concern the functions within the algorithm, the stochastic processes driving it, and – most important – the sequence of gains used. For details, see Ljung and Söderström (1987), chapter 4.3.

66 For a heuristic discussion see Ljung and Söderström (1987), pp. 146 ff. A „receipe" for the construction of the ODE is given on p. 151.

67 The ODE can be used to analyse the convergence and stability of the underlying SRA, but it describes the actual development of the variables in the SRA only approximately. The variables in the ODE are therefore distinguished from their counterparts by the subscript *D*, and a different time scale τ is introduced.

68 See Marcet and Sargent (1989), proposition 3 on p. 351. The result relies on some regularity assumptions and some restrictions on the sets of permissible values for *z*, R and ϕ. Global convergence is more difficult to determine. See proposition 1 in the same article.

equation, stability and convergence is determined only by the behaviour of $T\left(\phi_D(\tau)\right)-\phi_D(\tau)$.[69]

This leads to the criterion of *expectational stability* or *E-stability* with the following definition: Consider an economic model, an REE solution to the model and a PLM which nests this solution. Thus the REE solution is characterized by a vector of parameters of the agents' PLM, such as ϕ in the example above. Assume that the agents try to learn the REE value of ϕ from past experiences by an adaptive learning algorithm. As before the mapping from PLM to ALM may be denoted by $T(\cdot)$. With this mapping the ODE

$$(36)\quad \frac{d}{d\tau}\phi(\tau)=T\left(\phi(\tau)\right)-\phi(\tau)$$

is constructed. It is clear that if (36) converges, it can only converge to REE values of ϕ, because $T\left(\phi(\tau)\right)=\phi(\tau)$ only for an REE. An REE is then called *expectationally stable* or *E-stable* if the ODE (36) is locally asymptotically stable for the REE value of ϕ, i.e. if in the limit (36) converges to the REE value when starting from within a neighbourhood of it.[70] The intuition behind that criterion is that an expectations equilibrium is expectationally stable if the following condition is fulfilled: If there is a (small) random deviation of the expectation function from the expectations equilibrium, then the process of expectations updating (or learning) returns the expectations towards the equilibrium.[71] Thus the ODE can be interpreted as a "stylized learning process in notional time"[72]. For a given model structure and a PLM, the stability of (36) around the REE in question depends on the actual parameters of the model. Therefore, with the help of this ODE it is possible to determine conditions on the model's parameters which ensure E-stability of the corresponding REE.

---

69 Provided R_D is invertable along the trajectory. See Evans and Honkapohja (2001), p. 38.

70 The concept of E-stability was introduced by Evans, first in an iterative form which corresponds to iterative learning mechanisms as described in the previous section, and later in the form given here suitable for adaptive learning. For the former see Evans (1985), the latter is discussed in detail in Evans and Hinkapohja (2001). For the following results, see there.

71 See Evans (1985), p. 1221.

72 Evans (1989), p. 299.

In many systems of a model and a learning algorithm there is an equivalence between E-stability of an REE and the stability of the REE under learning. Such an equivalence means that if the ODE (36) is stable for a set of model parameter values, then for the same set the REE is stable under learning. Moreover, by virtue of the learning process, the agents' PLM will then converge to the REE solution almost surely, if the initial value of ϕ is sufficiently close to the REE value (local convergence). In other words, the REE is learnable. For some models and REEs this convergence can also be shown to hold for arbitrary initial values of ϕ (global convergence). Likewise the equivalence between E-stability and stability under learning implies: If (36) is unstable for a set of model parameter values, then for the same set the REE is unstable under learning and the agents' PLM will converge on the REE solution with probability zero. *Figure 2.1* illustrates the E-stability approach and the relationships described above.

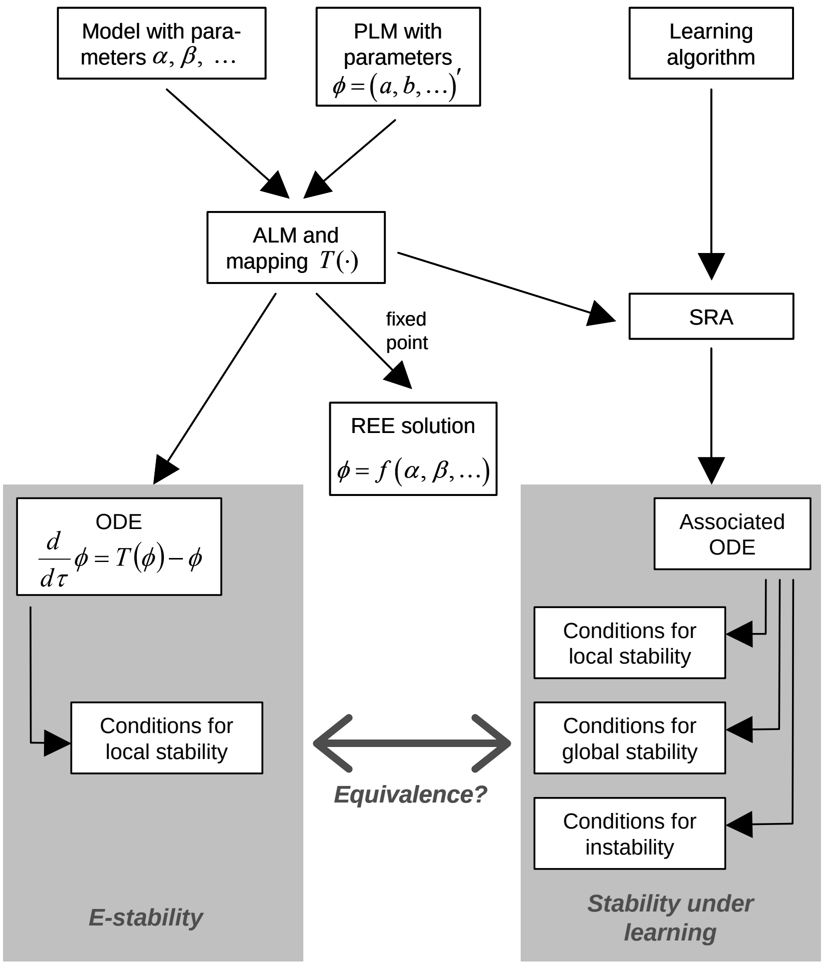

Figure 2.1: *E-stability and stability under learning*

This equivalence has been shown to obtain for a wide range of models with a learning algorithm,[73] but for some cases it could not be proven generally so far, while in others it was shown that it does not hold. This depends

- on the class of models which is examined,
- on the PLM employed by the agents and
- on the learning algorithm.

To show the equivalence for a class of models with a suitable PLM and a learning algorithm, the conditions for convergence under learning are derived by stochastic approximation and are then compared with the conditions for E-stability. However, in some cases the conditions for the application of stochastic approximation are not fulfilled, so the technique cannot be used.[74]

A crucial point is the type of learning algorithm. One main condition for the equivalence is that the gain of the learning algorithm has to be decreasing to zero in the limit. Otherwise the agents will forever continue to adapt the learned parameter values to random fluctuations of the observed data, and the system will never "settle down".[75] Nevertheless there are some advantages of using a learning algorithm whose gain does not decrease to zero. This will be discussed below.[76] A possibility that is often used are learning algorithms with a small but constant gain. But even among the learning algorithms with gains decreasing to zero there are examples for which the equivalence between E-stability and stability under learning is not guaranteed. While the equivalence holds for many models with RLS learning, it has been shown that for

---

[73] Evans and Honkapohja (2001) examined the conditions for this in detail, especially for least squares learning.

[74] For example if the REE in question is an unbounded continuum of solutions rather then a single point. See Evans and Honkapohja (2001), p. 192.

[75] While the vector of parameter estimates will thus not converge to a point, it may still be possible that it converges to a limiting probability distribution. See Evans and Honkapohja (2001), p. 163.

[76] See the discussion of constant gain learning in section 2.5.1.

stochastic gradient learning[77] a system can converge to solutions which are not E-stable, and that there are E-stable solutions which are not stable under learning.[78]

In the literature a distinction is made between a weak and a strong form of E-stability.[79] In the description above it was tacitly assumed that not only the PLM nests the REE in question, but that it also does not contain additional variables which are not part of the REE solution (that it is not "overparameterized"). In this case E-stability is called *weak*. Additionally, if it is possible to find conditions on the model's parameters which ensure E-stability for an overparameterized PLM as well, these are called *strong E-stability conditions*.[80] Intuitively it is plausible that convergence of the resulting SRA to the REE gets more "difficult" if additional unnecessary variables are included in the PLM. Consequently, the conditions on the model's parameters for E-stability tend to be more restrictive in this case. Under strong E-stability the estimated parameters for the superfluous variables in the PLM approach zero in the process of convergence. It is thus learned by the agents that – in contrast to their initial perceptions – these variables are actually not relevant.

For different possible REE solutions (such as the MSV solution, explosive or non-explosive bubble solutions) there can be different conditions on the model parameters which ensure E-stability of the respective REE. Some REE solutions turn out not to be E-stable at all, and if the equivalence between E-stability and stability under learning holds then these REEs are never learnable as well. In this case the REE in question will occur with probability zero. Even if the agent's expectations formation coincides initially with such an REE, already minimal random deviations suffice to lead the system away from the REE. Economically, such REEs are therefore not relevant. The E-stability criterion can consequently be used to select economically meaningful REEs among the set of possible REEs if there are multiple equilibria as described in section 2.3. Since the MSV solution is of special interest in this context, it has been

---

77 Stochastic gradient (SG) learning (also called *least mean squares* learning) resembles RLS learning, but it does not use the matrix of second moments as a weighting matrix. For an analysis of SG learning and an extension of the algorithm, see Evans, Hokapohja and Williams (2005).

78 See Heinemann (2000).

79 See Evans and Honkapohja (2001), for example p. 140 f.

80 Since there are several possibilities to overparameterize the PLM, strong E-stability conditions are usually specific to a class of PLM. See Evans and Honkapohja (2001), p. 141.

examined for some classes of linear models whether this solution is always E-stable. Although it can be shown that for some parameter values the MSV-solution is not E-stable,[81] McCallum (2002) argues that these parameter values are economically not relevant.[82] He demands that an economic model has to be "well formulated"[83] in order to be economically plausible. According to the definition given in the paper, for a model to be well formulated there must be restrictions on the admissible parameter values which exclude infinite discontinuities, but still "admit a large open set of values" including zero. Although no comprehensive results exist so far, it has been shown that at least for a class of well formulated univariate linear models the MSV-solution is always E-stable and therefore learnable.[84]

Let us get back to the initial question whether RLS and related learning algorithms can be used for continuous learning, although the ALM producing the data basis for it is changing from period to period. If the SRA (respectively the associated ODE) converges to stable values for ϕ and R, then the learning algorithm allows the parameter estimates to converge to the REE values in the limit,[85] even if the algorithm used for estimation (e.g. RLS) does not give unbiased estimates along the way. It depends on the model and on the model's actual parameter values whether the system converges to a given REE solution and whether convergence is local or global. If the equivalence between E-stability and stability under learning holds for the system, then convergence can be determined more easily by examining under which conditions E-stability occurs.

---

81 See Evans and Honkapohja (2001), p. 183 ff and 198 ff.

82 See McCallum (2002), p. 13 ff.

83 The characteristic of being "well formulated" within the context of learning is introduced in McCallum (2002), p. 17.

84 See McCallum (2002), p. 18 ff.

85 Provided either the REE is not only locally but also globally stable, or that the algorithm starts within an appropriate neighbourhood of the REE.

2.5. Constant Gain Learning, Misspecification, and Escape Dynamics

2.5.1. Constant Gain Learning

The behaviour of a system which includes a learning algorithm depends to a considerable degree on the development of the updating of the learned parameters over time. This is controlled by the *gain* parameter of the learning algorithm[86], often denoted by γ. In many learning algorithms the sequence of gains decreases to zero over time. Examples are the stochastic gradient algorithm or the RLS algorithm, in which the gain is given by $\gamma_t = t^{-1}$. This sequence of gains ensures in the RLS algorithm that every observed data point and thus every observed forecasting error have the same weight in the calculation of the current estimates, similar to the ordinary least squares algorithm which also associates the same weight to every data point.[87]

There are several consequences from such a decreasing gain: On the one hand, it allows the system to converge over time.[88] While the gain decreases, the impact of new observations, which are still influenced by random disturbances, gets arbitrarily small. Hence it gets increasingly less likely that the learned parameters can be deviated from their optimal values by a disturbance, and in the limit the system will "settle down" to the optimum estimates. But on the other hand, this is only advantageous if the actual model parameters are constant. If for some reason the parameters vary over time, either continuously or intermittently, the influence of all former parameter values will forever be present in the calculation of the actual estimates. The current estimates will therefore always be biased with respect to the current parameter values. That means that the learning algorithm is not able to track parameter changes of the model it is supposed to estimate.

An alternative is to use a constant gain learning algorithm. This means to use a (small) constant value as the gain parameter $\gamma_t = \bar{\gamma}$. It is clear that with such a gain the agents'

---

[86] See the discussion of the learning algorithm on p. 26 above.

[87] This can be shown by simply transforming the OLS algorithm in recursive form. See Ljung and Söderström (1987), section 2.2.1.

[88] Provided the appropriate equilibrium is E-stable. See section 2.4.4 above for convergence properties of the SRA resulting from the learning algorithm.

observations are no longer equally weighted. Instead, the last observed values are given the greatest weight, and the weights of earlier observations decrease geometrically with their age.[89] The influence of old data gets thus deliberately small over time. If a parameter of the actual model changes, the algorithm will ultimately "forget" the data generated by the former parameter value, and it will adapt to the new value by forming its estimations predominantly from recently generated data. This property of constant gain algorithms is called *tracking*.

Moreover, a constant gain algorithm can be superior to a decreasing gain algorithm when it is used to learn the parameters of a misspecified model.[90] Suppose the agents' PLM is misspecified in the sense that an exogenous variable which is present in the ALM is omitted in the PLM. Then this variable may influence the outcomes on which the agents' estimates are based without them recognising this influence.[91] Depending on the evolution of this variable it can be advantageous to employ an algorithm which can adapt repeatedly to such unaccounted fluctuations to estimate the PLM's parameters.

Intuitively, both reasons to employ a constant gain algorithm are related. Within the class of models described by a misspecified PLM, the data generated by the ALM with the additional influencing variable can look like stemming from a model with constantly changing parameters. Therefore in some cases a tracking algorithm which can adapt to such changes may yield better forecasts than a decreasing gain algorithm.

An important property of a tracking algorithm is that in a stochastic environment the algorithm will never converge to equilibrium estimates or to another point, not even in the limit, because new – randomly disturbed – data will always have a certain influence and will deviate the estimates away from their equilibrium values. That means that a model with a constant gain algorithm will exhibit "persistent learning dynamics"[92], which represent a further source of dynamics in the model, in addition to the usual stochastic disturbances. The size and behaviour of this deviations depend, apart from

---

89 See Marcet and Nicolini (2003), p. 1484, or Carceles-Poveda and Giannitsarou (2007), p. 2680 f. for details.

90 See Williams (2004), p. 6 as well as Evans and Honkapohja (2001), p. 332.

91 This will be examined in more detail in section 2.5.2.

92 In the words of Evans and Hokapohja (2001).

the model structure, mainly on the size of the constant gain parameter $\bar{\gamma}$ and on the distribution of the disturbances. The lower the value of the gain, and the tighter the distribution of the disturbances, the smaller the deviations get and the more the learned estimations stay near their equilibrium values. For some models it has already been shown that with a small gain and a bounded support of the disturbances, in the limit the estimates are distributed stochastically in an unbiased and tight normal distribution around the equilibrium values,[93] but more complicated dynamics are possible as well.[94]

As a consequence of such learning dynamics, it is possible that agents are forced to coordinate on one learning algorithm: If some agents choose to use a constant gain learning algorithm (as opposed to decreasing gain learning), they can cause the system to exhibit additional dynamics, under which a constant gain algorithm is indeed the one who yields the best forecasts, so rational agents would choose this kind of learning algorithm. Along the lines of an analogous argument agents can also be forced to coordinate on the same optimal value for the constant gain parameter $\bar{\gamma}$.[95] The choice of constant gain algorithm for learning can thus be a kind of "self-fulfilling prophecy".

Above it was suggested that one possible motivation to use a constant gain learning algorithm is the possibility that agents use a misspecified model for their forecasts. The consequences of a misspecified PLM, especially in connection with a constant gain learning algorithm, will be described in the following two sections.

2.5.2. Self-Confirming Equilibria and Misspecification

In this section the type of equilibrium which can be reached with a misspecified PLM will be characterized. From the different possibilities for a misspecification, here and in the following the discussion will focus on the case that a variable which influences the outcomes in the ALM is omitted from the PLM. As described in section 2.2, the original

---

93 E-stability appears to be a necessary condition for this convergence to take place. In addition to that, some further assumptions and conditions are required for convergence to such a distribution. See Evans and Honkapohja (2001), p. 333 ff.

94 See section 2.5.3 below.

95 Evans and Honkapohja (2001), p. 346 ff. gives a short analysis of this line of thought. See also Evans (2001), section 4.

condition of Muth which should be satisfied by "rational"[96] expectations was that the agents do not perceive any contradiction between their expectations and the observed data generated by the true model. But an REE is only one special possibility to satisfy this condition, namely the one requiring that agents base their expectations on beliefs about the true model which are complete and correct. This would be the case if the PLM were well specified (so it nests the REE) and all parameter estimates of the PLM were correct. But there may be other beliefs satisfying the condition as well although they are not correct. To satisfy the condition, it is only necessary that these beliefs cause the actual true model to generate data which are consistent with the beliefs. In other words, it is necessary that the agents cannot detect any correlations between the forecasting errors and the data used in their subjective model. Then the agents attribute the forecasting errors to noise and do not have any reason to change their beliefs. Subsequently they do not change their behaviour and the model continues to generate data which do not contradict their beliefs.

So this condition describes an equilibrium in the self-referential process of expectation formation and data generation in a more general way than an REE. Such an equilibrium is called *self confirming equilibrium* (SCE).[97] An REE is therefore the special case of a SCE in which all beliefs are correct. But apart from that, a SCE can also exist with incorrect, misspecified beliefs. It is sufficient that the agents' beliefs (and their resulting expectations) are correct about events that are observed sufficiently often, and that they are incorrect only about events that are rarely observed.[98]

If the agents' PLM is misspecified, then it does not nest the REE. Therefore no matter how the learning of the PLM's parameters takes place, it is impossible that the agents' beliefs will ever be completely correct, and their expectations will never become rational expectations. Nevertheless it is still possible that by learning the system converges to a SCE. In the case of a misspecified PLM such a SCE is also called more specifically *restricted perceptions equilibrium* (RPE).[99] As in every SCE, in an RPE

---

96 Here, the term "rational" is used in the broader sense, not in the specific meaning of RE.

97 The term was first coined in game theory for an analogous situation. See Fuldenberg and Levine (1998).

98 See Hansen and Sargent (2001), p. 532 ff., Williams (2004), p. 1., or Cho and Sargent (2006).

99 See Evans and Hokapohja (2001), p. 57 and p. 317 ff.

agents also do not perceive any correlation between their forecasting errors and the observed data they use in their model.

2.5.3. Escape Dynamics

As suggested in the section 2.5.1, constant gain learning can lead to more complicated dynamics than estimates simply converging once and for all to a distribution around the equilibrium values. A description of the conditions for these dynamics and their properties will be given in this section.

When agents learn, they estimate the parameters of their PLM on the basis of past observations. Since these observations are influenced by random shocks, the sequence of estimates which results from learning is not always moving smoothly in the direction of the equilibrium values. Instead the estimates can repeatedly be moved away from the equilibrium values if large values are observed for the disturbance variables in a model. With decreasing gain learning, the influence of newly observed data gets smaller as agents learn, so the probability and size of such movements away from the equilibrium decreases. In the limit the system settles down at the equilibrium values.

In contrast, if agents use a constant gain learning algorithm, these deviations from the equilibrium will occur repeatedly forever, because past observations are discounted and newly observed data will always have a positive weight. The movement away from equilibrium can become considerable due to an accumulation of stochastic shocks. Such movements are called *escapes*.[100] The particular path followed by the agents' beliefs during an escape is the result of the interaction of two opposing dynamics.

Consider first a system in which agents use a constant gain learning algorithm, but with an infinitesimal small gain. So the influence of a single random shock on the agents' beliefs approaches zero, and escapes cannot occur. The remaining dynamics of the agents' beliefs are described by the mathematical expectation of the learning dynamics. That is because, when agents are learning with an infinitesimal small gain, they average over an infinite number of observances, so deviations due to random shocks are

---

[100] Escape dynamics in learning models have been analyzed formally in Williams (2004). The following results are taken from there, but are described here in a non-formal, intuitive way.

averaged out (law of large numbers). This expected evolution of the agents estimates is called *mean dynamics*.[101] If the equilibrium is learnable, then the mean dynamics pull the agents' estimates smoothly towards the equilibrium values.

But with a positive constant gain, the evolution of agents' beliefs consist of the mean dynamics superimposed with the influence of random shocks. This influence, separated from the mean dynamics, is called *escape dynamics*. Although the escape dynamics are caused by random shocks, it is possible to characterize the most likely path followed by the agents' estimates in case of an escape, called the *dominant escape path*.[102] The higher the value for the gain parameter, the more likely escapes can be observed for a given time interval.

Depending on the model, an escape path can move the agents' beliefs away from the equilibrium to a noticeable degree and for a comparatively long time, if the system exhibits features that reinforce the escape dynamics or dampen the mean dynamics along the escape path, at least temporarily. This phenomenon has been shown in different types of models. First, escapes can cause agents to switch between multiple equilibria. For example, in a model in Evans and Honkapohja (2001) there exist two stable equilibria. Large escapes can cause the agents to switch from the neighbourhood of one equilibrium, switch to the other and remain around there until another escape cause them to switch back.[103] Such large escapes can occur in the model if the value of the gain parameter or the support of the disturbances are large enough. Once escaped from the first equilibrium, the agents estimates are moved further away and towards the second equilibrium by the mean dynamics due to which the second equilibrium is stable, and vice versa.

A different but related situation is considered in an illustrative model in Williams (2004). Here, there is a unique equilibrium, but a second point describes

---

[101] In the case of decreasing gain learning, the agents estimates follow similar dynamics in the limit. Recall that with the stochastic approximation technique we also used the mathematical expectation of the exogenous disturbances to construct the associated ODE which describes the evolution of the estimates.

[102] For details, see Williams (2004).

[103] See Evans and Honkapohja (2001), p. 337 ff.

“nearly” an equilibrium. That is, this point would be an equilibrium in a deterministic version of the model, and in the stochastic version the learning dynamics are nevertheless comparatively weak around it. Here the dominant escape path leads from the unique equilibrium towards this point. During an escape agents beliefs follow most likely this path away from the equilibrium and remains around the “near equilibrium” for some time because of the weak mean dynamics there, but eventually return to the equilibrium.

A third mechanism which can reinforce the escape dynamics is a combination of misspecification and feedback of agents’ actions on the model dynamics. An example is the model used in Bullard and Cho (2005), which will be described here without going much into the details. The model consists of a central bank and private agents. The central bank uses a Taylor rule but adapts the inflation target of the rule to the target which is expected by the agents. The agents have misspecified perceptions of the central banks policy in the sense that they expect the inflation target to be independently set by the central bank but unknown to them. Instead, they infer this inflation target by trying to learn the parameters of the central bank’s Taylor rule with constant gain learning. The system is constructed to exhibit the following feedback: If one of the estimated parameters of the Taylor rule (as perceived by the private agents) escapes towards lower values, then the agents infer a lower target inflation rate. This in turn is incorporated into the central banks rule, which generates observations pushing said estimated parameter further down, until the system comes to a rest in a point with much lower inflation than in the equilibrium. However, this point is no equilibrium, and the system eventually returns to the equilibrium it started from.

In the model, the escape dynamics are made possible by the constant gain learning algorithm and are reinforced by the combination of misspecification and feedback in the model. As a result, the agents beliefs escape repeatedly for longer periods towards the low inflation point. The authors emphasize that the escape dynamics are not generated by a change of central bank policy, and that the misspecification is quite subtle and not detectable by the agents in the SCE. The resulting escapes are nevertheless relatively huge and will come as a surprise for the agents as well as for the central bank, since from their point of view nothing special happened to trigger it.

This last point is true for escapes in general: They represent distinctive shifts in beliefs and consequently shifts in model outcomes. For the agents in the model these shifts seem to come "out of nothing". In this respect escapes may help describe sudden changes within empirically observed time series that could not be explained by models up to now.

2.5.4. Plausibility and Empirical Support

While not many empirical studies have been conducted so far, constant gain learning and the resulting dynamics seem to be a more natural description of actual real-world learning behaviour than decreasing gain learning. Several theoretical arguments seem to support this view. First, agents cannot be sure that the environment is not changing. Therefore it may be rational to allow for a changing environment when choosing a learning algorithm. Likewise, agents cannot be sure to employ a correctly specified PLM. As explained in the previous sections, this could also be a motivation to use a constant gain algorithm. Both arguments are valid especially if agents are risk-averse. Confronted with an unknown environment, agents use a learning algorithm to derive information about it, but since "unknown environment" includes the fact that it is also unknown to them whether it is changing or not and whether they use a correctly specified model of the environment, it would be rational to allow for these uncertainties by using a constant gain algorithm, while accepting that the resulting estimates may not converge completely. That kind of reasoning could be well expressed by the old adage that "it is better to be approximately right than exactly wrong".

Apart from that and seen from a technical point of view, in a decreasing gain algorithm the starting point (the time of the first data point used for learning) of the algorithm has a significance it does not seem to deserve. While data from this point and all points later enter the algorithm with equal weight, all data before this point are discarded completely. So if an agent uses decreasing gain learning, he deliberately partitions data into one part to use and one part to ignore, which forms a sudden break in the treatment of the data generated by the economy.[104] In contrast to that, constant gain learning admittedly also has a starting point, but nevertheless the first data included in the

[104] The same argument applies if it is not the agent who decides on the starting point, but the fact that data are available only from a certain date on.

learning process always have the least weight, approaching zero in the limit. The transition between the treatment of data before and after the starting point is thus smoothed.

Considering the level of single agents, a related conjecture can be made. Suppose that agents all use decreasing gain learning algorithms, but are heterogeneous in so far as they use data from different starting points on. Then the probability that a certain data point is used for an agent's estimations decreases the older it is. It seems plausible that a lower weight for older observations can be the result of the aggregation of such agents. Even if the resulting distribution of weights may not be exactly described by the weights that result from constant gain learning, it can be conjectured that such a learning algorithm, applied to a (fictitious) representative agent, may mimic the heterogeneous agents' behaviour better than decreasing gain learning.

In the empirical studies conducted so far, constant gain learning appears to provide a reasonably accurate description of real-world behaviour. The consistency of the theoretical learning model with empirical results can be examined in two ways: Either directly by comparing survey results about subjective expectations of individuals with the results suggested by learning algorithms, or indirectly by comparing economic time series which may result from learning behaviour with a theoretical model with bounded rationality and learning. The latter approach can be implemented by fitting such a model to economic time series or by comparing the characteristics of numerical simulations of the model with the characteristics of the respective time series.

Both approaches are used in Branch and Evans (2006). The authors use a simple VAR model with two alternative types of learning (decreasing and constant gain) to forecast inflation and output growth. They also compared the forecasts made by the model to the ones of professional forecasters collected in a survey. The model with constant gain learning performed significantly better in forecasting out of sample than the one with decreasing gain. Constant gain learning (albeit with slightly different values for the gain parameter) also provided the best fit with the professional forecasts. In a related study Weber (2007) used a similar model to examine how households and professional forecasters learn about inflation evolution in different countries of the European

Currency Union. Again, the forecasting both of households and of professionals can best be described by constant gain learning. The constant gain parameter was generally higher with professional forecasters than with households. This could suggest a greater awareness for a possibly changing environment on the side of the professionals.

It is sometimes difficult to explain sudden shifts and changes in economic time series by macroeconomic models, if they occur seemingly without being triggered by detectable changes in behaviour or the environment. Especially since the application of rational expectations has become standard practice in macroeconomic models, it has been a common problem that those models often do not exhibit sufficient dynamics in comparison to the time series empirically observed. The persistent learning dynamics and the possibility of recurrent escapes in models with constant gain learning therefore represent an intriguing feature. In some studies it has been examined how simulated escapes can mimic the pattern and stylized facts of periods of high inflation, which sometimes seem to occur without a particular cause as well.

In Albert and Nicolini (2003) the authors develop a model to emulate stylized facts of recurrent hyperinflations in South American countries. Those periods can occur without a change in fiscal or monetary policy and after a long period of low inflation. In the model it is possible that agents beliefs about expected inflation escape a low-inflation equilibrium so that the model enters an unstable area in which hyperinflations are likely to occur. In periods of hyperinflations the agents' estimates are driven further away from the equilibrium by a positive feedback between actual and perceived inflation until the government forces inflation to stop by imposing an exchange rate rule. The agents in the model use two learning rules alternatively. They switch from RLS to constant gain learning when they detect a changing environment, that is, when their prediction error exceeds a threshold. In periods of hyperinflation this is the case, and the use of constant gain learning increases the probability that after a period of hyperinflation the system enters the unstable area again. Those escapes into hyperinflation are therefore caused at least in part by the persistent learning dynamics due to the constant gain algorithm. The authors show in the article that their model matches the stylized facts about those hyperinflations better than standard inflation models, and that their model differs from those only in the use of learning instead of rational expectations.

The study of Sargent, Williams and Zha (2006) extends the examination of South American inflation. The model used in this paper is in many points similar to the one of Marcet and Nicolini (2003). As in this model, escapes from the equilibrium are possible and can trigger a period of hyperinflation, which is eventually stopped by government reforms. A prominent difference is that agents in this model always use a constant gain learning algorithm. In contrast to most other studies concerning the empirical significance of learning algorithms, which compare only stylized facts or some characteristic figures of the model and the empirical data, the authors actually estimate their model for five countries and compare the fit of their model with the fit of a standard autoregressive model. Their finding is that their model fits the data very well, especially in the hyperinflation episodes, where it performs even better than alternative models.[105]

Finally, the model in the already mentioned article by Bullard and Cho (2005) can mimic the evolution of inflation towards a low level, like for example in Japan during the 1990s. Although the authors stress that the article's theme is a purely theoretical one, their model is motivated by the experiences of Japan. They do not apply it to empirical data, but the general framework of the model and the resulting evolution of inflation resemble the basic picture of the Japanese data. The sudden shift to low inflation rates and the seemingly unexplainable inefficiency of a previously well working central bank policy can be well described by a learning model featuring escapes.

2.6. Summary

In this chapter the basic concepts of adaptive learning have been described. Learning was introduced as a dynamic formation of expectations. It has been shown how rational expectations can be regarded as an equilibrium point of recursive least squares learning and how the dynamics and the stability properties of the learning process can often be captured by the E-stability concept. In the last section constant gain learning was discussed as an alternative to RLS learning. The studies cited in this section, as well as

---

105 See Sargent, Williams and Zha (2006), p. 27 ff. This is in particular rewarding because from a theoretical point of view a standard AR model may fit the data well, but does not explain the underlying mechanisms in a satisfying way, in contrast to a model like the one used there.

theoretical arguments, support the view that constant gain learning, also in combination with misspecification, can be a rewarding method to extend macroeconomic models: While being able to explain expectation formation in a more realistic way than standard rational expectations, these methods can also improve the empirical fit of macroeconomic models. However, constant gain learning can be seen as being based on RLS learning and as a refinement of it. As described above, models with constant gain learning exhibit additional dynamics with respect to RLS learning, because stochastic disturbances will continue to influence the agents' learning forever due to the constant gain. This means that in general those models tend to be less stable than their counterparts which use RLS learning. Moreover, the dynamics of models with constant gain learning are more complicated to describe mathematically, so the basic consequences of the introduction of learning instead of rational expectations in a model can be shown less clearly.

For these reasons RLS learning is still the method of choice in many of the more abstract models, albeit constant gain learning is employed more and more in empirical works. Roughly speaking, it is in general more convenient to use RLS learning in a basic, abstract model to examine its basic learning properties first in a straightforward way. Once it has been established that the model's behaviour under RLS learning is unproblematic, constant gain learning can be used as a more refined variant to render the model more realistic. This work examines some basic consequences of the introduction of learning into a dynamic optimization model by the example of a basic and abstract consumption model. It is therefore reasonable to employ the recursive least squares learning method in order to concentrate on the main features of the model and to render the results as clear as possible.

3. Learning in a Dynamic Optimization Model

3.1. A Basic Consumption Model

3.1.1. Dynamic Optimization

The RLS learning technique described in the previous chapter will now be applied to a basic consumption model with dynamically optimizing agents.[106] The model is a very elementary and rudimentary example of a micro-founded dynamic general equilibrium model, but as the dynamic optimization approach to determine the level of consumption is one of the building blocks of many more comprehensive models, it is a useful object to demonstrate the impact of learning. The consumption model is used here as a concrete example, but the results of the discussion could be applied to other dynamic optimization problems with a similar structure. Although well known, the canonical solving technique of the model will be discussed in detail in order to emphasize some of its properties which will be important when introducing learning. While in its usual form the model uses a representative agent, I will allow for heterogeneity from the beginning. This will make it easier to examine the impact of the representative agent assumption later on. Throughout this work, heterogeneity should be understood as the possibility of heterogeneous values for the agent's variables like income or consumption and also for his beliefs and expectations, but not for his fundamental parameters like his temporal discount factor or the shape of his utility function, which are still assumed to be the same for all agents.

The model describes a closed economy which is inhabited by a number of (possibly heterogeneous[107]) infinitely-lived private consumers, indexed by $i \in \mathbb{N}$. In this section, I will consider the individual consumer's dynamic optimization problem, leaving the effects of aggregation aside for now. Each consumer tries to maximize his expected life-time utility, subject to his budget constraints. The problem of consumer i is then given as

---

[106] The model is close to those used in the related literature, like Honkapohja, Mitra and Evans (2003) or Preston (2005a). For methodological background, see for example Woodford (2003), chapter 2 or Obstfeld and Rogoff (1996), chapters 2 and 5.

[107] The extent to which the agents are allowed to be heterogeneous will be defined more precisely below.

(37) $\max_{\{C_T^i\}} U^i = \hat{E}_t^i \sum_{T=t}^{\infty} \beta^{T-t} U(C_T^i)$

subject to the flow budget constraint

(38) $B_{t+1}^i = R_t (Y_t^i + B_t^i - C_t^i)$

and the transversality condition

(39) $\lim_{T \to \infty} \hat{E}_t^i \left\{ R_{t,T+1}{}^{-1} B_{T+1}^i \right\} \geq 0$

where $R_{t,T} \equiv \prod_{s=t}^{T-1} R_s$ and $R_{t,t} \equiv 1$.

The operator $\hat{E}_t^i$ describes consumer *i*'s expectation in period *t*, which is based on his subjective beliefs. The circumflex denotes that in this general notation the expectation may be not rational. β denotes the temporal discount factor which is assumed to be identical for all agents and to satisfy $0 < \beta < 1$. $U(\cdot)$ is the instantaneous utility function, assumed to be identical for all agents as well, to satisfy $U_C(\cdot) > 0$ and $U_{CC}(\cdot) < 0$ and to have a constant intertemporal elasticity of substitution.[108] C_t^i, Y_t^i, and B_t^i designate the consumer's consumption in period *t*, his income in *t*, which may be stochastic, and his net wealth in the form of bonds at the beginning of period *t,* respectively. Those bonds are riskless and the only available asset in the economy.[109] R_t is the gross one-period rate of return for this asset from period *t* to *t*+1 which is determined in period *t* by market clearing on the bonds market. The information set of the agents comprises information about past periods and the present one, but only expectations of future periods' variables.

---

[108] $U_C(\cdot)$ denotes the first derivative of *U* with respect to *C*, $U_{CC}(\cdot)$ the second derivative.

[109] Thus the model's asset markets are incomplete and consumers cannot insure against individual risks.

Since the model describes a closed economy, the aggregated net wealth in each period must be zero:

$$(40)\quad \sum_{i=1}^{n} B_t^i = 0 .$$

The flow budget constraint (38) must hold in all periods. Iterating it over the infinite horizon of the consumer and using the fact that in the optimum (39) holds as an equality gives the usual intertemporal budget constraint

$$(41)\quad \hat{E}_t^i \sum_{T=t}^{\infty} R_{t,T}{}^{-1} C_T^i = B_t^i + \hat{E}_t^i \sum_{T=t}^{\infty} R_{t,T}{}^{-1} Y_T^i .$$

In order to close the model it is necessary to specify the process driving the exogenous state variables, here the income Y_t^i. I will assume that each agent's income follows the process

$$(42)\quad Y_t^i = M^i \, Y_{t-1}^i{}^{\rho} e^{v_t^i}$$

with $|\rho| < 1$, $M^i > 0$ and where $\{ v_t^i \}$ is a bounded i.i.d. white noise process with zero mean which is similar for all agents.[110]

The consumer's problem can be solved by maximizing (37) subject to (38) and (39). This yields the stochastic Euler equation:[111]

$$(43)\quad \hat{E}_t^i \{ U_C(C_T^i) \} = \beta \hat{E}_t^i \{ R_T \, U_C(C_{T+1}^i) \} \quad \forall T \geq t .$$

This equation, albeit familiar, deserves some remarks. First, it is useful to clarify the interpretation of the expectations operator $\hat{E}_t^i$. As mentioned above, it denotes the agent's expectation, based on his subjective beliefs. Optimally, these beliefs are based

---

[110] A process like (42) is frequently used for convenience, because it results in a AR(1) process when log-linearized.

[111] The derivation can be done by dynamic programming.

on the complete set of information available to him in period *t*. If the operator is applied to future variables within the agent's influence like here to C_{t+1}^{i}, this set of information comprises the knowledge about his own behaviour. Accordingly, the expectation should be interpreted as the value which the agent expects he will choose as the optimum in the respective future period, given the information available to him in period *t*.[112] The fact that the variable is within his influence means that the agent knows according to what reasoning he will behave in period *t*+1, and he will use this knowledge to form his expectations in *t* as well as possible. It is not plausible if an agent has an expectation of his future consumption which is contradicted by the way how he actually determines his consumption in the current period. This is called "internally consistent" expectations by Preston.[113]

Furthermore, it is important to note that the Euler equation (43) determines the optimal consumption in *t* only if it holds for all periods $T \geq t$. This is crucial because C_{t}^{i}, determined according to (43), is only optimal on the condition that the expected value for C_{t+1}^{i} is indeed optimal as well, given the information in *t*. The expected optimal value for C_{t+1}^{i}, in turn, can only be determined by the Euler equation if the expected value for C_{t+2}^{i} is optimal, and so on. The Euler equation is constructed recursively, so it is actually an infinite set of conditions, each linking expected optimal consumption in one period to expected optimal consumption in the next one. This is an elegant way to express the optimality conditions despite the "infinite-horizon"-nature of the problem. However, when applying the Euler equation one should still bear in mind that each instance of (43) for a particular *T* is only part of solving the consumer's problem *simultaneously* for the entire future consumption path. In the example above this means that for every period *T* the entire future optimal consumption path is reflected in the expected optimal value for C_{T+1}^{i}.[114] So, if an optimizing agent uses the Euler equation to determine his optimal consumption in *t*, this only makes sense if his expectation of his

---

[112] This interpretation requires that the utility function is not subject to dynamic inconsistency, which is ensured by the utility function used in the example.

[113] See Preston (2005), p. 97.

[114] If one solves the dynamic optimization problem by dynamic programming, this property of the Euler equation becomes especially apparent in the use of the "value function".

consumption in the next period is as good as possible, given his information set in *t*. The Euler equation as an optimality condition therefore implies that the expectations operator used in (43) is interpreted as being "internally consistent" as described in the paragraph above.[115]

Moreover, the Euler equation is only a necessary condition which does not always ensure optimality on its own. It imposes an optimality condition on the consumption in one period relative to the values in other periods but not on the absolute level of the consumption. Consequently, with only the Euler equation as a condition, it may be possible for a consumer to spend more or less than his intertemporal budget allows. In order to guarantee optimality and especially feasibility of the consumption path, equation (43) must therefore be combined either with (38) and (39) or with (41). In other words, the Euler equation determines the *shape* of the consumption path, as expected in *t*, while the intertemporal budget constraint determines its *level*.

The main result from the reasoning above is the following: An agent's dynamic optimization problem similar to the one outlined here can generally be solved by making use of the intertemporal budget constraint (41) as well as the Euler equation (43) for all periods $T \geq t$. That means that an agent can find the optimal current consumption if he employs all relevant state variables – their actual values or their expectations for future periods – which are used in this equations and which are beyond the agent's influence. In the model above these comprise B_t^i as well as $\left\{ Y_T^i, R_T^i \middle| T \geq t \right\}$.

The considerations above describe how the optimal current consumption can be determined in a dynamic optimization environment in general without reference to a specific model, that is without regard to the concrete assumptions concerning the stochastic processes of the exogenous variables or the homogeneity or heterogeneity of the agents. However, that does not necessarily mean that an agent must always consider all of these equations and all of these variables explicitly in his decision. Depending on the actual model it may be possible to omit some of the equations or variables while still arriving at a REE solution. For example, if the expectations of Y_T^i and R_T^i are related to

115 Again, barring the case of a dynamically inconsistent utility function.

the variables of the current period in some stable way, it may still be possible to capture their influence by considering the influence of the variables of the current period and thus to determine the optimal current consumption by using just the current state variables. Whether this is indeed possible and whether the resulting solution is stable under learning depends on the way the actual model is constructed. In other words, it depends on the modeller's decisions about how to close the model in terms of expectations and additional assumptions, such as the way the agents form their expectations or the scope of the representative agent assumption. I will argue that in some model types, namely in rational expectations models with a representative agent, it is possible to use simplifying shortcuts when solving the dynamic optimization problem which cannot be used generally. As a consequence, some forms of learning will lead to rational expectations only when they are applied to those special kinds of models, while others can be used generally without restrictions. I will argue that this is at the heart of the debate about the differences between Euler equation learning and infinite horizon learning.

3.1.2. Steady State and Log-Linearization of the Model

A frequently used technique to solve a model like the one above is to transform it in a log-linear form. This makes it possible to use purely linear methods, and as a further benefit it allows a solution without having to specify the instantaneous utility function in detail. To arrive at the log-linear form, the model's variables are first substituted using the logarithm of their deviation from their steady-state value. The model is then linearized around the non-stochastic steady state by way of a Taylor expansion. The log-linear form is consequently an approximation to the original model, which can nevertheless be employed provided the model's variables remain in a neighbourhood of their steady-state values.[116] Throughout this work it will be assumed that the stochastic

[116] However, it may be questioned what consequences a log-linearization has for the notion of an expectation. It may make a difference whether an agent forms an expectation of a future absolute value or whether he forms an expectation of a future logarithmic deviation from a steady state, which he may not know. But this question is left aside here.

disturbances, like $\{v_t^i\}$ in the exogenous income process (42), are bounded in such a way that the system remains sufficiently near its steady state.[117]

The non-stochastic steady state of the income and the rate of return as implied by (42) and (43) are

$$(44)\quad \overline{Y}^i = M^{i\frac{1}{1-\rho}},\ \overline{R} = \beta^{-1}$$

where steady-state values are designated by an upper bar. The non-stochastic steady state of the consumer's net wealth and consumption, as implied by (38) and (39), would in a general notation – without assuming a representative agent – be

$$(45)\quad \overline{B}^i = B_t^i,\ \overline{C}^i = \overline{Y}^i + \frac{\overline{R}-1}{\overline{R}}\overline{B}^i = \overline{Y}^i + (1-\beta)\overline{B}^i\,.$$

Note that these steady state values depend on the agent's wealth in the current period *t*. That means that if the agent's wealth varies in some future period, the steady state, as seen from this future period, would be different. However, in order to make variables of different periods compatible, they all have to be log linearized around the same steady state. I define this "reference steady state" as the steady state of period *t*=0, so that $\overline{B}^i = B_0^i$ and the correspondent steady states of consumption will be used as linearization points for all periods.

The model will now be log-linearized at this steady state. In the following, lower case variables indicate the natural logarithm of the deviation from the steady state, for example $y_t^i \equiv \ln\left(Y_t^i / \overline{Y}^i\right)$ and so on, with the exception of the agent's wealth where the log deviation will be defined as $b_t^i \equiv \ln\left(\left[B_t^i - \overline{B}^i + \overline{C}^i\right] / \overline{C}^i\right)$.[118] The log-linearization of the income process (42) for the agent gives

[117] For further discussion, see Woodford (2003), chapter 2 and in particular appendix A.3.

[118] This is done to avoid a technical problem. Were the log deviation of wealth taken as $b_t^i \equiv \ln\left(B_t^i / \overline{B}\right)$ it would not be defined for a zero steady state wealth. The definition used here circumvents this problem

(46) $y_t^i = \rho\, y_{t-1}^i + v_t^i$.

Using a first-order Taylor approximation,[119] the stochastic Euler equation becomes

(47) $\hat{E}_t^i c_T^i = \hat{E}_t^i c_{T+1}^i - \sigma\, \hat{E}_t^i r_T \quad \forall T \geq t$ with $\sigma \equiv -\dfrac{U_C(\overline{C}^i)}{U_{CC}(\overline{C}^i)\overline{C}^i} > 0$.

Here, σ denotes the intertemporal elasticity of substitution at the steady state. Since the instantaneous utility function is assumed to exhibit a constant intertemporal elasticity of substitution, this value is independent from the agent's steady state consumption and consequently the same for all agents. Finally, log-linearizing the flow budget constraint (38) yields

(48) $b_{t+1}^i = \dfrac{\overline{B}^i}{\overline{C}^i} r_t + \beta^{-1} \dfrac{\overline{Y}^i}{\overline{C}^i} y_t^i + \beta^{-1} b_t^i - \beta^{-1} c_t^i$.

Log linearizing the model in this way gives rise to two technical difficulties. If individual agents are considered, it is frequently necessary to aggregate their individual variables over all agents within the economy. After the model has been log-linearized, each variable is expressed as a deviation from its steady state, which cannot be simply summed up over all agents. However, as it is shown in appendix A.2, the variables can nevertheless be aggregated by constructing the arithmetic mean of the individual log deviations, provided that the number of agents is sufficiently large and that the log

while having the log deviation b_t^i still centred around the steady state. For the special case of a representative agent wealth is *always* zero, so it could as well be omitted from all equations before the log linearization, but this is no longer possible when individual agents are considered. The definition used here is possible as long as $B_t^i > \overline{B}^i - \overline{C}^i$. It is assumed that this condition is covered by the requirement mentioned above which demands that the system remains sufficiently close to its steady state.

119 Because of Jensen's inequality, it would be more adequate to use a second-order Taylor approximation when linearizing an equation which contains an expected value of a (possibly) non-linear function like the Euler equation here. See for example Obstfeld and Rogoff (1996), p. 504. With a second-order approximation, equation (47) would be augmented by a constant term which depends on the variance of the random variables. This would result in a slightly different value for the equilibrium interest rate, but it would not qualitatively change the results this work focuses on. It is therefore neglected here. Another effect, that has been neglected here as well, can arise from the fact that both future consumption and future interest rates are random variables, so there can be a further influence from the covariance between those variables. See Mason and Wright (2001) for a detailed discussion.

deviation of the variable in question and its steady state are not correlated over the population of agents. Due to the definition of the log-linearized variables as percentage deviations from the respective steady state, and as long as the stochastic variables driving the model are distributed appropriately, this condition often holds.[120] Throughout this work it will be assumed that no log-linearized variable is correlated with the respective steady state value over the population of agents, so that the variables can be approximately aggregated by the average of their log deviations. Possible correlation between the log-linearized variables are, of course, not ruled out by this assumption. Similarly it is assumed that the initial beliefs of the agents are independently distributed over the population of agents and are not correlated with the log-linearized variables either.

Because of the log-linearization as it is used in this work, it will sometimes also be necessary to calculate the average ratio of the steady state wealth and the steady state consumption $n^{-1}\sum_{i=1}^{n}\left(\bar{B}^i/\bar{C}^i\right)$ in the course of an aggregation. The average of the steady state wealth over all agents in a closed economy must always be zero, but – depending on the distribution of wealth and consumption over the agents – this cannot be taken for granted for the ratio. To simplify the calculation it will be assumed in the following that wealth and consumption are distributed in such a way that the average of their ratio can nevertheless be approximated by zero. The average ratio of steady state income to steady state consumption, which must sometimes be calculated as well, is then

(49) $$n^{-1}\sum_{i=1}^{n}\left(\bar{Y}^i/\bar{C}^i\right)=1-\left(1-\beta\right)n^{-1}\sum_{i=1}^{n}\left(\bar{B}^i/\bar{C}^i\right)\approx 1 .$$

Note that these difficulties could both be circumvented by defining the log deviations differently when linearizing the model. In some works[121] the log-linearized variables of individual agents are defined as logarithmic deviation from the respective steady state

---

[120] For example, in (47) the consumption's deviation from the steady state is determined by the actual interest rate, which is common to all agents and can therefore not be a cause of correlation, and by the agent's expectation of next period's consumption *log-deviation*, which can normally be assumed to be not correlated to his steady state consumption *level*. For the agent's income (46) it is immediately clear that correlation can be ruled out because it is driven by the independent white noise process ν_t.

[121] For example Preston (2005a).

values of the *representative* agent. Alternatively, sometimes even several different variables are all defined as logarithmic deviations from the representative agent's steady state income. With these approaches the aggregation problems described above would not occur, but on the other hand it may be difficult to argue that the variables in question will always stay sufficiently close to the steady state value used as a linearization point, especially if the agents are heterogeneous. In contrast to that, the approach used here requires the additional assumptions concerning the distributions of the steady state variables over the population of agents as described above, thus rendering the results somewhat less general, but the advantage is that it can also be used if the individual agents' steady state values are very different from each other.[122]

3.1.3. The Canonical Solution: The Rational Expectations Equilibrium with a Representative Agent

In this section the canonical solution will be discussed, which serves as a reference for the variations in the following sections. For this solution rational expectations and the existence of a representative agent are assumed, which are widespread assumptions when solving models like the one discussed here. A usual approach is to assume that the economy consists of only one representative agent. This is the limiting case of the slightly less restrictive assumption that the economy consists of a number of agents with identical characteristics, facing identical exogenous variables and, in particular, having identical starting conditions in t. As a result, all agents will exhibit identical behaviour. In models using this assumption the economy's aggregated variables are usually given as per capita values. The population of such an economy can then be normalized to one, yielding the "single representative agent" economy.

Using a representative agent, the superscripted i can consequently be dropped from all variables in the model equations. Furthermore, in a closed economy with a

[122] The approach used here can easily be transformed into one of the described alternatives. If the log-linearized variables would be defined as log-deviations from the respective steady state values of a representative agent, then all superscripts of the steady state variables could simply be dropped, the steady state wealth would be zero and the steady state consumption would be equal to steady state income. That way the aggregates of the ratios mentioned in the text would always be zero respectively one.

representative agent, the agent can never possess any net wealth, so the steady state values in (45) simplify to[123]

(50) $\overline{B}=0$, $\overline{C}=\overline{Y}$.

The representative agent versions of the log-linearized income process (46) and the Euler equation (47) are accordingly

(51) $y_t=\rho y_{t-1}+v_t$,

(52) $\hat{E}_t c_T=\hat{E}_t c_{T+1}-\sigma\hat{E}_t r_T \quad \forall T\geq t$.

The last element necessary to solve the model is the market clearing condition $c_T=y_T, \quad \forall T\geq t$. The REE for the interest rate can now be obtained by taking the linearized Euler Equation for T=t and imposing the market clearing condition:

(53) $y_t=c_t=\hat{E}_t c_{t+1}-\sigma r_t$.

Given the current income, the interest rate thus depends on next period's expected consumption. The rational expectation for next period's consumption is determined by the rational expectation for the income in t+1 as implied by the process (51) and again the market clearing condition: [124]

(54) $E_t c_{t+1}=E_t y_{t+1}=\rho y_t$.

It follows that the REE for the interest rate is

(55) $r_t=-\sigma^{-1}(1-\rho)y_t$.

[123] Note that thus in the non-stochastic steady state market clearing is equivalent to the flow budget constraint in conjunction with the representative agent assumption.

[124] In order to indicate that these are rational expectations, here the expectations operator is written without a circumflex.

While the procedure outlined above may appear straightforward and innocuous, there are several points worth mentioning. First, note that the rational expectations for income and consumption only describe which possible expectations would be consistent with a REE. In this context, they do not yet imply whether and how these expectations can actually be attained by a dynamic process. Neither do they imply any reasoning by the agents or any process which could induce the agents to form these expectations. Apart from that, there are two possibilities of how the expectations of the representative agent can be interpreted. Taken literally, they can be thought to express the individual expectations of a single representative agent. However, since the representative agent is just a fiction to simplify the solution, what they actually describe are the aggregated expectations of the individual agents in the economy.

This distinction has significant consequences for the application of the market clearing condition. It may seem trivial, but it is important to note that this condition must of course be fulfilled for the economy as a whole but, in general, not for an individual agent. This means that it can be applied without any problem for the actual consumption as in (53), because this is after all the aggregated consumption of the whole economy, no matter whether it is normalized to a population of one or not. But when expectations are concerned as in (54), one must be more careful. Generally, expectations do not need to fulfil any market clearing condition at all. For *rational* expectations this is different: Expectations can only be rational if every agent perceives the resulting states of the economy as indistinguishable from his expectations, and this implies that the agent's aggregated expectations fulfil the appropriate market clearing condition. The market clearing condition therefore constitutes a necessary condition for expectations to be rational, although it must only be fulfilled on the aggregated, macroeconomic level. However, if the expectations are taken as the rational expectations of a *representative* agent, the market clearing condition must be satisfied by construction even on the individual agent's level: If a market consists only of one or several representative agents, which by definition are all identical, the market clearing condition implies that in the equilibrium their individual excess demand (or excess supply) is zero.[125] In the model above, for example, the market clearing condition implies that they are not able

[125] This is the well known feature that many equilibria are "no-trade equilibria" when representative agents are considered.

to borrow or to lend, so none of them can possess any wealth. Consequently, even their individual expectation of next period's consumption can only be rational if it is the same as the rational expectation of next period's representative income.[126]

However, this is true only for the special case of the representative agent. If the expectations are interpreted as the aggregation of individual expectations which are not necessarily representative, the market clearing condition only makes sense on the aggregated level. In the model discussed here, it merely requires that the aggregate of the expectations of next period's consumption matches the aggregate of ρ times this period's income.[127] But that does not mean that on the individual level an agent's expectations satisfies a similar requirement. In fact, as it will be shown in sections 3.2.2 and 3.2.3, there is a continuum of individual agents' expectations, even including sunspot solutions, which all satisfy the requirement only on the aggregated level. The reasoning above can be summed up as follows: Expectations in general do not need to satisfy any market clearing condition. *Rational* expectations must satisfy an appropriate market clearing condition on the aggregate level as a necessary condition. On the individual agent's level rational expectations must satisfy a market clearing condition only in the special case of a representative agent.

Whether the market clearing condition can be used on an individual level has consequences for the solution of a model like the one discussed here. The reasoning above has shown that for the individual (not-representative) agent's consumption expectations the market clearing condition is not binding. Being able to buy or sell bonds, he also does not have to match his expected (or actual) income with his expected (or actual) consumption in any period. But of course that does not mean that he can borrow and consume arbitrarily because he is still restricted by the intertemporal budget constraint (41). As it has been pointed out when the significance of the Euler equation was discussed in section 3.1.1, solving the individual agent's dynamic optimization problem requires in general to employ the Euler equation for all periods $T \geq t$ as well as the intertemporal budget constraint, using all state variables in these equations. Yet, in

---

[126] Of course it is still possible for a representative agent to have different consumption expectations, but these would not be rational.

[127] Otherwise, at least one of the underlying individual expectations cannot be rational.

contrast to this finding, in the rational expectations solution for the representative agent above, the intertemporal budget constraint was not used at all and the Euler equation was only used for the actual period. Omitting the intertemporal budget constraint from the solution and restricting the Euler equation to the actual period represent two shortcuts which simplify the solution considerably.[128] The solution is nevertheless correct and can indeed be derived like shown above.

The underlying reason why both simplifications are possible is that the representative agent assumption allows the market clearing condition to be applied to the individual agent's rational expectations. Basically, each period's Euler equation is linked to the next period's one by the expectation of next period's consumption. A rational expectation of next period's consumption can only be formed when the next period's Euler equation is taken into account, including the expectation for the following period's consumption, and so on to infinity. The dynamic optimization problem must thus be solved simultaneously for all periods. But if a representative agent is considered, the market clearing condition provides an alternative way to determine next period's expected consumption. The market clearing condition connects the consumption expectation to a state variable (i.e. the income) which is exogenous for the agent and beyond his influence. Since the expectation of such a state variable is independent from the agent's future behaviour, a partial solution for a "single period" Euler equation becomes possible. In this way the market clearing condition effectively makes the link between the Euler equations redundant by providing a peg for the expectation of next period's consumption.

Concerning the intertemporal budget constraint, a similar argument can be made. In the general solution to the optimization problem, the intertemporal budget constraint must be used to determine the level of the future consumption path, because the infinite chain of Euler equations can only determine its shape. When the market clearing condition relates the expected consumption to the (exogenous) income, the level of consumption is immediately determined, and the intertemporal budget constraint is made redundant. Again, this can only be done for the representative agent. To put it more technically, the

[128] When infinite horizon learning is demonstrated in section 3.3, it will also be shown how the REE solution can be derived without employing these simplification.

representative agent assumption together with the market clearing condition already imply that the flow budget constraint and the transversality condition are always respected (because a representative agent cannot borrow or lend), so the intertemporal budget constraint becomes redundant.

The main point derived from these considerations is the following: The above mentioned shortcuts are possible if the market clearing condition can be applied to individual expectations. But this in turn hinges on two conditions: That representative agents are considered, and that the expectations in question are assumed be rational.

3.1.4. The Rational Expectations Solution with Individual Heterogeneous Agents

The reasoning in the preceding section has shown that the representative agent assumption allows a simplified solution method which may not be feasible if individual, possibly heterogeneous agents were considered.[129] However, that does not necessarily mean that the solution for an aggregate variable (such as the interest rate) obtained by this method is not valid for the underlying problem with individual agents. In this section it will be considered in more detail how the solution of the model changes if the representative agent assumption is dropped. In particular, the differences between the consumption of a representative agent and the consumption which would be obtained for a population of individual agents are examined, and it will be discussed how the individual agent solution should be interpreted.

In order to determine the significance of the representative agent assumption, one again has to distinguish between the aggregate, macroeconomic level and the individual, microeconomic level. It will turn out that on the aggregate level, at least for the present model, the representative agent solution is the same as the solution for a group of individual agents. This is not surprising because the representative agent, by definition, represents the aggregate behaviour of a whole class of agents. The representative agent is essentially a macroeconomic concept, even if his behaviour is described by

[129] The term “heterogeneous” is meant here only in the limited sense that agents may exhibit differences in their state variables like income or wealth, but not in their parameters concerning the utility functions or their time preference. See also the description of the agents’ optimizing problem in section 3.1.1.

microeconomic functions. Applying the representative agent assumption to a model thus means that the individual, possibly heterogeneous agents' behaviour is aggregated into one or several fictitious homogeneous representative agents, and it is plausible that in a linearized model this does not make any difference for the economy's aggregate variables. Indeed, in the model discussed here the REE solution for the interest rate would be the same if individual agents would be considered instead of a representative agent. At the same time this aggregation means that the agents' individual differences, like differences in initial wealth or in income, are ignored when the representative agent assumption is used. The consequence is that on the microeconomic, individual level the representative agent solution cannot be used to describe an individual agent's consumption.

These considerations will now be demonstrated formally by deriving the conditions for a REE solution without resorting directly to the representative agent assumption. In contrast to the previous section the Euler equation and the income process are now allowed to be heterogeneous

$$(56)\quad c_t^i = \hat{E}_t^i c_{t+1}^i - \sigma r_t ,$$

$$(57)\quad y_i^i = \rho y_{t-1}^i + v_t^i .$$

In order to determine the interest rate, equation (56) must be aggregated over all n agents in the economy. As pointed out in section 3.1.2, this is done by constructing the arithmetic mean of the individual log deviations. Thus, averaging over all n agents in the economy and applying the market clearing condition to the average consumption and average income gives

$$(58)\quad r_t = \sigma^{-1}\left(\frac{1}{n}\sum_{i=1}^{n} \hat{E}_t^i c_{t+1}^i - \frac{1}{n}\sum_{i=1}^{n} y_t^i \right).$$

Now, if the expectations of next period's consumption are rational for all agents, its sum must be equal to the sum of the rational expectation of next period's individual income, which is $E_t^i y_{t+1}^i = \rho y_t^i$. This step is again nothing else but applying the market clearing condition to expectations, which is unproblematic here, because it is done only on the aggregate level and because rational expectations are considered. Using the rational expectation of each agent's income and aggregating over all agents we get the REE solution for the interest rate as

$$(59)\quad r_t = \sigma^{-1}\left(\frac{1}{n}\sum_{i=1}^{n} E_t^i y_{t+1}^i - \frac{1}{n}\sum_{i=1}^{n} y_t^i\right) = \sigma^{-1}(\rho-1)\frac{1}{n}\sum_{i=1}^{n} y_t^i = \sigma^{-1}(\rho-1)\tilde{y}_t$$

with $\tilde{y}_t \equiv n^{-1}\sum_{i=1}^{n} y_t^i$ being the average income.

If all agents had the same income, as it would be the case with a representative agent, then $\tilde{y}_t = y_t^i$, and then this solution is the same as (55). Inserting (59) into the individual agent's linearized Euler equation yields the REE for the agent's current consumption as

$$(60)\quad c_t^i = E_t^i c_{t+1}^i + (1-\rho)\tilde{y}_t .$$

If the economy consisted of representative agents one could use the market clearing condition on their individual expectations, so this would simply amount to $c_t^i = y_t^i$, which for a representative agent would be required anyway by the market clearing condition. But if individual agents are considered, this is not possible. Instead, one has to use the Euler equation for all future periods, calculating the rational expectation for the income, interest rate and consumption in each future period up to infinity. Since $|\rho| < 1$ this yields[130]

$$\begin{aligned} c_t^i &= \lim_{s\to\infty}\left[E_t^i c_{s+1}^i + \sum_{T=t}^{s}(1-\rho)\tilde{y}_T\right] \\ (61)\qquad &= \lim_{s\to\infty}\left[E_t^i c_{s+1}^i + (1-\rho)\sum_{T=t}^{s}\rho^{T-t}\tilde{y}_t\right] \\ &= \lim_{s\to\infty}\left[E_t^i c_{s+1}^i\right] + \tilde{y}_t . \end{aligned}$$

[130] This requires that the law of iterated expectation holds, which for rational expectations must be the case.

The current consumption is thus determined by the average current income and the expectation of consumption in the infinite future. Note that, in general, this expectation does not need to be zero. Again, if the economy consisted of representative agents one could use the market clearing condition on this expectation, so it could be replaced by the rational expectation of income in the infinite future, which would indeed be zero. Since in this case all agents had the same income, this solution again amounts to $c_t^i = y_t^i$, the same as for the representative agent.

But with individual, not representative agents the situation is different. Current consumption is deviated from its steady state because of an income shock in the economy. If this income shock is of different magnitude for different agents, then they react by borrowing or lending to each other, accumulating individual wealth or debt. This immediately causes a shift in the individual consumption's steady state, which in turn depends on an agent's wealth,[131] because in future the agent pays or receives interest on his debt or wealth. But since the system was log linearized with respect to a reference steady state in the past that means that in the infinite future, when the short term effects of the shock have died down and the new steady state is reached, this shift in the steady state is represented by a permanent log-deviation of consumption from the reference steady state. The term $\lim_{s\to\infty}\left[E_t^i c_{s+1}^i\right]$ is thus the rational expectation of the shift in the agent's steady state. The amount of the shift depends on the rationally expected wealth or debt the agent accumulates over the future periods, and it depends therefore on the agent's individual income and the average income. The shift would only be zero if the income shock would have been the same[132] for every agent so that no lending would take place.

To calculate a closed form solution for the current consumption it would be necessary to determine the effect on the steady state, which in turn would require to use the

[131] As shown in equation (45)

[132] Note that we are considering log deviations, so that in absolute values the income shock would have to be proportional to each agent's steady state. Strictly speaking, the log deviation would have to be the same for all agents only if in the period which serves as a reference for the log linearization no agent possesses any wealth and consumption is equal to income in the steady state. Otherwise, in order to prevent any lending, the income shock would have to be adapted for each agent to the difference between his income and his consumption which results from the accumulated wealth.

intertemporal budget constraint.[133] As the solution will emerge anyway when infinite horizon learning is discussed in section 3.3, this calculation will not be derived here in detail but merely presented.[134] The resulting solution under rational expectations for individual consumption is

$$(62)\quad c_t^i = (1-\beta)b_t^i + \frac{\overline{Y}^i - \overline{C}^i}{\overline{C}^i}\left[\lambda y_t^i + (1-\lambda)\sigma^{-1}\widetilde{y}_t\right] + \lambda\left(y_t^i - \widetilde{y}_t\right) + \widetilde{y}_t$$

$$\text{with } \lambda = \frac{\overline{R}-1}{\left(\overline{R}-1\right)+\left(1-\rho\right)},\ 0 < \lambda < 1.$$

It is straightforward to show that if an agent's consumption can be described by this equation, then it satisfies the Euler equation under rational expectations.[135] The agent's expectation is thus "internally consistent" in the sense that his expectation of his own behaviour is not contradicted by the decision rule he will actually use in the next period to determine his consumption.[136]

Comparing equation (62) with (61) implies that the expected permanent shift in the steady state of consumption, expressed as log deviation from the current steady state, is

$$(63)\quad \lim_{s\to\infty}\left[E_t^i c_{s+1}^i\right] = (1-\beta)b_t^i + \frac{\overline{Y}^i - \overline{C}^i}{\overline{C}^i}\left[\lambda y_t^i + (1-\lambda)\sigma^{-1}\widetilde{y}_t\right] + \lambda\left(y_t^i - \widetilde{y}_t\right).$$

Equations (56) to (63) hold not only for period t but also for all later periods. Note that nevertheless all steady states and consequently all log linearized variables are defined with respect to the beginning of period $t = 0$. This is relevant because for an individual agent, in contrast to a representative agent, the steady state of wealth and consumption changes when an income shock occurs. To interpret equation (63) assume that a single income shock takes place in period t and that no additional shocks occur after that period. The permanent shift is expected immediately when the shock occurs, and under

[133] This again underlines that for the solution to the individual agent's dynamic optimization problem, it is indispensable to include information about all future periods as well as the intertemporal budget constraint, as pointed out in section 3.1.1.

[134] A detailed derivation can also be found in Appendix A.5.

[135] See appendix A.3.1.

[136] See the discussion in section 3.1.1, p. 49.

rational expectations it comprises all influences on the steady state of consumption which will accumulate over the future periods. Therefore it should not change while the shock settles down. Indeed, it can be shown that in the absence of further additional shocks (63) remains constant over time as individual and average income approach their steady state values again.[137]

The term $(1-\beta)b_t^i$ represents the additional interest payments the agent receives (or must pay) for the wealth (debt) he has accumulated up to the current period while the income shock dies down. In the period of the shock it is zero because in this period the agent's wealth has not yet changed with respect to the steady state. All other terms on the right hand side of (63) depend on individual or average income and represent the sum of expected additional interest payments in future periods. While the agent's wealth or debt increases as he lends or borrows each period, these terms decrease from period to period with the rate ρ. The resulting change in the agent's interest payments $(1-\beta)b_t^i$ and the decrease in the expected interest payments (the income terms of (63)) thus cancel out each other.

For an illustration, assume that the agent has no wealth at the beginning of period *t*, so that $\overline{C}^i = \overline{Y}^i$ and $\overline{B}^i = 0$. Then the income terms of (62) simplify to a weighted average of the individual and the average income:

$$(64)\quad c_t^i = (1-\beta)b_t^i + \lambda\left(y_t^i - \tilde{y}_t\right) + \tilde{y}_t .$$

The amount lent (or, if negative, borrowed) by the agent in one period is then

$$(65)\quad y_t^i - c_t^i + (1-\beta)b_t^i = (1-\lambda)\left(y_t^i - \tilde{y}_t\right).$$

Inserting into (64) and recognizing that $\lambda(1-\lambda)^{-1} = \left(\overline{R}-1\right)(1-\rho)^{-1}$ yields

$$(66)\quad c_t^i = (1-\beta)b_t^i + \left(\overline{R}-1\right)\frac{1}{1-\rho}\left(y_t^i - c_t^i + (1-\beta)b_t^i\right) + \tilde{y}_t .$$

[137] This is shown formally in appendix A.3.2.

Under rational expectations, the individual and the average income decrease over time with the rate ρ, which means that the value determined in (65), the amount lent (or borrowed) by the agent, decreases with this rate as well. Thus, for an arbitrary future period T we have

$$(67)\quad E_t\left\{y_T^i - c_T^i + (1-\beta)b_T^i\right\} = \rho^{T-t}\left(y_t^i - c_t^i + (1-\beta)b_t^i\right).$$

Hence equation (66) can alternatively be written as

$$(68)\quad \begin{aligned} c_t^i &= (1-\beta)b_t^i + \left(\bar{R}-1\right)\frac{1}{1-\rho}\left(y_t^i - c_t^i + (1-\beta)b_t^i\right) + \tilde{y}_t \\ &= (1-\beta)b_t^i + \left(\bar{R}-1\right)\sum_{T=t}^{\infty} E_t\left\{y_T^i - c_T^i + (1-\beta)b_T^i\right\} + \tilde{y}_t. \end{aligned}$$

From this notation it becomes apparent that the second term on the right hand side represents the sum of all expected interest payments which will eventually take place because of additional lending or borrowing in future periods.

If the agent possesses some wealth or debt at the beginning of the period t, whose steady state serves as linearization point, then the solution is augmented by two additional effects, represented by the terms in square brackets on the right hand side of equation (62). The first effect is that with initial wealth steady state income is not equal to steady state consumption, therefore log deviations of consumption and income are expressions on different scales. The first term accounts for the necessary correction. The second term indicates how the temporary change of the interest rate, caused by the income shock, changes the interest payments which are received (or paid) by the agent on his initial wealth (debt). For a representative agent, however, it can immediately be seen from (62) that all described effects are cancelled out, because he cannot possess any wealth and his income is identical to average income.

Summing up, it turns out that for the REE solution of the interest rate, this being a variable on the aggregate level, it is irrelevant whether one considers individual agents or a representative one. The underlying reason is that in a linear system like the one considered here a representative agent can be defined by simply taking the average of the individual agents, so that it does not make any difference whether the individual agents' equations are aggregated first and solved for the aggregate variable afterwards, or whether the system is solved for the representative agent from the beginning. However, being an aggregation, the representative agent assumption cannot be used in general if one intends to determine individual variables of an agent, like in this model his consumption, although the representative agent solution is still valid for individual agents in the special case that each agent has identical income and identical (zero) wealth.

The solutions which have been discussed in this section are REE solutions, that is they are solutions which solve the agents' optimization problems under rational expectations. But this means merely that the consumption rules and the resulting interest rate, together with their expectations, would be consistent with a REE. However, it gives no answer to the question whether and how these REE solutions could be the outcome of the agents' behaviour. The basic concept of rational expectations, as it was initially formulated by Muth and others, was motivated by the idea that some dynamic process could lead to the elimination of irrational expectations, leaving only rational ones. In this view an REE is not just a theoretical construct but more of a natural equilibrium towards which an economy could be driven by the behaviour of its agents. Such a behaviour can be realized in the form of learning. It is therefore of interest whether an REE can only be posited as a possible equilibrium or whether it is also learnable by the agents. In the following sections this will be examined for different approaches of RLS learning, with a special focus on possible differences between the learning process of a representative agent and the learning processes of individual agents.

3.2. Euler Equation Learning

As described in chapter 2.4, adaptive learning is a possible way to explain how rational expectations could be attained. Learning describes the process of how expectations, being at first not rational, evolve over time, possibly converging towards the REE solution. When rational expectations are replaced by learning, the expectation in question is substituted by a function which derives the expectation from state variables within the bounded information set of the agent. The process of learning consists of revising the parameters of this expectation function each period as additional information arrives.

When learning is introduced in a dynamic optimization model, this is usually done in the form of "Euler equation learning" (EE learning in the following). An alternative method is "infinite horizon learning" (IH learning) which was established by Bruce Preston.[138] Presently, the majority of articles and papers uses EE learning. In this section EE learning will be demonstrated for the case of a representative agent. Next, it will be shown how this method can lead to a continuum of non-rational solutions and in some cases can even lead to sunspot solutions when individual agents are considered. The results will then be discussed in section 3.2.4.

3.2.1. EE Learning with a Representative Agent

Consider again the case of a representative agent who is trying to find the optimal value for his current consumption. In section 3.1.3 the solution was determined for the case that the agent has rational expectations about the future. In this section the agent will instead use a PLM to form his expectations. By revising the parameters of the PLM he is trying to improve his expectations. The question is whether such a learning process will converge towards rational expectations or not.[139]

The learning process used in this section is called "Euler equation learning" because the agent uses the Euler equation for the current period (53) as a basis to determine his optimal consumption, which is repeated here for convenience:

$$(69)\quad c_t = \hat{E}_t c_{t+1} - \sigma r_t \,.$$

The interest rate of the current period is determined by applying the market clearing condition $y_t = c_t$, so that it depends of the current income and of the expectation for next period's consumption:

---

[138] See for example Preston (2005) or Preston (2005a). The terms "Euler equation learning" and "infinite horizon learning" were coined in Honkapohja, Mitra and Evans (2003).

[139] A similar system was examined in Honkapohja, Mitra, and Evans (2003).

(70) $r_t = \sigma^{-1}\left(\hat{E}_t c_{t+1} - y_t\right).$

The PLM which is examined here is linear and has the form

(71) $c_t = m + l y_{t-1} + e\varepsilon_{t-1} + \upsilon_t$

where *m*, *l*, *e* are the parameters which represent the agent's beliefs about the world he lives in and υ_t is a zero mean error term. Besides nesting the REE solution (54) the PLM contains a constant and a extraneous sunspot variable ε_t and is thereby overparameterized.[140] This allows to test whether the learning system is strongly E-stable. To test for weak E-stability it suffices to fix *m* and *e* to zero, so that the resulting PLM contains no more variables than the REE solution. The sunspot variable is common knowledge to all agents and is assumed to follow the AR(1) process

(72) $\varepsilon_t = \mu\varepsilon_{t-1} + \pi_t$

where $|\mu| \leq 1$ and $\{\pi_t\}$ is a bounded i.i.d. white noise process with zero mean. The agent's expectations is formed according to this PLM

(73) $\hat{E}_t c_{t+1} = m + l y_t + e\varepsilon_t$

so that the resulting interest rate is

(74) $r_t = \sigma^{-1}\left(m + (l-1)y_t + e\varepsilon_t\right).$

Since we are considering a representative agent, this interest rate will always induce the agent to consume exactly his income (the market clearing condition is fulfilled on the individual level), so that the ALM is simply

(75) $c_t = y_t = \rho y_{t-1} + \nu_t$.

---

[140] For the definition of overparameterization, see p. 32.

The T-mapping from the PLM to the ALM is then

$$(76)\quad T\begin{pmatrix} m \\ l \\ e \end{pmatrix} = \begin{pmatrix} 0 \\ \rho \\ 0 \end{pmatrix}$$

where the only fixed point at $m = e = 0, l = \rho$ is the REE solution. Convergence under adaptive learning can be shown by determining E-stability of the system, which is determined by the asymptotic stability of the ODE

$$(77)\quad \frac{d}{d\tau}\begin{pmatrix} m \\ l \\ e \end{pmatrix} = T\begin{pmatrix} m \\ l \\ e \end{pmatrix} - \begin{pmatrix} m \\ l \\ e \end{pmatrix} = \begin{pmatrix} -m \\ \rho - l \\ -e \end{pmatrix}.$$

The derivatives of the right hand side are all negative, so the equilibrium at the fixed point is globally asymptotically stable. The solution is therefore strongly E-stable. For this model class E-stability implies stability under learning.[141] Thus, under learning the system converges globally towards the REE solution with probability 1, even if the PLM contains a sunspot variable.

3.2.2. EE Learning with Individual Agents

In order to find out whether the results of the previous section are dependent on the representative agent assumption, the same setting will now be examined for the case of individual agents. For simplification of the analysis, it is at first assumed that the income shocks are similar[142] for all agents. Again each agent uses the Euler equation to determine the optimal current consumption, now having individual expectations. The derivation of the REE solution for individual agents in section 3.1.4 has shown that according to (58) the interest rate of the current period is

---

141 See Evans and Honkapohja (2001), section 8.3.2.

142 “Similar” here should be understood in the sense that the shocks are proportional to each agent’s steady states before the log linearization, which means that they are identical for all agents after log linearization.

$$(78)\quad r_t = \sigma^{-1}\left(\frac{1}{n}\sum_{i=1}^{n}\hat{E}_t^i c_{t+1}^i - y_t\right).$$

The parameters of the PLM can now be different for each agent:

$$(79)\quad c_t^i = m^i + l^i y_{t-1} + e^i \varepsilon_{t-1} + \upsilon_t^i .$$

If each agent uses his PLM to form his expectations we get

$$(80)\quad \hat{E}_t^i c_{t+1}^i = m^i + l^i y_t + e^i \varepsilon_t$$

and the resulting interest rate is

$$(81)\quad r_t = \sigma^{-1}\left(\tilde{m} + (\tilde{l} - 1) y_t + \tilde{e}\varepsilon_t\right)$$

where parameters with a tilde (~) denote the average of the individual agent's parameters. In contrast to the previous section, the market clearing condition does not hold any more on the individual level, and the ALM for an individual agent is now:

$$(82)\quad c_t^i = \left(m^i - \tilde{m}\right) + \left(l^i - \tilde{l} + 1\right)\left(\rho\, y_{t-1} + \nu_t\right) + \left(e^i - \tilde{e}\right)\left(\mu\,\varepsilon_{t-1} + \pi_t\right).$$

The *T*-mapping from the PLM to the ALM is

$$(83)\quad T\begin{pmatrix} m^i \\ l^i \\ e^i \end{pmatrix} = \begin{pmatrix} m^i - \tilde{m} \\ \left(l^i - \tilde{l} + 1\right)\rho \\ \left(e^i - \tilde{e}\right)\mu \end{pmatrix}$$

and for the average parameters we have

$$(84)\quad T\begin{pmatrix} \tilde{m} \\ \tilde{l} \\ \tilde{e} \end{pmatrix} = \begin{pmatrix} 0 \\ \rho \\ 0 \end{pmatrix},$$

the same as for the representative agent. There is now a continuum, indexed by m^i and possibly e^i, of fixed points which can be described by the conditions

$$\text{(85)}\quad \begin{aligned} &\tilde{m} = 0 \\ &l^i = \tilde{l} = \rho \\ &\begin{cases} e^i = \tilde{e} = 0 & \text{if } \mu < 1 \\ \tilde{e} = 0 & \text{if } \mu = 1. \end{cases} \end{aligned}$$

Accordingly a continuum of self confirming equilibria exists. Although every agent must learn the right value ρ for the parameter l^i, it does not matter what parameter he learns for m^i, as long as its average over all agents is zero. The same holds for e^i, but only if the sunspot variable follows a unit root process ($\mu = 1$), otherwise the agent must learn the correct value, zero. In the present model, a sunspot variable may therefore only be part of the PLM in an equilibrium if it follows a unit root process.

To examine the E-stability of these equilibria it is convenient to split (83) into its three components. The dynamic of the system stems from the interaction between the individual agents and the average of all agents. E-stability is determined by the asymptotic stability of the ODEs for each parameter and its average[143]

$$\text{(86)}\quad \frac{d}{d\tau}\begin{pmatrix} m^i \\ \tilde{m} \end{pmatrix} = \begin{pmatrix} -\tilde{m} \\ -\tilde{m} \end{pmatrix},$$

$$\text{(87)}\quad \frac{d}{d\tau}\begin{pmatrix} l^i \\ \tilde{l} \end{pmatrix} = \begin{pmatrix} \rho\left(l^i - \tilde{l}\right) + \rho - l^i \\ \rho - \tilde{l} \end{pmatrix},$$

$$\text{(88)}\quad \frac{d}{d\tau}\begin{pmatrix} e^i \\ \tilde{e} \end{pmatrix} = \begin{pmatrix} \mu\left(e^i - \tilde{e}\right) - e^i \\ -\tilde{e} \end{pmatrix}.$$

The dynamic of the average parameters is independent from the individual ones. As in the case of the representative agent, all average parameters are globally asymptotically stable at their fixed points. Inserting their equilibrium values into the ODEs for the

[143] Instead of looking at one agent and the average, an alternative way to examine the E-stability would be to determine the stability for n agents simultaneously. This approach is shown in appendix A.4.

individual parameters shows that the derivative of the ODE for l^i on the right hand side of (87) is negative, so the equilibrium value for this parameter is globally asymptotically stable as well. The same holds for e^i as long as $\mu = 1$. The derivative of the ODE for m^i and for e^i in the case of $\mu = 1$ are zero. This means that the equilibrium is stable, but not attractive and therefore not asymptotically stable. If the parameters are at any point within the equilibrium continuum, then the dynamics of the ODE do not drive the parameters away from this point, but if they are pushed out if it by the learning algorithm, as a reaction to exogenous stochastic influences, they do not necessarily return to the same equilibrium values. Thus, they do not converge towards a particular equilibrium, but are moved around within the equilibrium continuum by the influence of the exogenous stochastic variables.[144]

If the ALM (82) was averaged over all agents it would simply be $\tilde{c}_t = y_t$, as it must be the case because of market clearing. For the individual agent the ALM at an equilibrium point is

$$(89)\quad c_t^i = m^i + y_t + e^i \varepsilon_t .$$

Compare this to the rational expectation solution, which can be obtained by adapting (62) to the case of common income shocks:

$$(90)\quad c_t^i = (1-\beta) b_t^i + \frac{\overline{Y}^i - \overline{C}^i}{\overline{C}^i}\left[\lambda + (1-\lambda)\sigma^{-1}\right] y_t + y_t .$$

Consider first the case that in t=0 (the reference period for the log linearization) the agent has no wealth so that $\overline{Y}^i = \overline{C}^i$. If we fix m^i and e^i to zero to rule out overparameterization, then the ALM is the same as the REE solution: The agent consumes always exactly his income and his wealth will always remain zero ($b_t^i = 0$). But if we allow for overparameterization, then the values for m^i (and possibly e^i) can be

144 For an unbounded continuum of solutions, equivalence of E-stability and stability under least squares learning could not be shown so far because the SRA convergence theorems do not apply to such cases. However, simulations suggest that the equivalence holds nevertheless. See Evans and Honkapohja (2001), p. 192.

different from zero so that the agent may consume too much or too little. Further, if the agent already holds some wealth in t=0, then his consumption under learning is never optimal in any case, no matter if overparameterization is ruled out or not.[145] In periods during the learning process (with or without overparameterization) where an equilibrium is not yet reached, the agents' consumption will be determined according to the ALM (82). In those periods individual consumption is also too high or too low, so that the agents will accumulate wealth or debt, which again will lead to a non-optimal consumption even when the equilibrium is reached eventually. In any of these cases, the agents do not react to over- or underconsumption because they only consider their Euler equation, but not their intertemporal budget constraint. In an equilibrium, the Euler-equation based PLM is confirmed by the ALM, even if it violates the intertemporal budget constraint and is therefore not optimal.

However, if some equilibrium is reached by the learning process, then the resulting interest rate is nevertheless the same as under rational expectations, although the individual expectations lead the agents to choose not optimal values for their consumption. This can be seen by inserting the average equilibrium parameter values into (81). The reason is that the average parameters are still learned correctly and that the differences between the learned individual parameters and the "right" parameters cancel each other out in the aggregate.

So, if individual agents with similar income shocks are considered, and if the PLM is not overparameterized and contains only the common income variable, then the learning process converges towards the global equilibrium with probability 1. But if a constant is included, the learning process converges not to a single point any more, but only ensures that the agent's PLMs are such that they are rational in the average, while still allowing for individual non-rationality. The same result obtains if a sunspot variable following a unit root AR(1) process is part of the PLM. Thus, the solution is weakly E-stable, but not strongly any more. Moreover, the solution is different from the solution under rational expectations, except for the particular case in which the agent has no initial wealth, does not accumulate any wealth during the learning process and has a PLM which is not overparameterized. In all other cases the agent's consumption

145 This case is discussed in Preston (2005a), section 2.3.

will be permanently too high or too low. Nevertheless, the interest rate follows the same law of motion as it does under rational expectations.

3.2.3. EE Learning with Individual Agents and Heterogeneous Income Shocks

The analysis of the previous section will now be repeated for the case that the income shocks are no longer common to all agents. Each agent's income now varies independently from the other agents' income, according to equation (57). The rational expectations solution for this situation was already discussed in section 3.1.4., where it was shown that interest rate of the current period is

$$(91)\quad r_t = \sigma^{-1}\left(\frac{1}{n}\sum_{i=1}^{n}\hat{E}_t^i c_{t+1}^i - \tilde{y}_t\right).$$

The agent's PLM is constructed analogous to the PLM of a representative agent, but focuses now on his individual income, just as his expectation:

$$(92)\quad c_t^i = m^i + l^i y_{t-1}^i + e^i \varepsilon_{t-1} + \upsilon_t^i,$$

$$(93)\quad \hat{E}_t^i c_{t+1}^i = m^i + l^i y_t^i + e^i \varepsilon_t.$$

The resulting interest rate is

$$(94)\quad r_t = \sigma^{-1}\left(\tilde{m} + \frac{1}{n}\sum_{i=1}^{n}\left(l^i - 1\right)y_t^i + \tilde{e}\varepsilon_t\right).$$

and the ALM for an individual agent yields

$$(95)\quad \begin{aligned} c_t^i &= \left(m^i - \tilde{m}\right) + \left(e^i - \tilde{e}\right)\left(\mu\varepsilon_{t-1} + \pi_t\right) \\ &+ \left(l^i - n^{-1}\left(l^i - 1\right)\right)\left(\rho y_{t-1}^i + \nu_t^i\right) \\ &- n^{-1}\sum_{\substack{j=1 \\ j\neq i}}^{n}\left(l^j - 1\right)\left(\rho y_{t-1}^j + \nu_t^j\right). \end{aligned}$$

The T-mapping from the PLM to the ALM is for each agent

$$(96)\quad T\begin{pmatrix} m^i \\ l^i \\ e^i \end{pmatrix} = \begin{pmatrix} m^i - \widetilde{m} \\ \left(\left(1-n^{-1}\right)l^i + n^{-1}\right)\rho \\ \left(e^i - \widetilde{e}\right)\mu \end{pmatrix}.$$

Note that in this approach the agent can only use his own income to learn. Therefore, the influence of the other agents' income on his consumption via the interest rate is just noise in his learning process. The T-map resembles the one of the previous section for common income shocks. Accordingly the continuum of equilibria, which is indexed by m^i and e^i and which occurs with an overparameterized PLM, is similar to the case of common income shocks and has the same stability properties. Only the equilibrium value for the parameter l^i is different. The fixed point for this parameter is

$$(97)\quad l^i = \frac{\rho}{n(1-\rho)+\rho}, \quad \lim_{n\to\infty} l^i = 0$$

and E-stability is determined by the ODE

$$(98)\quad \frac{d}{d\tau} l^i = \left(\rho - \rho n^{-1} - 1\right) l^i + \rho n^{-1}.$$

This ODE is globally asymptotically stable at the fixed point, so E-stability obtains. Intuitively, this equilibrium value for the parameter l^i highlights the fact that with a PLM based merely on the Euler equation, an agent's income determines his consumption only via the interest rate. The more agents exist in the economy, the less influence each agent's individual income has on the interest rate, and the lower the equilibrium value for l^i which corresponds to this influence. In the limit with an infinite number of agents, each agent learns that his own income does not matter for the interest rate and, since his behaviour is based only on the Euler equation, it does not matter for his consumption either. He will therefore exclude it from his PLM, learn some equilibrium values for m^i and e^i instead and determine his consumption accordingly.

Inserting the equilibrium parameters into the ALM (95) yields the individual consumption in an economy with many agents as

$$(99)\quad c_t^i = m^i + \tilde{y}_t + e^i \varepsilon_t$$

which is similar to the case with a common income shock which was examined in the previous section. Note that in such an equilibrium the interest rate is not the same as under rational expectations any more. Inserting the equilibrium values into (94) yields $r_t = -\sigma^{-1}\tilde{y}_t$ for the interest rate. The reason is that in an equilibrium the average consumption expectation is zero, because all agents have eliminated the influence of their own income from their PLM.

Fixing m^i and e^i to zero yields the ALM which would be learned if the PLM was not overparameterized. In this case each agent would simply consume the average income in each period. In contrast to that, the solution under rational expectations which was discussed in section 3.1.4 is

$$(100)\quad c_t^i = (1-\beta)b_t^i + \frac{\overline{Y}^i - \overline{C}^i}{\overline{C}^i}\left[\lambda y_t^i + (1-\lambda)\sigma^{-1}\tilde{y}_t\right] + \lambda\left(y_t^i - \tilde{y}_t\right) + \tilde{y}_t .$$

It is evident that this solution could not be learned with the PLM used in this section. Even if the influence of a constant and a sunspot variable on the agent's expectations is ruled out and even if the agent has no initial wealth and does not accumulate any wealth during the learning process either, the resulting consumption is still different from the consumption under rational expectations. The agent is not able to learn the optimal solution under rational expectations, and in the equilibrium he disregards his individual income completely. As in the case of common income shocks with an overparameterized PLM in the previous section, the learning process does not detect any resulting over- or underconsumption of the agent and cannot react to it, because it is based only on the Euler equation.

3.2.4. Discussion of EE Learning

In the preceding sections EE learning was examined for a representative agent as well as for individual agents. In the literature learning is usually modelled with a representative agent. However, while the representative agent assumption may be suitable to determine the REE solution for the aggregate variables, this is different if the expectations themselves and learning are considered. As it has been demonstrated in section 3.1.4, in a linear model like this the representative agent assumption does not change the REE solutions for the interest rate as an aggregate variable (even though the solutions for the individual variables are different), because it does not matter for the aggregate how diverse the individual agents are. If a representative agent is interpreted as the average of the underlying class of agents, then the consumption or the income of a representative agent can be reasonably interpreted as the average consumption or the average income of a population. But such an interpretation is not suitable any more if rational expectations and learning are considered. The idea behind rational expectations is that agents could not afford to form non-rational expectations in the long run. Yet, if the beliefs and the expectations of a representative agent are rational, this does only mean that the average of the individual agents' beliefs and expectations is confirmed by the actual outcomes, while the individual ones could still be non-rational. This contradicts the actual idea of rational expectations. Similarly, the aim of the introduction of learning into a model is to find a possible explanation how a rational expectations equilibrium could be the result of optimizing behaviour by the agents. A learning process of a representative agent, which would drive his beliefs towards rational expectations, could be interpreted as the average of the individual agents' beliefs and their evolution through learning. In this case it would be perfectly possible that the average learning process would converge towards rational expectations while allowing the individual beliefs to remain non-rational, as long as their irrationality cancels out in the average.

An alternative interpretation of a learning representative agent would be to assume that all individual agents in the economy are alike, with the same income shocks and the same initial beliefs, thereby following the same learning path. But if one applies such a strong assumption to the initial beliefs of the agents, then this is hardly different from the assumption that all agents have already rational beliefs from the beginning. It is just the interesting feature of learning that different agents may be able to coordinate their different beliefs and that these diverse beliefs can converge towards rational expectations.

The aim of a micro-founded model with individual agents is to derive the behaviour of aggregate variables from individual behaviour. In order to ensure the micro-foundation of the model's results, the individual agent's behaviour, including his entire optimization problem and his individual reactions to shocks, must be described with respect to this agent's information set and within the framework set by his microeconomic behavioural functions, regardless of whether his parameters coincide with those of a fictitious representative agent or not. Loosely speaking one could say that even if an individual agent *is* representative, he cannot know this fact and cannot use it to simplify his optimization problem.[146] Consequently, individual behaviour must be described without resorting to the representative agent assumption and the simplifications this assumption makes possible. Thus, the application of the representative agent assumption is only suitable if it is just a technical simplification and does not change the underlying behaviour nor the results of such a model. This is usually true for the variables on the aggregated level. But as pointed out above, the learnability of an equilibrium is by definition a property of individual agents, so it can be misleading if the learnability of an equilibrium is determined by means of a representative agent.

This was demonstrated in the sections above. The results are shown in comparison in table 3.1 on page 81. It turned out that the simple PLM which was perfectly learnable and strongly E-stable for the representative agent could not be learned any more by individual agents with heterogeneous income shocks, although in the REE their average consumption was still the same as the consumption of a representative agent. The reason why a representative agent is able to learn where individual agents are not, is that the market clearing condition can be applied to his behaviour, which cannot be done for

[146] As an example of how these simplifications cannot be used by individual agents, note that a representative agent would not need to solve a dynamic optimization problem at all, because it is immediately clear from the market clearing condition that in each period he consumes exactly his income.

the individual agents.[147] The market clearing condition forces his consumption to follow his income, so it allows the representative agent to learn the rational expectation for his future consumption as being the same as the rational expectation of his future income. For an individual agent this is different. In this case, the market clearing condition imposes merely restrictions on the average consumption, but not on the individual ones. Since thus a direct feedback is missing in the individual agent's learning process, a permanent over- or underconsumption or the influence of a sunspot variable becomes possible, and the expectations become non-rational. Note that the consumption expectations in these cases can be reasonably labelled as non-rational, even if they are confirmed by the actual law of motion for the consumption, because they are not consistent with the "right" solution of the complete dynamic optimization problem. This confirmation takes place only because in such a non-rational equilibrium each actual realization of the consumption is itself again influenced by the same non-rational expectation. In other words, for an individual agent the Euler equation allows for a self-fulfilling prophecy.

Some of these problems have been discussed by Preston.[148] In his articles he criticized the fact that Euler equation learning can lead to non-optimal consumption, because the individual agent's wealth is not taken into account, while the additional problems which arise from heterogeneous income shocks and from overparameterized PLMs have not been examined there. As an alternative method Preston has proposed infinite horizon learning. This will be discussed in the following section.

[147] See also the discussion in section 3.1.4 for the differences between the REE solution between a representative agent and individual agents.

[148] See Preston (2005, 2005a).

	Represen-tative Agent	Individual Agents with Common Income Shocks	Individual Agents with Heterogeneous Income Shocks
REE Solution	$c_t = y_t$	$c_t^i = (1-\beta) b_t^i + \frac{\overline{Y}^i - \overline{C}^i}{\overline{C}^i}\left[\lambda + (1-\lambda)\sigma^{-1}\right] y_t + y_t$	$c_t^i = (1-\beta) b_t^i + \frac{\overline{Y}^i - \overline{C}^i}{\overline{C}^i}\left[\lambda y_t^i + (1-\lambda)\sigma^{-1} \widetilde{y}_t\right] + \lambda\left(y_t^i - \widetilde{y}_t\right) + \widetilde{y}_t$
EE Learning with correctly specified PLM	$c_t = y_t$ weakly E-stable stable under learning	$c_t^i = y_t$ weakly E-stable stable under learning	$c_t^i = \widetilde{y}_t$ weakly E-stable stable under learning
	= REE	= REE only if no wealth	≠ REE
EE Learning with Constant and Sunspot Variable (over-parameterized PLM)	$c_t = y_t$ strongly E-stable stable under learning	$c_t^i = m^i + y_t + e^i \varepsilon_t$ with $\widetilde{m} = 0$ $\begin{cases} e^i = \widetilde{e} = 0 & \text{if } \mu < 1 \\ \widetilde{e} = 0 & \text{if } \mu = 1 \end{cases}$ not E-stable probably not stable under learning	$c_t^i = m^i + \widetilde{y}_t + e^i \varepsilon_t$ with $\widetilde{m} = 0$ $\begin{cases} e^i = \widetilde{e} = 0 & \text{if } \mu < 1 \\ \widetilde{e} = 0 & \text{if } \mu = 1 \end{cases}$ not E-stable probably not stable under learning
	= REE	≠ REE	≠ REE

Table 3.1: Results of Euler equation learning compared with the REE solution

3.3. Infinite Horizon Learning

The infinite horizon learning method (IH learning) was introduced by Bruce Preston to meet some of the shortcomings of EE learning as described in the previous section. In contrast to the latter it is based very closely on the underlying dynamical optimization. With the Euler equation learning method an agent uses only his Euler equation to determine his optimal current consumption. To this end he has to learn how to predict his own consumption in the next period. In contrast to that, IH learning uses another approach to determine the optimal consumption in the current period. It is based primarily on the intertemporal budget constraint instead of the Euler equation, but uses the Euler equation as well to ensure an optimal consumption path.[149] Here, the procedure will be shown immediately for the case of individual agents with heterogeneous income shocks, since the cases of representative agents or of individual agents with common income shocks can be derived easily from the results.[150]

Log-linearizing the intertemporal budget constraint (41) of an individual agent yields

$$(101)\quad 0 = \overline{C}^i b_t^i + \hat{E}_t^i \sum_{T=t}^{\infty} \overline{R}_{t,T}^{\,-1} \left[\overline{Y}^i y_T^i - \overline{C}^i c_T^i - r_{t,T}\left(\overline{Y}^i - \overline{C}^i\right) \right]$$

with $\overline{R}_{t,T} \equiv \prod_{s=t}^{T-1} \overline{R}_s = \beta^{-(T-t)}$ and $r_{t,T} \equiv \ln\left(R_{t,T} / \overline{R}_{t,T}\right) = \sum_{s=t}^{T-1} r_s$, which gives

$$(102)\quad 0 = \overline{C}^i b_t^i + \hat{E}_t^i \sum_{T=t}^{\infty} \beta^{T-t} \left[\overline{Y}^i y_T^i - \overline{C}^i c_T^i - \sum_{s=t}^{T-1} r_s \left(\overline{Y}^i - \overline{C}^i\right) \right].$$

In order to determine the optimal current consumption the agent predicts his income and the interest rate for all future periods. He then uses the expected path of the interest rate to determine the shape of his future consumption path by using the Euler equation for

---

[149] The procedure used here follows the IH learning method which is described for example in Preston (2005) or Preston (2005a). In contrast to these articles, I here apply it to agents with individual income shocks and allow for an overparameterized PLM with a sunspot variable. Moreover, the examination of E-stability is done here with a special focus on the learning process on the individual level to point out the interaction between the learning agents.

[150] A detailed derivation with intermediate steps can be found in appendix A.5.

all periods from now up to infinity. Given his beliefs about future interest rates, the resulting shape of the consumption path is thus optimal, and it is also beyond the influence of the agent, whereas the level of the consumption path can still be chosen by the agent. The optimal level is the one which is as high as possible without violating the intertemporal budget constraint, given the agent's beliefs about future interest rates and his future income. The agent's task is to find the single value for his current consumption which guarantees that the future consumption path is optimal and feasible.

Accordingly, consumption in every future period is here expressed as a function of current consumption and the future path of the interest rate. Formally, the Euler equation (52) is iterated from t to T which yields

$$(103)\quad \hat{E}_t^i c_T^i = c_t^i + \sigma \hat{E}_t^i \sum_{s=t}^{T-1} r_s \ .$$

Inserting into (102) gives

$$(104)\quad 0 = \overline{C}^i b_t^i + \hat{E}_t^i \sum_{T=t}^{\infty} \beta^{T-t} \left[\overline{Y}^i y_T^i - \overline{C}^i c_t^i + \sum_{s=t}^{T-1} r_s \left(-\overline{C}^i \sigma - \overline{Y}^i + \overline{C}^i \right) \right].$$

Solving for the current consumption and simplifying yields the agent's individual decision rule as

$$(105)\quad c_t^i = (1-\beta) b_t^i + \sum_{T=t}^{\infty} \beta^{T-t} \left[(1-\beta) \frac{\overline{Y}^i}{\overline{C}^i} \hat{E}_t^i y_T^i - \left(\sigma + \frac{\overline{Y}^i - \overline{C}^i}{\overline{C}^i} \right) \beta \hat{E}_t^i r_T \right].$$

By this rule the agent's optimal current consumption is determined as a function of his expectations for his future income and the future interest rate.

When he is learning, the agent has non-rational future expectations, which are based on his PLM. His income is an exogenous variable, independent from his behaviour or the behaviour of other agents, and following an AR(1) process. Therefore its law of motion can be learned by the agent by a simple regression over its past evolution, and the learning process will converge to rational expectations with probability one. The learning process for the interest rate, however, must be examined in more detail. The agent's PLM for the interest rate is assumed to be

$$(106)\quad r_T = m^i + l^i \tilde{y}_T + e^i \varepsilon_T + \upsilon_T^i \quad \forall T > t$$

where m^i, l^i, e^i are the parameters representing the agents individual beliefs and υ_T^i is an zero mean error term.[151] Similar to the PLM used in EE learning, this PLM nests the REE solution (59), but at the same time contains a constant and a sunspot variable ε. It is therefore overparameterized and allows to test whether the learning system is strongly E-stable, while weak E-stability can also be tested if m^i and e^i are fixed to zero, so that the resulting PLM contains no more variables than the REE solution. Again, the sunspot variable is common knowledge to all agents and is assumed to follow the AR(1) process

$$(107)\quad \varepsilon_t = \mu \varepsilon_{t-1} + \pi_t$$

where $|\mu| \leq 1$ and $\{\pi_t\}$ is a bounded i.i.d. white noise process with zero mean. To employ this PLM it must be assumed that the agent can observe the value of the average income $\tilde{y}_t$. As in the case of his individual income, he can learn the laws of motion for the average income and the sunspot variable by a regression over their past values, because the income processes and the process of the sunspot variable are exogenous. He will thus learn the true laws of motion of these variables with probability one. In the following it will therefore be assumed that these learning processes of exogenous variables have already converged, so that for these variables rational expectations can be used. The expectations of the agent are then

$$(108)\quad E_t^i y_T^i = \rho^{T-t} y_t^i,$$

$$(109)\quad \begin{aligned} \hat{E}_t^i r_T &= m^i + l^i E_t^i \tilde{y}_T + e^i E_t^i \varepsilon_T \\ &= m^i + l^i \rho^{T-t} \tilde{y}_t + e^i \mu^{T-t} \varepsilon_t . \end{aligned}$$

[151] Note that the agent can observe the current value of the interest rate, so he uses the PLM only for future periods $T>t$.

In the individual decision rule (105), the current value of the interest rate is separated, while the future values and the values for the individual income are replaced by the respective expectations. This is done because the current interest rate is observed, rather than expected, and the agents react immediately to it. After solving the infinitive sums, this procedure yields the current consumption of the individual agent

$$(110)\quad \begin{aligned} c_t^i &= (1-\beta)b_t^i + \frac{\overline{Y}^i}{\overline{C}^i}\frac{1-\beta}{1-\beta\rho}y_t^i \\ &-\left(\sigma + \frac{\overline{Y}^i - \overline{C}^i}{\overline{C}^i}\right)\left(\frac{\beta^2}{1-\beta}m^i + \frac{\beta^2\rho}{1-\beta\rho}l^i\tilde{y}_t + \frac{\beta^2\mu}{1-\beta\mu}e^i\varepsilon_t + \beta r_t\right)\end{aligned}$$

which is now aggregated[152] over all agents. The average consumption is then

$$(111)\quad \tilde{c}_t = \frac{1-\beta}{1-\beta\rho}\tilde{y}_t - \sigma\left(\frac{\beta^2}{1-\beta}\tilde{m} + \frac{\beta^2\rho}{1-\beta\rho}\tilde{l}\tilde{y}_t + \frac{\beta^2\mu}{1-\beta\mu}\tilde{e}\varepsilon_t + \beta r_t\right).$$

Applying the market clearing condition gives the equilibrium interest rate, or in other words, the ALM for the interest rate:

$$(112)\quad r_t = \frac{1}{1-\beta\rho}\left(\frac{\rho-1}{\sigma} - \beta\rho\tilde{l}\right)\tilde{y}_t - \frac{\beta}{1-\beta}\tilde{m} - \frac{\beta\mu}{1-\beta\mu}\tilde{e}\varepsilon_t.$$

The T-mapping from the PLM to the ALM is consequently

$$(113)\quad T\begin{pmatrix} m^i \\ l^i \\ e^i \end{pmatrix} = \left(-\frac{\beta}{1-\beta}\tilde{m},\quad \frac{\rho-1}{\sigma(1-\beta\rho)} - \frac{\beta\rho}{1-\beta\rho}\tilde{l},\quad -\frac{\beta\mu}{1-\beta\mu}\tilde{e}\right)'$$

and the T-mapping for the average parameters is the same. The only fixed point of the mapping is $(0,\quad \sigma^{-1}(\rho-1),\quad 0)'$ for the average parameters as well as for the individual ones. Only if every agent uses these values in his PLM and forms his expectations accordingly, then the resulting interest rate confirms their beliefs and that their

152 This is done by constructing the arithmetic mean of the individual log deviations. For the assumptions that are used to simplify the aggregation of the steady state values, see the discussion in section 3.1.2, especially p. 53, as well as appendix A.2.

expectations are rational. Indeed, the solution with these parameter values is the same as the REE for the interest rate in the canonical solution for a representative agent.

The stability of the system under learning can be determined by examining the ODEs for the individual and for the average PLM parameters[153]

$$(114)\quad \frac{d}{d\tau}\begin{pmatrix} m^i \\ \tilde{m} \end{pmatrix} = \begin{pmatrix} -\beta(1-\beta)^{-1}\tilde{m} - m^i \\ -(1-\beta)^{-1}\tilde{m} \end{pmatrix},$$

$$(115)\quad \frac{d}{d\tau}\begin{pmatrix} l^i \\ \tilde{l} \end{pmatrix} = \begin{pmatrix} \dfrac{\rho-1}{\sigma(1-\beta\rho)} - \dfrac{\beta\rho}{1-\beta\rho}\tilde{l} - l^i \\ \dfrac{\rho-1}{\sigma(1-\beta\rho)} - \dfrac{1}{1-\beta\rho}\tilde{l} \end{pmatrix},$$

$$(116)\quad \frac{d}{d\tau}\begin{pmatrix} e^i \\ \tilde{e} \end{pmatrix} = \begin{pmatrix} -\beta\mu(1-\beta\mu)^{-1}\tilde{e} - e^i \\ -(1-\beta\mu)^{-1}\tilde{e} \end{pmatrix}.$$

The derivatives of the right hand side are all negative, so the equilibrium is globally asymptotically stable and the solution is therefore strongly E-stable. Here again, E-stability implies stability under learning,[154] thus under learning the system converges globally towards the REE solution with probability 1, even if the PLM contains a sunspot variable.

The individual current consumption in the REE can be determined by inserting the learned equilibrium parameters values and the REE solution for the interest rate into the equation (110). This yields the solution

$$(117)\quad c_t^i = (1-\beta)b_t^i + \frac{\overline{Y}^i - \overline{C}^i}{\overline{C}^i}\left[\lambda y_t^i + (1-\lambda)\sigma^{-1}\tilde{y}_t\right] + \lambda\left(y_t^i - \tilde{y}_t\right) + \tilde{y}_t$$

$$\text{with } \lambda = \frac{1-\beta}{1-\beta\rho} = \frac{\overline{R}-1}{\left(\overline{R}-1\right)+\left(1-\rho\right)},\ 0<\lambda<1,$$

---

153 In the appendix stability is determined alternatively for *n* agents simultaneously by examining the ODEs for the whole economy rather than for one agent and the average. See appendix A.5.

154 See Evans and Honkapohja (2001), section 8.6.1.

which was already discussed in section 3.1.4.[155] This solution can easily be adapted to the case of common income shocks or to a representative agent, since the learning process concerns the interest rate, an aggregate variable, and does not rely on the individual properties of the agents. It suffices therefore to replace the individual income in the solution by a common income to arrive at the appropriate solution for common income shocks. Similarly, for a representative agent the solution would simply amount to $c_t = y_t$ since he does not possess any wealth and his steady state consumption is the same as his steady state income.

This shows that the IH learning method is more robust than the basic EE learning method with respect to the representative agent assumption. It allows even individual agents with heterogeneous income shocks to learn the REE, because it explicitly employs not only the Euler equation but also the intertemporal budget constraint. It is thus independent from the "shortcuts" in the form of the simplified solution techniques discussed in sections 3.1.3 and 3.1.4, such as the application of the market clearing condition to individual behaviour, which can only be applied if a representative agent is used.

Preston introduced the IH learning method as a remedy to the problem that EE learning does not take the agent's individual wealth into account. In his article he underlines the fact that the Euler equation is only an optimal decision rule if the expectation therein is taken with respect to the correct probability distribution.[156] He concludes that therefore any learning algorithm based on the Euler equation is inherently faulty, because during the learning process the expectations are non-rational. His main point is that with EE learning, an agent has to forecast his own behaviour, which means that he uses two different approaches towards consumption: On the one hand, he determines his current consumption via the Euler equation, while on the other hand he uses his PLM to form his expectation of consumption. During the learning process these two approaches lead to different results. In Preston's view this violates the "internal consistency"

[155] This equation could alternatively be calculated by replacing the expectations in (105) by the rational expectations for the income and the interest rate. Details can be found in the appendix A.5.

[156] See Preston (2005a), p. 11 ff. as well as Preston (2005), section 1.3, p. 94 ff.

requirement. In contrast to that, IH learning is based exclusively on expectations of exogenous variables, so the question of internal consistency does not impose itself.

However, it can be argued that in this particular respect EE learning in general is nevertheless plausible.[157] The Euler equation is a prerequisite of optimal consumption, so it is natural that an agent uses this comparatively simple relationship to determine his current consumption as well as possible, even if he cannot be sure about his expected consumption in the next period. At the beginning of the learning process, the agent believes that his consumption, as determined by the Euler equation, follows some law of motion.[158] In a stochastic environment, an agent will always experience differences between his expectations and the realization which will actually take place in the future. A priori he cannot know whether such differences are due to stochastic disturbances or a symptom of systematic internal inconsistency. This will only become clear to him when he has learned the actual law of motion of his consumption, but at this point the internal inconsistency has disappeared. In other words, the agent will tolerate a transient internal inconsistency, because he is aware that he is learning and does not know the actual law of motion.

In contrast to that, the critique that standard EE learning leads to non-optimal consumption rules is valid as soon as individual agents are considered. Preston mentions the absence of wealth in the learning process, but similar problems arise due to individual income shocks or the possibility of an overparameterized PLM, as shown above. The IH learning approach, which takes the intertemporal budget constraint explicitly into account, is immune against these problems. In their discussion about learning methods, Preston and Honkapohja/Mitra/Evans disagree about the significance of the market clearing constraint. The latter argue that the market clearing condition implies that the budget constraint of the learning agent is not violated.[159] This is true for the representative agent but does not hold any more if individual agents are considered. On the other hand, Preston argues that the market clearing condition may be part of the equilibrium which agents are trying to learn, so that they cannot know about such

---

[157] For a related argument see also Honkapohja, Mitra and Evans (2003), p. 5 f.

[158] Such a belief is perfectly reasonable since in an equilibrium this will indeed be the case.

[159] See Honkapohja, Mitra and Evans (2003), p. 6.

conditions during the learning process.[160] But actually the knowledge of the market clearing condition only does not matter if all agents have a common income, if none of them has any wealth and if they know that, too. Only in this case they can apply the market clearing condition to their individual consumption. But this means that the agents would be similar to a representative agent. The fundamental difference between the learning approaches is not so much the assumed knowledge of the agents about the market clearing condition, but the fact that the representative agent assumption, which implies that the market clearing condition can be used on the agents behaviour, is necessary for the EE learning approach but not for IH learning.

3.4. Extended Euler Equation Learning

In the preceding sections it has been demonstrated that EE learning is a valid approach for an economy of representative agents, but that it has some drawbacks as soon as individual agents with non-representative beliefs are considered. As an alternative one could use the IH learning method. This approach allows even individual agents with heterogeneous income shocks to learn the REE of the economy. However, although this approach is valid also for individual heterogeneous agents, it was criticized because in this method an agent, using expectations up to the infinite future, is "assumed to be extremely far-sighted even though he is boundedly rational".[161] Moreover, the approach by which an agent determines his optimal current consumption is quite complicated in IH learning. While it seems plausible that in the Euler equation method a household tries to balance his current consumption with his consumption in the next period given the current interest rate, it seems less natural that the household calculates infinite sums up to the future when deciding how much to consume in the current period.

It would be interesting to find a learning approach which is comparatively simple, uses only short term expectations but can nevertheless be applied to individual agents with non-representative beliefs, thus combining the desirable features of EE learning and IH learning. In this section I suggest such a learning approach which I will call "Extended Euler Equation Learning" (EEE learning). Like EE learning it is based on the Euler

---

[160] See Preston (2005), p. 99 or Preston (2005a), p. 12.

[161] See Honkapohja, Mitra and Evans (2003), p. 7.

equation, but in contrast to the former it is augmented by additional variables which allow the agents to learn the REE solution even if they experience heterogeneous income shocks or if they accumulate wealth.

Again, as in the case of EE learning, an agent uses the Euler equation for the current period (53) as a basis to determine his optimal consumption:

$$(118)\quad c_t = \hat{E}_t^i c_{t+1} - \sigma r_t .$$

However, in contrast to EE learning the PLM of his consumption is now extended and includes also his wealth and the average income:

$$(119)\quad c_t^i = m^i + l^i y_{t-1}^i + k^i \tilde{y}_{t-1} + p^i b_{t-1}^i + e^i \varepsilon_{t-1} + \upsilon_t^i .$$

As before, the PLM contains a sunspot variable[162] and a constant as well, so it can be tested whether this overparameterized PLM is strongly E-stable under learning. The agent's expectation of his consumption in the next period is then

$$(120)\quad \hat{E}_t^i c_{t+1}^i = m^i + l^i y_t^i + k^i \tilde{y}_t + p^i b_t^i + e^i \varepsilon_t$$

and the resulting individual decision rule for the current consumption is

$$(121)\quad c_t = m^i + l^i y_t^i + k^i \tilde{y}_t + p^i b_t^i + e^i \varepsilon_t - \sigma r_t .$$

Note that in contrast to IH learning the agent does not take his budget constraint into account or predicts variables more than one period ahead if he uses EEE learning. He simply forecasts his consumption in the next period and derives his optimal current consumption accordingly, given the current interest rate. He will then learn by observing his consumption and the relevant state variables in previous periods and revising his parameter estimates of his PLM accordingly. The following paragraphs describe the way how the economic environment, that is the agent's flow budget

[162] The properties of this variable are the same as in the preceding sections.

constraint, the market and the other agents, determine the outcome of this learning process. The agent himself does not need to be aware of any of these relationships.

Given the individual decision rules, the resulting equilibrium interest rate is determined by aggregating consumption over all agents and applying the market clearing condition, as it was done in section 3.2.3, so we have

$$(122)\quad r_t = \sigma^{-1}\left[\tilde{m} + n^{-1}\sum_{i=1}^{n} l^i y_t^i + (\tilde{k}-1)\tilde{y}_t + n^{-1}\sum_{i=1}^{n} p^i b_t^i + \tilde{e}\,\varepsilon_t\right].$$

Inserting the interest rate and the agent's consumption expectation into (121) yields the ALM of the agent's consumption

$$(123)\quad c_t^i = (m^i - \tilde{m}) + l^i y_t^i - n^{-1}\sum_{j=1}^{n} l^j y_t^j + (k^i - \tilde{k} + 1)\tilde{y}_t + p^i b_t^i - n^{-1}\sum_{j=1}^{n} p^j b_t^j + (e^i - \tilde{e})\varepsilon_t .$$

In order to determine the *T*-mapping from the PLM to the ALM, the variables within this equation which are part of the agent's PLM have to be related to the period $t-1$, since values from this period are used in the PLM. For the exogenous variables, namely the agent's individual income, the average income and the sunspot variable, the necessary replacements can be derived from their respective stochastic processes as

$$(124)\quad \begin{aligned} \varepsilon_t &= \mu\,\varepsilon_{t-1} + \pi_t ,\\ y_t^i &= \rho\, y_{t-1}^i + v_t^i ,\\ \tilde{y}_t &= \rho\,\tilde{y}_{t-1} + \tilde{v}_t . \end{aligned}$$

For the agent's individual wealth in the current period this is more complicated. At the beginning of period *t* it is predetermined by the flow budget constraint (38), which in log-linearized form is

$$(125)\quad b_t^i = \frac{\overline{B}^i}{\overline{C}^i} r_{t-1} + \beta^{-1}\frac{\overline{Y}^i}{\overline{C}^i} y_{t-1}^i + \beta^{-1} b_{t-1}^i - \beta^{-1} c_{t-1}^i .$$

This expression contains two variables, the interest rate and the consumption, which again depend on the agents' PLMs in period $t-1$. Because the behaviour of the system under learning shall be examined, this dependence must be made explicit, thus these variables have to be replaced by their respective actual laws of motion (122) and (123), which are moved backwards by one period. This yields for the current wealth

$$\begin{aligned} b_t^i &= \kappa^i \tilde{m} - \beta^{-1} m^i + \beta^{-1}\left(\frac{\overline{Y}^i}{\overline{C}^i} - l^i\right) y_{t-1}^i + \beta^{-1}\left(1-p^i\right) b_{t-1}^i \\ &\quad + \left(\kappa^i\left(\tilde{k}-1\right) - \beta^{-1} k^i\right)\tilde{y}_{t-1} + \left(\kappa^i \tilde{e} - \beta^{-1} e^i\right)\varepsilon_{t-1} \\ &\quad + \kappa^i n^{-1}\sum_{j=1}^{n} l^j y_{t-1}^j + \kappa^i n^{-1}\sum_{j=1}^{n} p^j b_{t-1}^j \end{aligned} \tag{126}$$

with $\kappa^i = \frac{\overline{B}^i}{\overline{C}^i}\sigma^{-1} + \beta^{-1}$.

Inserting this expression and (124) into the ALM for the agent's current consumption (123) finally results in the ALM as a function of the exogenous variables of the previous period:[163]

[163] This expression is slightly simplified. In equation (126), the parameters of the PLMs which represent the agents' beliefs are from period *t*-1 while the parameters in (123) are from period *t*. Under learning, the parameters evolve from period to period, so they should actually be noted in a dated form in (127) and should not be collected over different periods. However, when the SRA which represents the system under learning is analyzed and the associated ODE is constructed, the vector of parameters is fixed anyway at the equilibrium values so the different dates are not relevant for the analysis of the stability. See p. 28 as well as Evans and Honkapohja (2001), chapter 6. In order to facilitate the notation the parameters are therefore noted without a date and collected without regard to their period.

$$
\begin{aligned}
c_t^i = {} & \left(1-\beta^{-1}p^i\right)m^i - \left(1-\kappa^i\, p^i\right)\tilde{m} \\
& + \left[\left(\rho-\beta^{-1}p^i\right)l^i + \beta^{-1}p^i\frac{\overline{Y}^i}{\overline{C}^i}\right]y_{t-1}^i + l^i\nu_t^i \\
& + \left[\left(\rho-\beta^{-1}p^i\right)k^i - \left(\rho-\kappa^i\, p^i\right)\left(\tilde{k}-1\right)\right]\tilde{y}_{t-1} + \left(k^i-\tilde{k}+1\right)\tilde{\nu}_t \\
& + \left[\left(\mu-\beta^{-1}p^i\right)e^i - \left(\mu-\kappa^i\, p^i\right)\tilde{e}\right]\varepsilon_{t-1} + \left(e^i-\tilde{e}\right)\pi_t \\
& + \left(\beta^{-1}-\beta^{-1}p^i\right)p^i b_{t-1}^i \\
& - \left(\rho-\kappa^i\, p^i\right)n^{-1}\sum_{j=1}^{n} l^j y_{t-1}^j - n^{-1}\sum_{j=1}^{n} l^j\nu_t^j \\
& + \kappa^i\, p^i n^{-1}\sum_{j=1}^{n} p^j b_{t-1}^j - n^{-1}\sum_{j=1}^{n} p^j b_t^j .
\end{aligned}
\tag{127}
$$

In order to simplify the notation, the individual agent's contribution to the sums in the last two lines of this equation was not separated from the sums and collected with the corresponding terms in the second and fifth line.[164] Provided the economy contains a sufficiently large number of agents, the individual agent's influence on these sums is so small it will not be noticeable to the agent. Hence the value of these sums will appear to him as independent from his individual income or wealth and will not influence his learning of the associated parameters l^i and p^i, while they can nevertheless be important for his learning of the parameter k^i, which describes the influence of the average income (and possibly for the learning of the constant m^i, although it will turn out that this is not the case).

The *T*-mappings from the PLM to the ALM, their fixed points and the E-stability of the equilibrium will now be analyzed for each parameter. Again, E-stability implies stability under learning in such a model.[165] The parameter p^i which represents the perceived influence of the agent's wealth in the previous period on his current consumption will be considered first, because this parameter evolves independently from the others while all other parameters depend on it. The *T*-map for this parameter is

---

[164] The significance of the number of agents within the economy was shown in section 3.2.3, where in a similar equation the individual agent's contribution to the aggregate was explicitly accounted for. See equation (95) on p. 75 and the following. That approach could be used here as well, without any difference for the results in an economy with many agents.

[165] See Evans and Honkapohja (2001), section 8.6, p. 198 ff.

$$(128)\quad T(p^i) = (1 - p^i)\beta^{-1} p^i$$

which has two fixed points, at $p^i = 0$ and $p^i = 1 - \beta$. The stability of these fixed points is determined by the ODE

$$(129)\quad \frac{d}{d\tau} p^i = (\beta^{-1} - 1) p^i - \beta^{-1} {p^i}^2$$

with the derivation $\beta^{-1} - 1 - 2\beta^{-1} p^i$ which is positive for $p^i = 0$, implying instability, and negative for $p^i = 1 - \beta$, implying stability for this fixed point. This means that a negative starting value for this parameter will be getting even more negative under learning, while a positive value will converge $p^i = 1 - \beta$.

The instability result is intuitively plausible: If an agent believed that his consumption in the next period depended negatively on his current wealth, then the more he was indebted (respectively wealthy), the more (less) he would consume, accumulating more and more debt (wealth) with an increasing growth rate, which would result in an explosive consumption path. It is a natural assumption to rule out such behaviour. Even if an agent is not sure about the parameter values of his PLM and is trying to learn them, it is plausible to assume that nevertheless he supposes that more wealth would allow him to consume more. Moreover, it is likely that an agent will notice if his consumption follows an explosive path and will question the parameters of his PLM.[166] For the remaining of this section it will therefore be assumed that for every agent the value for p^i is restricted to be always greater than zero. In this case the parameter will converge towards the solution $p^i = 1 - \beta$ under learning with probability 1.

For the parameter l^i, which describes the influence of the individual income, the *T*-map is

$$(130)\quad T(l^i) = (\rho - \beta^{-1} p^i) l^i + \beta^{-1} p^i \frac{\overline{Y}^i}{\overline{C}^i}.$$

[166] A similar argument was suggested in Honkapohja, Mitra and Evans (2003), p. 6 for the case of an explosive solution as a result of a wrong learning approach.

Since the evolution of p^i is not dependent on any other parameter, it will converge to its equilibrium value independently from the way how l^i evolves under learning. Setting it to the equilibrium value yields as the only fixed point of the mapping

$$(131)\quad l^i = \frac{1-\beta}{1-\beta\rho}\frac{\overline{Y}^i}{\overline{C}^i} = \lambda\frac{\overline{Y}^i}{\overline{C}^i}.$$

E-stability can be determined by examining the ODE

$$(132)\quad \frac{d}{d\tau}l^i = \left(-\frac{1-\beta\rho}{\beta}\right)l^i + \frac{1-\beta}{\beta}\frac{\overline{Y}^i}{\overline{C}^i}$$

which is globally asymptotically stable because of the negative coefficient for l^i. So, provided that the initial value for p^i is greater than zero, p^i and l^i will converge to the REE values under learning almost surely, that is with probability one.

In contrast to these parameters which could be learned independently by each agent, the learning process for the remaining parameters is characterized by an interdependence between all agents. To clarify the basis for this remaining learning process, equation (127) is repeated here, with p^i and l^i replaced by their learned values. Note that as soon as the equilibrium value for p^i is learned, this parameter can be factored out from the aggregates of the individual wealth in the last line of (127), because the equilibrium value for p^i is the same for all agents. These aggregates can then be dropped because the average of the individual wealth is always zero in a closed economy. Similarly, when all agents have learned the equilibrium values for l^i, the aggregates in the last line but one of (127) can approximately be written in terms of average values.[167] Put less technically, these sums represent the aggregated influences of all individual agents' wealth and income on the interest rate and consequently on every agent's consumption. As soon as all agents have learned the equilibrium values for these parameters, the influences of their individual wealth cancel out each other, while the combined influences of their individual incomes appear as the influence of the average income to an observing agent. This yields

[167] Using the assumptions discussed in section 3.1.2.

$$
\begin{aligned}
(133)\qquad c_t^i &= \left(1-(1-\beta)\beta^{-1}\right)m^i - \left(1-(1-\beta)\kappa^i\right)\tilde{m} \\
&+ \frac{\overline{Y}^i}{\overline{C}^i}\lambda y_{t-1}^i + \frac{\overline{Y}^i}{\overline{C}^i}\lambda v_t^i \\
&+ \left[\left(\rho-(1-\beta)\beta^{-1}\right)k^i + \left(\rho-(1-\beta)\kappa^i\right)\left(1-\lambda-\tilde{k}\right)\right]\tilde{y}_{t-1} \\
&+ \left(k^i+1-\lambda-\tilde{k}\right)\tilde{v}_t \\
&+ \left[\left(\mu-(1-\beta)\beta^{-1}\right)e^i - \left(\mu-(1-\beta)\kappa^i\right)\tilde{e}\right]\varepsilon_{t-1} + \left(e^i-\tilde{e}\right)\pi_t \\
&+ (1-\beta)b_{t-1}^i .
\end{aligned}
$$

On the basis of this revised ALM the learning process for the parameters m^i, k^i, and e^i will be analyzed one by one. Using the definition for κ^i, the learning process for the constant m^i and its average over all agents is described by the T-maps[168]

$$
(134)\quad T\left(m^i\right) = \left(1-\frac{1-\beta}{\beta}\right)m^i - \left(1-\frac{1-\beta}{\beta}-(1-\beta)\sigma^{-1}\frac{\overline{B}^i}{\overline{C}^i}\right)\tilde{m}
$$

and

$$
(135)\quad T\left(\tilde{m}\right) = (1-\beta)\sigma^{-1}\left(n^{-1}\sum_{j=1}^{n}\frac{\overline{B}^i}{\overline{C}^i}\right)\tilde{m} .
$$

Using the assumptions discussed in section 3.1.2, the aggregate of the steady state wealth terms in (135) can be approximated by zero.[169] Then the fixed points for these parameters are $\tilde{m}=0$ for the average and consequently $m^i=0$. E-stability of this equilibrium can be determined by examining the ODEs

$$
(136)\quad \frac{d}{d\tau}m^i = -\frac{1-\beta}{\beta}m^i - \left(1-\frac{1-\beta}{\beta}-(1-\beta)\sigma^{-1}\frac{\overline{B}^i}{\overline{C}^i}\right)\tilde{m} ,
$$

$$
(137)\quad \frac{d}{d\tau}\tilde{m} = -\tilde{m}
$$

168 For details and an alternative calculation which directly uses the equations for all agents in the economy simultaneously and does not employ the average parameter, see appendix A.6.2.

169 See section 3.1.2, p. 53 ff.

where in the second equation the average of the steady state terms has already been approximated by zero as described above. The derivatives of both ODEs are negative, hence the equilibrium is globally asymptotically stable. Thus, under learning each agent's parameter values will converge to zero with probability 1.

Next, the *T*-maps for the parameter e^i and its average over all agents are

$$(138)\quad T\left(e^i\right)=\left(\mu-\frac{1-\beta}{\beta}\right)e^i-\left(\mu-\frac{1-\beta}{\beta}-(1-\beta)\sigma^{-1}\frac{\overline{B}^i}{\overline{C}^i}\right)\tilde{e},$$

$$(139)\quad T\left(\tilde{e}\right)=(1-\beta)\sigma^{-1}\left(n^{-1}\sum_{j=1}^{n}\frac{\overline{B}^i}{\overline{C}^i}\right)\tilde{e}, \text{ where } n^{-1}\sum_{j=1}^{n}\frac{\overline{B}^i}{\overline{C}^i}\approx 0$$

has been approximated by zero as before. The fixed point of the average parameter is $\tilde{e}=0$. Consequently the fixed point of the individual parameters is also $e^i=0$, except for the case that $\mu=\beta^{-1}$. In this case there would be a continuum of values, indexed by e^i, which would have to satisfy merely the requirement that their average would have to be zero in order to be a fixed point. However, since the sunspot variable ε was defined as following a non-explosive AR(1)-process with $|\mu|\leq 1$ this case is ruled out anyway. E-stability is determined by the ODEs

$$(140)\quad \frac{d}{d\tau}e^i=\left(\mu-\frac{1}{\beta}\right)e^i-\left(\mu-\frac{1-\beta}{\beta}-(1-\beta)\sigma^{-1}\frac{\overline{B}^i}{\overline{C}^i}\right)\tilde{e},$$

$$(141)\quad \frac{d}{d\tau}\tilde{e}=-\tilde{e}.$$

For $|\mu|\leq 1$ the derivatives of both ODEs are negative and the equilibrium is globally asymptotically stable. Under learning each agent's parameter values will converge to zero with probability 1, as before in the case of the constant m^i. In other words, the agents are able to learn that neither a constant nor a sunspot variable are part of their initially overparameterized PLM.

Finally the evolution of the parameter k^i which represents the influence of the average income is examined. This parameter and its average are described by the *T*-maps

$$(142)\quad T\left(k^i\right)=\left(\rho-\frac{1-\beta}{\beta}\right)k^i+\left(\rho-\frac{1-\beta}{\beta}-(1-\beta)\sigma^{-1}\frac{\overline{B}^i}{\overline{C}^i}\right)\left(1-\lambda-\tilde{k}\right),$$

$$(143)\quad T\left(\tilde{k}\right)=(1-\beta)\sigma^{-1}\left(n^{-1}\sum_{j=1}^{n}\frac{\overline{B}^i}{\overline{C}^i}\right)\tilde{k}+\left(\rho-\frac{1-\beta}{\beta}-(1-\beta)\sigma^{-1}n^{-1}\sum_{j=1}^{n}\frac{\overline{B}^i}{\overline{C}^i}\right)(1-\lambda),$$

where, with the approximation $n^{-1}\sum_{i=1}^{n}\left(\overline{B}^i/\overline{C}^i\right)\approx 0$ as before, the second equation simplifies to

$$(144)\quad T\left(\tilde{k}\right)=\rho-\lambda\,.$$

The fixed points for these mappings are

$$(145)\quad \tilde{k}=\rho-\lambda\,,$$

$$(146)\quad k^i=\rho-\lambda-(1-\lambda)(1-\beta)\sigma^{-1}\frac{\overline{B}^i}{\overline{C}^i}\,.$$

E-stability is determined by the ODEs

$$(147)\quad \frac{d}{d\tau}k^i=\left(\rho-\frac{1}{\beta}\right)k^i+\left(\rho-\frac{1-\beta}{\beta}-(1-\beta)\sigma^{-1}\frac{\overline{B}^i}{\overline{C}^i}\right)\left(1-\lambda-\tilde{k}\right),$$

$$(148)\quad \frac{d}{d\tau}\tilde{k}=\rho-\lambda-\tilde{k}\,.$$

The derivatives of both ODEs are negative, so the system is globally asymptotically stable. Under learning each agent's parameter values will converge to the equilibrium values with probability 1.

Summing up the results of the previous paragraphs, each agent will learn the equilibrium values

$$
\begin{aligned}
& m^i = 0 \\
& l^i = \lambda \frac{\overline{Y}^i}{\overline{C}^i} \\
(149)\quad & k^i = \rho - \lambda - (1-\lambda)(1-\beta)\sigma^{-1}\frac{\overline{B}^i}{\overline{C}^i} \\
& p^i = 1-\beta \\
& e^i = 0
\end{aligned}
$$

with probability 1, provided the parameter p^i is required to be strictly positive as discussed above. To show that with these learned equilibrium values for the parameters of the PLM the agent has indeed rational expectations, the values are inserted into equations (122), (120) and (121) which describe the interest rate, the agent's expectation and his current consumption. With the usual approximation for the aggregation, this yields for the interest rate the REE solution $r_t = \sigma^{-1}(\rho-1)\tilde{y}_t$. For the agent's expectation one gets

$$
(150)\quad \hat{E}_t^i c_{t+1}^i = (1-\beta)b_t^i + \lambda \frac{\overline{Y}^i}{\overline{C}^i} y_t^i + \left(\rho - \lambda - (1-\lambda)(1-\beta)\sigma^{-1}\frac{\overline{B}^i}{\overline{C}^i} \right)\tilde{y}_t .
$$

If the agent uses this expectation in his Euler equation, then he can determine his current consumption as

$$
(151)\quad
\begin{aligned}
c_t &= (1-\beta)b_t^i + \lambda \frac{\overline{Y}^i}{\overline{C}^i} y_t^i + \left(\rho - \lambda - (1-\lambda)(1-\beta)\sigma^{-1}\frac{\overline{B}^i}{\overline{C}^i} \right)\tilde{y}_t - \sigma r_t \\
&= (1-\beta)b_t^i + \lambda \frac{\overline{Y}^i}{\overline{C}^i} y_t^i + \left(1 - \lambda + (1-\lambda)\sigma^{-1}\frac{\overline{Y}^i - \overline{C}^i}{\overline{C}^i} \right)\tilde{y}_t
\end{aligned}
$$

which is the REE solution as it was already derived by IH learning and discussed in section 3.1.4.

Thus, if the EE learning method is extended by including additional variables in the way suggested in this section, then even individual agents with heterogeneous income shocks are able to learn the REE.[170] The solution for his consumption is the same as with the IH learning method, and in the equilibrium it is internally consistent. In the discussion of IH learning I have argued that internal inconsistency during the learning process is not relevant to the agent because he cannot detect it when he is still learning, as long as the expectation and the decision rule are internally consistent in the equilibrium. Thus EEE learning meets the critique which gave rise to the IH learning method. Moreover, EEE learning is much easier for an agent to apply and it is more plausible to argue that agents actually use a two period approach than an infinite horizon approach.

On the face of it this result seems to contradict the requirement that to solve the dynamic optimization problem it is necessary to take into account information about all future periods as well as the budget constraint, but in fact it fits well with this requirement. In the discussed model it is sufficient that the agents consider only variables of the current period in their expectations, because it is assumed that the exogenous variables evolve according to an AR(1)-process. This means that the information about the rationally expected future path of these variables is contained in their present values.[171] The agent does not explicitly considers this relationship, but it is implicitly present in the learning process and affects its outcome.

Further, although in the EEE learning approach an agent uses only the Euler equation to determine his current consumption, his budget constraint is nevertheless part of the learning process. Since his wealth is included in his PLM and since his wealth forcibly evolves according to the flow budget constraint, the influence of this constraint is present in the learning process and will serve as a corrective against permanent over- or underconsumption. This requires that the agent learns the right parameter p^i for the influence of wealth in his PLM. It must be admitted that in order to ensure that learning always converges to the REE in EEE learning and to rule out explosive solutions the parameter p^i is required to be greater than zero. The reason is that this learning method

---

170 It should be mentioned again that, due to the aggregation of individual agents, the results are based on the assumptions made during the log-linearization concerning the distribution of the log-linearized variables and the steady states over the population of agents. For details, see section 3.1.2.

171 Otherwise the agent would not be able to learn the optimal current consumption when using this learning method.

uses the flow budget constraint (albeit not known to the agent, instead he simply experiences its effects on his state variables and learns from them), but not the no-ponzi condition.[172] However, as described above, it suffices to assume that the agent must know beforehand that a positive relation exists between his wealth and his consumption to rule out ponzi type solutions.

It is interesting to note that the average of the learned parameters is not zero for the parameter p^i, although for a representative agent the aggregated wealth, being always zero, does not play any role for his consumption. If a representative agent would directly use EEE learning, he would not be able to determine any influence from this zero wealth on his consumption. In contrast to that, if the representative agent is considered as the average of the individual agents which have learned their parameters individually, then his parameter would be the same as for every individual agent, namely $\tilde{p} = 1 - \beta$. This shows again that the learning process of a representative agent cannot generally be used to emulate the aggregate of the learning processes of individual agents.

[172] In contrast to that, the no-ponzi condition is included in the IH learning approach because it is part of the intertemporal budget constraint.

4. Conclusions

One of the basic building blocs of micro-founded economic models are dynamic optimization models. In such models agents try to maximize some target function over an infinite time horizon subject to a constraint. A common example is the consumption model in which agents maximize their discounted utility subject to their budget constraints. Those models are usually solved by assuming a representative agent and rational expectations. While the micro-founded models stand mainly in the tradition of neoclassical theory and are therefore already relatively abstract and idealized compared to purely macroeconomic models like the Keynes-Hicks IS-LM framework, these assumptions render the models even more academic. By replacing rational expectations by learning algorithms, learning theory tries to explain how rational expectations, which else are rather regarded as a hypothetical equilibrium, could be the outcome of optimizing behaviour. Moreover, even if in reality rational expectation could not be perfectly reached by a learning process, learning theory can give hints about how beliefs and expectations could evolve over time in response to the dynamics of a model. Thus learning can bring an abstract model again a little closer to reality.

In this respect it is especially interesting whether a model is stable under learning and to what equilibrium expectations will converge. In this work this question has been studied by the example of a basic consumption model. To take the model a further step towards a more realistic situation it was in particular examined how the representative agent assumption affects the results of different recursive least squares learning approaches and how these results change when individual agents with individual beliefs and possibly heterogeneous income shocks are considered. The model was analyzed in a log-linearized form and the results are based on the assumptions which were made in the linearization process.

First the REE solutions for a representative agent and for individual agents have been established as a point of reference and their respective derivation techniques have been compared. The discussion has stressed the fact that the representative agent assumption allows for a simplified solution of the underlying dynamic optimization problem which takes only the Euler equation into account. That is because for a representative agent the market clearing condition can be applied to his behaviour, especially to his rational expectations. This is not possible for individual agents, and the solution for them must therefore take the complete dynamic optimization problem into account.

When learning is introduced into a model it is done to explain how a dynamic optimizing process can lead to a (possibly rational) expectations equilibrium. However, this makes sense only if individual agents are considered. Learning as an evolution of beliefs in reaction to experiences is originally an individual concept. If a learning process allows the beliefs of a representative agent to converge to a rational expectations equilibrium this can nevertheless mean that the underlying individuals which are represented by the representative agent still have non-rational or even diverging beliefs. Moreover, if the micro-foundations of a model are such that the behaviour of an agent is influenced by his expectations, this relationship cannot always be aggregated: It cannot be taken for granted that the aggregate of the reaction to the individual expectations is the same as the reaction of a single representative agent to the aggregate expectations, especially not when a recursive process like learning is considered. This means that before a learning method is applied to a representative agent it must be made sure that it yields the same results with individual agents.

This has been examined for the Euler equation learning method first. The method is the most usual way of how learning is included in a model, and it is usually applied to a model with a representative agent. While it is well known that this learning method leads to the rational expectations for a representative agent, this is not the case any more for individual agents, because their individual wealth is not included in the agents' perceived law of motion. In this work it has additionally been shown that for individual agents this learning method allows for a continuum of solutions and also for the inclusion of a sunspot variable into the agents' expectations if the PLM is overparameterized. The learning method leads the economy into this continuum of solutions but fails to converge to a particular point. The underlying reason is that the Euler equation which forms the basis for the learning method allows the agents to form self-fulfilling prophecies. Moreover, this work examines the case of heterogeneous individual incomes. If the Euler equation learning method focuses here on the individual income, then even the resulting interest rate, being an aggregate variable, fails to converge to the rational expectations equilibrium. The reason why the method works well for a representative agent can again be found in the fact that for a representative agent the market clearing condition can be applied to his behaviour. This condition restricts the possible outcome for the representative agent in each period, while it is not applicable to individual agents, whose consumptions are therefore more free to be varied by their beliefs and expectations. The consequence is that it is misleading if Euler equation learning is applied to a representative agent to test the learnability of an equilibrium.

The infinite horizon learning method yields the correct equilibria for both a representative agent and individual agents. This method was suggested by Preston as an alternative method which is based closely on the complete dynamic optimization problem. In this work the learning method was applied to the case of heterogeneous individual income in connection with an overparameterized PLM which allows for a constant as well as a sunspot variable. The result is strongly E-stable and the individual agents can learn the correct beliefs. The difference between this method and Euler equation learning lies in the way how the agents solve their optimization problem. In the EE learning approach they use the Euler equation to determine their consumption, so they have to predict their own behaviour. This can lead to self-fulfilling prophecies in the absence of a restriction as it would be the market clearing constraint for the representative agent. In contrast to that, the IH learning approach lets the agents solve their optimization problems in such a way that the necessary expectations, which they need to determine their optimal consumption, are all expectations of variables beyond their immediate influence. This is the reason why this method is valid for individual agents as well as for a representative one. It does not matter that one of the variables is the interest rate, which as an endogenous variable is influenced indirectly by the agents: Since this variable is per se an aggregate one, it makes again no difference if it is determined by an aggregate of individual agents or by a representative agent.

However, the infinite horizon learning approach has been criticized for two main reasons. Since learning implies boundedly rational agents, it seems not natural to assume that they consider expectations over an infinite horizon to determine their current consumption. Moreover, while the Euler equation is a simple relationship which could be interpreted as a “natural” desire to balance consumption over time, the calculations on which the consumption is based in the infinite horizon approach are more complicated. In this work an extension to the original Euler equation approach has been proposed as a more simple alternative. In this extended Euler equation approach the agents’ PLM also includes the average income and the individual wealth in addition to the individual income. It has been shown that for the model discussed here this learning approach converges to the rational expectations equilibrium for individual agents with heterogeneous income provided that the agents are aware of the fact that wealth and consumption are positively related and that they restrict the respective parameter in their PLMs appropriately. As a further evidence for the result that the representative agent assumption can be misleading when learning is considered, it has been shown that the average of the perceived influence of wealth on an agents consumption is a well defined positive value. A representative agent as the average of the underlying individual agents would have to exhibit this value, while a representative agent which would directly learn on the basis of his Euler equation would not be able to detect any influence of wealth on his consumption. It is interesting that in this learning approach it is not necessary that the agents consider their budget constraints explicitly. As long as they include their wealth in their PLMs appropriately, they get the necessary feedback from the budget constraint via this value.

The comparison of these results demonstrate two main points: In order to explain how expectations are formed in a model it is not sufficient to replace simply rational expectations by a learning algorithm. When this is done it must also be ensured that the learning algorithm works well with individual agents and that the results are not changed by the representative agent assumption, as learning algorithms can be sensitive to this assumption. In addition to that it can be seen that Euler equation learning can still be a valid learning algorithm provided that it is applied to the appropriate perceived law of motion.

In future research, the examination as it was done in this work could be extended in some aspects. First, the considered learning methods are all variants of recursive least squares, because this is the most generic and up to now also the most common learning algorithm, as well as yielding the clearest results. It would be interesting to conduct a similar study for other learning algorithms like stochastic gradient learning or constant gain learning. Further, it could be tried to increase the degree of heterogeneity of the agents by including either individual temporal discount factors, individual coefficients for the autocorrelated income processes or different intertemporal elasticities of substitution for the utility function. Moreover, while there are already some experimental and empirical studies about individual learning behaviour, it would be desirable to establish further whether and how individual persons do learn in the real world.

Finally, the results of this work have been demonstrated by the example of a basic dynamic optimization model. Because this model considers only the consumer's problem, positing his income as an exogenous process, it is only a partial model of a microeconomic flavour. As a next step the results could be examined for more complicated models which build on such dynamic optimizations, such as fully-fledged micro-founded macroeconomic models. Because of the additional interdependencies in such models, it can of course not be taken for granted that the results of this work hold also for those more complex models. However, in future research it could be tried to apply the approach used in this work to such models, introducing learning as well as heterogeneous agents. In this respect the present work hopes to contribute to the task of rendering micro-founded models closer and closer to the economic reality.

References

Aoki, Masanao, and Canzoneri, Matthew (1979): “Reduced Forms of Rational Expectations Models”, *Quarterly Journal of Economics*, vol. 93, February 1979, pp. 59-71.

Blanchard, Oliver J. (2000): “What do we know about Macroeconomics that Fischer and Wicksell did not?”, *Quarterly Journal of Economics*, vol. 115, November 1979, pp. 1375-1409.

Blanchard, Oliver J., and Kahn, Charles M. (1980): “The Solution of Linear Difference Models under Rational Expectations”, *Econometrica*, vol. 48, no. 5, July 1980, pp. 1305-1311.

Blinder, Alan S. (1987): „Keynes, Lucas, and Scientific Progress”, *The American Economic Review*, vol. 77, no. 2, May 1987, pp. 130-136.

Branch, William A. (2006): “Restricted Perceptions Equilibria and Learning in Macroeconomics”, in: “Post Walrasian Macroeconomics: Beyond the Dynamic Stochastic General Equilibrium Model”, ed. David Colander, Cambridge University Press, 2006.

Branch, William A., and Evans, George W. (2006): “A Simple Recursive Forecasting Model”, *Economics Letters*, 2006, vol. 91, pp. 158-166.

Brannon, Ike (2006): “Remembering the Man Behind Rational Expectations”, *Regulation: the Cato review of business and government*, vol. 29, Spring 2006, Washington: The Cato Insitute, pp. 18-22.

Bray, Margaret (1982): “Learning, Estimation, and the Stability of Rational Expectations”, *Journal of Economic Theory*, vol. 26, 1982, pp. 318-339

Bray, Margaret M., and Savin, Nathan E. (1986): “Rational Expectations Equilibria, Learning, and Model Specification”, *Econometrica*, vol. 54, no. 5, September 1986, pp. 1129-1160.

Burg, Klemens, Haf, Herbert, and Wille, Friedrich (2002): “Höhere Mathematik für Ingenieure. Band 3: Gewöhnliche Differentialgleichungen, Distributionen, Integraltransformationen”, 4th edition, Vieweg+Teubner Verlag, Wiesbaden, 2002.

Burmeister, Edwin; Flood, Robert P., and Garber, Peter M. (1983): “On the Equivalence of Solutions in Rational Expectations Models”, *Journal of Economic Dynamics and Control*, vol. 5, February 1983, pp. 311-321

Bullard, James, and Cho, In-Koo (2005): “Escapist policy rules”, *Journal of Economic Dynamics and Control*, vol. 29, November 2005, pp. 1841-1865.

Cagan, Phillip (1956): "The Monetary Dynamics of Hyperinflation," in: Friedman, Milton, ed., “Studies in the Quantity Theory of Money”, Chicago: University of Chicago, 1956.

Carceles-Poveda, Eva, and Giannitsarou, Chryssi (2007): “Adaptive learning in practice”, *Journal of Economic Dynamics & Control*, vol. 31, 2007, pp. 2659-2697.

Cass, David, and Shell, Karl (1983): “Do Sunspots Matter?”, *Journal of Political Economy*, vol. 91, April 1983, pp. 193-227.

Chiappori, Pierre A., and Guesnerie, Roger (1991): “Sunspot Equilibria in Sequential Markets Models”, in: Hildenbrand, W. and Sonnenschein, H., eds., “Handbook of Mathematical Economics”, North-Holland, Amsterdam, 1991, vol. 4, pp 1683-1762.

Dimand, Robert W. (1999): “Irving Fisher and the Fisher Relation: Setting the Record Straight”, *The Canadian Journal of Economics / Revue canadienne d'Economique*, vol. 32, no. 3, May 1999, pp. 744-750.

DeCanio, Stephen J. (1979): “Rational Expectaions ad Learning from Experience”, *The Quarterly Journal of Economics*, February 1979, pp. 47-57.

De Grauwe, Paul (2006): “Central banking model for neither gods nor monkeys”. Commentary in The Financial Times, July 26, 2006.

De Vroey, Michel (2006): “Getting rid of Keynes? A reflection on the history of macroeconomics”, Discussion paper 2006-51, Department of Economics, Université catholique de Louvain.

Dieckheuer, Gustav (2003): “Makroökonomik. Theorie und Politik”, 5. Auflage, Springer-Verlag, Berlin, 2003.

Evans, George W. (1986): “Selection Criteria for Models with Non-Uniqueness”, *Journal of Monetary Economics*, vol. 18, 1986, pp. 147-157.

Evans, George W. (1989): “The Fragility of Sunspots and Bubbles”, *Journal of Monetary Economics*, vol. 23, 1989, pp. 297-317.

Evans, George W. (2001): “Expectations in Macroeconomics: Adaptive versus Eductive Learning”, *Revue Economique*, vol. 52, no. 3, 2001, pp. 573-582.

Evans, George W. and Honkapohja, Seppo (2001): “Learning and Expectations in Macroeconomics”, Princeton, Princeton University Press, 2001.

Evans, George W., Honkapohja, Seppo, and Williams, Noah (2005): “Generalized Stochastic Gradient Learning”, *NBER Technical Working Papers*, no. 0317, October 2005, National Bureau of Economic Research.

Evans, George W. and Guesnerie, Roger (1993): “Rationalizability, Strong Rationality, and Expectational Stability”, *Games and Economic Behavior*, vol. 5, 1993, pp. 632-646.

Felderer, Bernhard, and Homburg, Stefan (1993): “Makroökonomik und neue Makroökonomik”, 6th edn., Springer, Berlin, 1993.

Fuldenberg, Drew, and Levine, David K. (1998): “The Theory of Learning in Games”, MIT Press, Cambridge, Massachusetts, 1998.

Hansen, Lars P., and Sargent, Thomas J. (2001): “ Acknowledging Misspecification in Macroeconomic Theory”, *Review of Economic Dynamics*, vol. 4, 2001, pp. 519-535.

Harvey, Andrew C. (1995): “Zeitreihenmodelle”, 2nd edition, Oldenbourg, München.

Heinemann, Maik (2000): “Convergence of Adaptive Learning and Expectational Stability: The Case of Multiple Rational-Expectations Equilibria”. *Macroeconomic Dynamics*, vol. 4, 2000, pp. 263-288.

Honkapohja, Seppo, Mitra, Kaushik, and Evans, George W. (2003): “Notes on Agents’ Behavioral Rules Under Adaptive Learning and Recent Studies of Monetary Policy”, Working Paper, Revision from March 6, 2003, available at http://darkwing.uoregon.edu/~gevans/learn-notes.pdf

Hoover, Kevin D. (2008): “Idealizing Reduction: The Microfoundations of Macroeconomics”, Working Paper, Departments of Economics and Philosophy, Duke University, Durham, NC.

Keuzenkamp, Hugo A. (1991): “A precursor to Muth: Tinbergen's 1932 model of rational expectations”, *The Economic Journal*, vol. 101, September 1991, pp. 1245-1253.

Kirman, Alan P. (1992): „Whom or What Does the Representative Individual Represent?”, *Journal of Economic Perspectives*, vol. 6, no. 2, Spring 1992, pp 117-136.

Ljung, Lennart, and Söderström, Torsten (1987): “Theory and Practice of Recursive Identification”. 1st MIT Press paperback edition, 1987, MIT Press, Cambridge, Massachusetts.

Lucas, Robert (1975): “An Equilibrium Model of the Business Cycle”, *Journal of Political Economy*, vol. 83, no. 6, December 1975, pp. 1113-1144.

Marcet, Albert, and Nicolini, Juan P. (2003): “Recurrent Hyperinflations and Learning”, *The American Economic Review*, vol. 93, no. 5, December 2003, pp. 1476-1498.

Marcet, Albert, and Sargent, Thomas J. (1989): “Convergence of Least Squares Learning Mechanisms in Self-Referential Linear Stochastic Models”, *Journal of Economic Theory*, vol. 48, 1989, pp. 337-368.

Mason, Robin and Wright, Stephen (2001): “The effects of uncertainty on optimal consumption”, *Journal of Economic Dynamics and Control*, vol. 25, 2001, pp. 185-212.

McCallum, Bennett T. (1983): “On non-uniqueness in rational expectations models: An Attempt at Perspective”, *Journal of Monetary Economics*, vol. 11, 1993, pp. 139-168.

McCallum, Bennett T. (1998): “Solutions to linear rational expectations models: a compact exposition”, *Economic Letters*, vol. 61, 1998, pp. 143-147.

McCallum, Bennett T. (1999): “Role of the Minimal State Variable Criterion in Rational Expectations Models”, *International Tax and Public Finance*, vol. 6, 1999, pp. 621-639.

McCallum, Benett T. (2002): “Consistent Expectations, Rational Expectations, Multiple-Solution Indeterminacies, and Least-Squares Learnability”, NBER Working Paper no. 9218, September 2002, National Bureau of Economic Research.

Muth, John F. (1961): “Rational Expectations and the theory of price movements”, *Econometrica*, vol. 29, July 1961, pp. 315-335.

Obstfeld, Maurice and Rogoff, Kenneth (1996): “Foundations of International Macroeconomics”, MIT Press, Massachusetts Institute of Technology, Cambridge, MA.

Plosser, Charles (1989): “Understanding Real Business Cycles”, *Journal of Economic Perspectives*, American Economic Association, vol. 3, Summer 1989, pp. 51-77.

Preston, Bruce (2005): “Learning about Monetary Policy Rules when Long-Horizon Expectations Matter”, *International Journal of Central Banking*, September 2005, pp. 81-126.

Preston, Bruce (2005a): “Adaptive Learning in Infinite Horizon Decision Problems”, mimeo, Revision from September 22, 2005, available at: http://www.columbia.edu/~bp2121/infinite.pdf

Sargent, Thomas J. (1993), “Bounded Rationality in Macroeconomics”, Oxford University Press, Oxford 1993.

Sargent, Thomas J. (1996): “Expectations and the nonneutrality of Lucas”, *Journal of Monetary Economics*, vol. 37, June 1996, pp. 353-548.

Sargent, Thomas, Williams, Noah, and Zha, Tao: "The Conquest of South American Inflation", Working Paper 2006-20, Federal Reserve Bank of Atlanta, November 2006.

Simon, Herbert A. (1957): "Models of Man: Social and Rational; Mathematical Essays on Rational Human Behavior in Society Setting", John Wiley, New York 1957.

Taylor, John B. (1977): "Conditions for Unique Solutions in Stochastic Macroeconomic Models with Rational Expectations", *Econometrica*, vol. 45, no. 6, Semptember 1977, pp. 1377-1385.

Tinbergen, Jan (1932): "Ein Problem der Dynamik", *Zeitschrift für Nationalökonomie*, vol. 3, no. 2, April 1932, pp. 169-184.

Weber, Anke (2007): Heterogenous expectations, learning and European inflation dynamics. Discussion paper no. 16/2007, Deutsche Bundesbank, Discussion paper series 1: Economic studies.

Williams, Noah (2004): "Escape Dynamics in Learning Models", Working Paper (submitted to *Econometrica*), Revision from June 15, 2004, available at http://www.princeton.edu/~noahw/edlm-new8.pdf

Woodford, Michael (1999): "Revolution and Evolution in Twentieth-Century Macroeconomics", mimeo, presented at a conference "Frontiers of the Mind in the Twenty-First Century", U.S. Library of Congress, Washington, D.C., June 1999. Available at http://www.columbia.edu/~mw2230/macro20C.pdf

Woodford, Michael (2003): "Interest and Prices: Foundations of a theory of Monetary Policy", Princeton University Press, Princeton, NJ, 2003.

Wren-Lewis, Simon (2006): "Are There Dangers in the Microfoundation Consensus?", Working Paper, available at http://www.econ.ox.ac.uk/members/simon.wren-lewis/docs/dangers_in_the_microfoundations_consensus.pdf

A. Appendix

A.1. Adaptive expectations as a special case of learning

Without regard to a specific model, let an agent's PLM be:

$$(1) \quad x_{t+1} = c_t z_t + u_t$$

with u being a disturbance with zero means. His expectation then takes the form

$$(2) \quad E_t x_{t+1} = c_t z_t$$

Let the agent's estimate of the parameter c be adjusted according to stochastic gradient learning:

$$(3) \quad c_t = c_{t-1} + \gamma_t z_{t-1}(x_t - c_{t-1} z_{t-1})$$

Setting the exogenous variable z as constant 1 and inserting into (2) yields

$$(4) \quad E_t x_{t+1} = c_{t-1} + \gamma_t (x_t - c_{t-1})$$

and since from (2) $E_{t-1} x_t = c_{t-1}$ we arrive at

$$(5) \quad E_t x_{t+1} = E_{t-1} x_t + \gamma_t (x_t - E_{t-1} x_t)$$

For a constant gain γ this gives the familiar adaptive expectation. Note that while in most models (1) with $z = 1$ would be a misspecification, learning with a constant gain still allows the expectation to track real data.

A.2. Aggregation of log-linearized variables

Consider an individual agent's variable X_t^i which shall be aggregated over all *n* agents in the economy. The aggregate is here denoted by X_t and describes the arithmetic mean[1] of all agents' X_t^i, so that we have

(6) $$X_t \equiv \frac{1}{n}\sum_{i=1}^{n} X_t^i$$

Log-linearize both sides around each variables steady state $\overline{X}$ resp. $\overline{X}^i$ to get

(7) $$\overline{X} + \overline{X}\,x_t = \frac{1}{n}\sum_{i=1}^{n} \overline{X}^i + \frac{1}{n}\sum_{i=1}^{n} \overline{X}^i x_t^i$$

where lower case variables denote the log deviation from their respective steady states. The problem is now to find the mean of the product $\overline{X}^i x_t^i$ on the right hand side, which can be calculated as

(8) $$\frac{1}{n}\sum_{i=1}^{n} \overline{X}^i x_t^i = \left(\frac{1}{n}\sum_{i=1}^{n} \overline{X}^i\right)\left(\frac{1}{n}\sum_{i=1}^{n} x_t^i\right) + \frac{1}{n}\sum_{i=1}^{n}\left[\left(\overline{X}^i - \frac{1}{n}\sum_{i=1}^{n} \overline{X}^i\right)\left(x_t^i - \frac{1}{n}\sum_{i=1}^{n} x_t^i\right)\right]$$

Consider a variable of a single agent as a random variable with a certain distribution. If the number of agents within the economy is sufficiently large, the law of large numbers allows to approximate the average of the variable over all agents by the expected value of its distribution. The second term on the right hand side of (8) can then be approximated by the covariance between $\overline{X}^i$ and x_t^i. Thus, if the distribution of x_t^i over the agents in the economy is not correlated with the distribution of the steady states $\overline{X}^i$ and assuming that the number of agents in the economy is sufficiently large, this term is approximately zero and we can approximate the mean of the product by the product of the respective means:

---

[1] The same method could as well be applied if the aggregate was defined as the sum instead of the arithmetic mean. The result would be the same.

$$(9)\quad \overline{X}+\overline{X}\,x_t=\frac{1}{n}\sum_{i=1}^{n}\overline{X}^i+\frac{1}{n}\sum_{i=1}^{n}\overline{X}^i\,\frac{1}{n}\sum_{i=1}^{n}x_t^i$$

Recognizing that (6) holds for the steady states as well and simplifying, we get

$$(10)\quad x_t=\frac{1}{n}\sum_{i=1}^{n}x_t^i$$

Thus, if in an economy with many agents the steady state of the agent's variable in question is uncorrelated with the variable's log deviation, the log-linearized variable can be aggregated by simply taking the mean of the log deviations.

A.3. The REE solution for heterogeneous agents

A.3.1. Optimal Consumption and the Euler Equation

Consider the consumption model described in the text and a world of rational expectations. Suppose that the agent's consumption is given by

$$(11)\quad c_t^i=(1-\beta)\,b_t^i+\frac{\overline{Y}^i-\overline{C}^i}{\overline{C}^i}\left[\lambda\,y_t^i+(1-\lambda)\sigma^{-1}\tilde{y}_t\right]+\lambda\left(y_t^i-\tilde{y}_t\right)+\tilde{y}_t$$

To show that the agent's consumption and his expectation of his own consumption satisfy the Euler equation under rational expectations, form the rational expectation of (11) for the next period which is

$$(12)\quad E_t^i c_{t+1}^i=(1-\beta)\,E_t^i b_{t+1}^i+\frac{\overline{Y}^i}{\overline{C}^i}\lambda\,E_t^i y_{t+1}^i+\left[\frac{\overline{Y}^i-\overline{C}^i}{\overline{C}^i}\sigma^{-1}+1\right](1-\lambda)\,E_t^i\tilde{y}_{t+1}$$

where b_t^i is predetermined by the variables from the current period according to the log-linearized flow budget constraint

$$(13)\quad b_{t+1}^i=\frac{\overline{B}^i}{\overline{C}^i}r_t+\beta^{-1}\frac{\overline{Y}^i}{\overline{C}^i}y_t^i+\beta^{-1}b_t^i-\beta^{-1}c_t^i$$

Since we are considering a world of rational expectations, the current interest rate in this expression can be expressed by the REE solution $r_t = \sigma^{-1}\left(1-\rho\right)\tilde{y}_t$. Inserting (13) and the rational expectations for the individual and the average income into (12) yields

$$
\begin{aligned}
(14)\quad E_t^i c_{t+1}^i =& \left(1-\beta\right)\left(\frac{\overline{B}^i}{\overline{C}^i}\sigma^{-1}\left(1-\rho\right)\tilde{y}_t + \beta^{-1}\frac{\overline{Y}^i}{\overline{C}^i}y_t^i + \beta^{-1}b_t^i - \beta^{-1}c_t^i\right) \\
& + \frac{\overline{Y}^i}{\overline{C}^i}\lambda\,\rho y_t^i + \left[\frac{\overline{Y}^i - \overline{C}^i}{\overline{C}^i}\sigma^{-1} + 1\right]\left(1-\lambda\right)\rho\tilde{y}_t
\end{aligned}
$$

According to the Euler equation this expectation can be used to determine the current consumption as

$$(15)\quad c_t^i = E_t^i c_{t+1}^i + (1-\rho)\tilde{y}_t$$

where the current interest rate as again been replaced by the REE solution. Solving for the current consumption and recognizing that $\left(1-\beta\right)\overline{B}^i = -\left(\overline{Y}^i - \overline{C}^i\right)$ gives

$$(16)\quad c_t^i = \frac{\overline{Y}^i - \overline{C}^i}{\overline{C}^i}\sigma^{-1}\left(\beta-\lambda\beta\rho\right)\tilde{y}_t + \frac{\overline{Y}^i}{\overline{C}^i}\left(1-\beta+\lambda\,\beta\rho\right)y_t^i + \left(1-\beta\right)b_t^i + \left(\beta-\lambda\beta\rho\right)\tilde{y}_t$$

Since $\beta\rho\lambda + 1 - \beta = \lambda$ and $\beta - \beta\rho\lambda = 1-\lambda$ we have indeed

$$(17)\quad c_t^i = \frac{\overline{Y}^i - \overline{C}^i}{\overline{C}^i}\sigma^{-1}\left(1-\lambda\right)\tilde{y}_t + \frac{\overline{Y}^i}{\overline{C}^i}\lambda y_t^i + \left(1-\beta\right)b_t^i + \left(1-\lambda\right)\tilde{y}_t$$

which is the same as (11).

■

A.3.2. Permanent shift of the individual agent’s consumption steady state

The optimal consumption of an individual agent in a world of rational expectations is given by

$$(18)\quad c_t^i = (1-\beta)\, b_t^i + \frac{\overline{Y}^i - \overline{C}^i}{\overline{C}^i}\left[\lambda\, y_t^i + (1-\lambda)\sigma^{-1}\tilde{y}_t\right] + \lambda\left(y_t^i - \tilde{y}_t\right) + \tilde{y}_t$$

$$\text{with } \lambda = \frac{\overline{R}-1}{\overline{R}-\rho} = \frac{1-\beta}{1-\beta\rho},\ 0<\lambda<1$$

where all terms on the right hand side except for the last one represent the expected permanent shift of the consumption steady state, caused by changes in the agent's wealth which will accumulate while the income shock settles down:

$$(19)\quad \begin{aligned} \lim_{s\to\infty}\left[E_t^i c_{s+1}^i\right] &= (1-\beta)\, b_t^i + \frac{\overline{Y}^i - \overline{C}^i}{\overline{C}^i}\left[\lambda\, y_t^i + (1-\lambda)\sigma^{-1}\tilde{y}_t\right] + \lambda\left(y_t^i - \tilde{y}_t\right) \\ &= (1-\beta)\, b_t^i + \frac{\overline{Y}^i}{\overline{C}^i}\lambda\, y_t^i + \frac{\overline{Y}^i - \overline{C}^i}{\overline{C}^i}(1-\lambda)\sigma^{-1}\tilde{y}_t - \lambda\,\tilde{y}_t \end{aligned}$$

To show that this expression remains constant over time provided no additional shocks occur, determine next period's wealth by log linearizing the flow budget constraint:

$$(20)\quad b_{t+1}^i = \frac{\overline{B}^i}{\overline{C}^i} r_t + \beta^{-1}\frac{\overline{Y}^i}{\overline{C}^i} y_t^i + \beta^{-1} b_t^i - \beta^{-1} c_t^i$$

where consumption can be expressed as

$$(21)\quad c_t^i = \lim_{s\to\infty}\left[E_t^i c_{s+1}^i\right] + \tilde{y}_t$$

Advancing (19) by one period and inserting (20) and (21) gives

$$(22)\quad \begin{aligned} \lim_{s\to\infty}\left[E_{t+1}^i c_{s+1}^i\right] = (1-\beta)\frac{\overline{B}^i}{\overline{C}^i} r_t &+ \frac{1-\beta}{\beta}\frac{\overline{Y}^i}{\overline{C}^i} y_t^i + \frac{1-\beta}{\beta} b_t^i \\ &- \frac{1-\beta}{\beta}\lim_{s\to\infty}\left[E_t^i c_{s+1}^i\right] - \frac{1-\beta}{\beta}\tilde{y}_t \\ &+ \frac{\overline{Y}^i}{\overline{C}^i}\lambda\, y_{t+1}^i + \frac{\overline{Y}^i - \overline{C}^i}{\overline{C}^i}(1-\lambda)\sigma^{-1}\tilde{y}_{t+1} - \lambda\,\tilde{y}_{t+1} \end{aligned}$$

Because additional future shocks are ruled out here, we have $y_{t+1}^i = \rho y_t^i$ and $\tilde{y}_{t+1} = \rho\tilde{y}_t$, and since we are considering a world of rational expectations the interest rate in *t* is

$r_t = -\sigma^{-1}(1-\rho)\tilde{y}_t$. Inserting these terms into (22), recognizing that $(1-\beta)\bar{B}^i = -(\bar{Y}^i - \bar{C}^i)$ and sorting then yields:

$$\lim_{s\to\infty}\left[E_{t+1}^i c_{s+1}^i\right] = \left(1-\frac{1}{\beta}\right)\lim_{s\to\infty}\left[E_t^i c_{s+1}^i\right] + \frac{1}{\beta}\left[(1-\beta)b_t^i + \frac{\bar{Y}^i}{\bar{C}^i}(\beta\rho\lambda + 1 - \beta)y_t^i\right] + \frac{1}{\beta}\left[\frac{\bar{Y}^i - \bar{C}^i}{\bar{C}^i}(\beta - \beta\rho\lambda)\sigma^{-1}\tilde{y}_t - (\beta\rho\lambda + 1 - \beta)\tilde{y}_t\right] \tag{23}$$

Since $\beta\rho\lambda + 1 - \beta = \lambda$ and $\beta - \beta\rho\lambda = 1 - \lambda$ we have

$$\lim_{s\to\infty}\left[E_{t+1}^i c_{s+1}^i\right] = \left(1-\frac{1}{\beta}\right)\lim_{s\to\infty}\left[E_t^i c_{s+1}^i\right] + \frac{1}{\beta}\left[(1-\beta)b_t^i + \frac{\bar{Y}^i}{\bar{C}^i}\lambda\, y_t^i + \frac{\bar{Y}^i - \bar{C}^i}{\bar{C}^i}(1-\lambda)\sigma^{-1}\tilde{y}_t + \lambda\,\tilde{y}_t\right] = \lim_{s\to\infty}\left[E_t^i c_{s+1}^i\right] \tag{24}$$

■

A.4. E-stability of EE Learning for Individual Agents

Consider the learning model of section 3.2.2 with the mapping from the PLM to the ALM

$$T\begin{pmatrix} m^i \\ l^i \\ e^i \end{pmatrix} = \begin{pmatrix} m^i - \tilde{m} \\ (l^i - \tilde{l} + 1)\rho \\ (e^i - \tilde{e})\mu \end{pmatrix} \tag{25}$$

In a world of n individual agents E-stability is determined by the stability of the ODEs

$$(26)\quad \frac{d}{d\tau}\begin{pmatrix} m^1 \\ m^2 \\ \vdots \\ m^n \end{pmatrix} = -\frac{1}{n}\underset{n\times n}{\boldsymbol{J}}\begin{pmatrix} m^1 \\ m^2 \\ \vdots \\ m^n \end{pmatrix}$$

$$(27)\quad \frac{d}{d\tau}\begin{pmatrix} l^1 \\ l^2 \\ \vdots \\ l^n \end{pmatrix} = \rho\underset{n\times 1}{\boldsymbol{J}} + \left[(\rho-1)\underset{n\times n}{\boldsymbol{I}} - \frac{1}{n}\rho\underset{n\times n}{\boldsymbol{J}}\right]\begin{pmatrix} l^1 \\ l^2 \\ \vdots \\ l^n \end{pmatrix}$$

$$(28)\quad \frac{d}{d\tau}\begin{pmatrix} e^1 \\ e^2 \\ \vdots \\ e^n \end{pmatrix} = \left[(\mu-1)\underset{n\times n}{\boldsymbol{I}} - \frac{1}{n}\mu\underset{n\times n}{\boldsymbol{J}}\right]\begin{pmatrix} e^1 \\ e^2 \\ \vdots \\ e^n \end{pmatrix}$$

where $\boldsymbol{I} = \begin{pmatrix} 1 & 0 & \cdots & 0 \\ 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 1 \end{pmatrix}$ and $\boldsymbol{J} = \begin{pmatrix} 1 & 1 & \cdots & 1 \\ 1 & 1 & \cdots & 1 \\ \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & \cdots & 1 \end{pmatrix}$.

These ODEs are asymptotically stable if all eigenvalues of their Jacobian matrices have negative real parts. For (27) this condition holds, since the matrix term in square brackets has one distinct eigenvalue of $\lambda = -1$ and an $(n-1)$-fold eigenvalue of $\lambda = \rho - 1 < 0$.[2] We get the same result for (28) if $\mu < 1$. However, for (26) as well as for (28) in the case of $\mu = 1$, the Jacobian has one distinct eigenvalue of $\lambda = -1$ and an $(n-1)$-fold eigenvalue of $\lambda = 0$. The eigenvectors corresponding to this zero eigenvalue are the vectors that satisfy the condition that the sum of their elements is zero, hence the corresponding eigenspace has the dimension $(n-1)$. So, for this eigenvalue its algebraic multiplicity is the same as the geometric multiplicity, hence the ODEs are stable, but not attractive and therefore not asymptotically stable.[3] With respect to learning this means that the ODE allows these parameters to stay at an equilibrium point, but if they are pushed out of it by exogenous stochastic influences,

[2] For some rules on the calculation of these eigenvalues, see appendix A.6

[3] See Burg, Haf, and Wille (2002), Satz 1.1.4, p. 110.

they do not necessarily return to the same equilibrium point. Thus, they do not converge towards a particular equilibrium, but are moved around within the equilibrium continuum by the exogenous stochastic variables.

A.5. Details for the Derivation of Infinite Horizon Learning and of the REE solution for Individual Agents

Consider an economy of individual agents with heterogeneous income shocks which use the infinite horizon learning approach as described in section 3.3. The derivation starts with the intertemporal budget constraint for an individual agent

$$(29)\quad \hat{E}_t^i \sum_{T=t}^{\infty} R_{t,T}{}^{-1} C_T^i = B_t^i + \hat{E}_t^i \sum_{T=t}^{\infty} R_{t,T}{}^{-1} Y_T^i$$

Log linearizing yields

$$(30)\quad 0 = \overline{C}^i b_t^i + \hat{E}_t^i \sum_{T=t}^{\infty} \overline{R}_{t,T}{}^{-1} \left[\overline{Y}^i y_T^i - \overline{C}^i c_T^i - r_{t,T} \left(\overline{Y}^i - \overline{C}^i \right) \right]$$

Use the definitions $\overline{R}_{t,T} \equiv \prod_{s=t}^{T-1} \overline{R}_s = \beta^{-(T-t)}$ and $r_{t,T} \equiv \ln\left(R_{t,T} / \overline{R}_{t,T} \right) = \sum_{s=t}^{T-1} r_s$ to write this constraint as

$$(31)\quad 0 = \overline{C}^i b_t^i + \hat{E}_t^i \sum_{T=t}^{\infty} \beta^{T-t} \left[\overline{Y}^i y_T^i - \overline{C}^i c_T^i - \sum_{s=t}^{T-1} r_s \left(\overline{Y}^i - \overline{C}^i \right) \right]$$

To ensure optimality of the consumption in future periods, iterate the log linearized individual Euler equation to get

$$(32)\quad \hat{E}_t^i c_T^i = c_t^i + \sigma \hat{E}_t^i \sum_{s=t}^{T-1} r_s$$

Inserting into (31) gives

$$(33)\quad \begin{aligned} 0 &= \overline{C}^i b_t^i + \hat{E}_t^i \sum_{T=t}^{\infty} \beta^{T-t} \left[\overline{Y}^i y_T^i - \overline{C}^i \left(c_t^i + \sigma \sum_{s=t}^{T-1} r_s \right) - \sum_{s=t}^{T-1} r_s \left(\overline{Y}^i - \overline{C}^i \right) \right] \\ &= \overline{C}^i b_t^i + \hat{E}_t^i \sum_{T=t}^{\infty} \beta^{T-t} \left[\overline{Y}^i y_T^i - \overline{C}^i c_t^i + \sum_{s=t}^{T-1} r_s \left(-\overline{C}^i \sigma - \overline{Y}^i + \overline{C}^i \right) \right] \end{aligned}$$

Now solve for the current consumption. Rearranging we get

$$(34)\quad \begin{aligned} 0 &= \overline{C}^i b_t^i + \sum_{T=t}^{\infty} \beta^{T-t} \overline{Y}^i \hat{E}_t^i y_T^i - \sum_{T=t}^{\infty} \beta^{T-t} \overline{C}^i c_t^i \\ &\quad + \left(-\overline{C}^i \sigma - \overline{Y}^i + \overline{C}^i \right) \sum_{T=t}^{\infty} \left[\beta^{T-t} \sum_{s=t}^{T-1} \hat{E}_t^i r_s \right] \end{aligned}$$

Simplify this expression using $\sum_{T=t}^{\infty} \left[\beta^{T-t} \sum_{s=t}^{T-1} \hat{E}_t^i r_s \right] = \sum_{T=t}^{\infty} \beta^{T-t} (1-\beta)^{-1} \beta \hat{E}_t^i r_T$ and $\sum_{T=t}^{\infty} \beta^{T-t} = (1-\beta)^{-1}$ to arrive at

$$(35)\quad \begin{aligned} 0 &= \overline{C}^i b_t^i + \sum_{T=t}^{\infty} \beta^{T-t} \left[\overline{Y}^i \hat{E}_t^i y_T^i + (1-\beta)^{-1} \left(-\overline{C}^i \sigma - \overline{Y}^i + \overline{C}^i \right) \beta \hat{E}_t^i r_T \right] \\ &\quad - (1-\beta)^{-1} \overline{C}^i c_t^i \end{aligned}$$

and solve for the current consumption

$$(36)\quad c_t^i = (1-\beta) b_t^i + \sum_{T=t}^{\infty} \beta^{T-t} \left[(1-\beta) \frac{\overline{Y}^i}{\overline{C}^i} \hat{E}_t^i y_T^i - \left(\sigma + \frac{\overline{Y}^i - \overline{C}^i}{\overline{C}^i} \right) \beta \hat{E}_t^i r_T \right]$$

For IH learning, the expectations in (36) must be replaced by the rational expectations for the income (because this law of motion can be learned independently) and the non rational expectation for the interest rate. The latter are, based on the agent's PLM:

$$(37)\quad r_T = m^i + l^i \tilde{y}_T + e^i \varepsilon_T + \upsilon_T^i \qquad \forall T > t$$

and the expectations are accordingly

$$(38)\quad E_t^i y_T^i = \rho^{(T-t)} y_t^i$$

$$(39)\quad \hat{E}_t^i r_T = m^i + l^i E_t^i \tilde{y}_T + e^i E_t^i \varepsilon_T = m^i + l^i \rho^{T-t} \tilde{y}_t + e^i \mu^{T-t} \varepsilon_t$$

In the agent's decision rule (36), the current value of the interest rate is separated because it is not expected, while the future values and the values for the individual income are replaced by the respective expectations:

$$(40)\quad \begin{aligned} c_t^i &= (1-\beta) b_t^i + \sum_{T=t}^{\infty} \beta^{T-t} \left[(1-\beta) \frac{\overline{Y}^i}{\overline{C}^i} \rho^{T-t} y_t^i \right] \\ &- \sum_{T=t+1}^{\infty} \beta^{T-t} \left[\left(\sigma + \frac{\overline{Y}^i - \overline{C}^i}{\overline{C}^i} \right) \beta \left(m^i + l^i \rho^{T-t} \tilde{y}_t + e^i \mu^{T-t} \varepsilon_t \right) \right] \\ &- \left(\sigma + \frac{\overline{Y}^i - \overline{C}^i}{\overline{C}^i} \right) \beta r_t \end{aligned}$$

We have

$$(41)\quad \begin{aligned} c_t^i &= (1-\beta) b_t^i + \sum_{T=t}^{\infty} (\beta\rho)^{T-t} \left[(1-\beta) \frac{\overline{Y}^i}{\overline{C}^i} y_t^i \right] \\ &- \sum_{T=t+1}^{\infty} \beta^{T-t} \left[\left(\sigma + \frac{\overline{Y}^i - \overline{C}^i}{\overline{C}^i} \right) \beta m^i \right] \\ &- \sum_{T=t+1}^{\infty} (\beta\rho)^{T-t} \left[\left(\sigma + \frac{\overline{Y}^i - \overline{C}^i}{\overline{C}^i} \right) \beta l^i \tilde{y}_t \right] \\ &- \sum_{T=t+1}^{\infty} (\beta\mu)^{T-t} \left[\left(\sigma + \frac{\overline{Y}^i - \overline{C}^i}{\overline{C}^i} \right) \beta e^i \varepsilon_t \right] \\ &- \left(\sigma + \frac{\overline{Y}^i - \overline{C}^i}{\overline{C}^i} \right) \beta r_t \end{aligned}$$

and solving the infinite sums yields

$$(42)\quad \begin{aligned} c_t^i &= (1-\beta) b_t^i + \frac{\overline{Y}^i}{\overline{C}^i} \frac{1-\beta}{1-\beta\rho} y_t^i \\ &- \left(\sigma + \frac{\overline{Y}^i - \overline{C}^i}{\overline{C}^i} \right) \left(\frac{\beta^2}{1-\beta} m^i + \frac{\beta^2 \rho}{1-\beta\rho} l^i \tilde{y}_t + \frac{\beta^2 \mu}{1-\beta\mu} e^i \varepsilon_t + \beta r_t \right) \end{aligned}$$

This expression is aggregated over all agents by building the arithmetic mean.[4] Recognizing that $\tilde{b}_t = 0$ and that the mean of the steady state consumption is equal to the mean of steady state income, we get

$$(43)\quad \tilde{c}_t = \frac{1-\beta}{1-\beta\rho}\tilde{y}_t - \sigma\left(\frac{\beta^2}{1-\beta}\tilde{m} + \frac{\beta^2\rho}{1-\beta\rho}\tilde{l}\tilde{y}_t + \frac{\beta^2\mu}{1-\beta\mu}\tilde{e}\varepsilon_t + \beta r_t\right)$$

Solving for the interest rate and applying the market clearing condition yields the ALM

$$(44)\quad \begin{aligned} r_t &= \frac{1}{\sigma\beta}\left(\frac{1-\beta}{1-\beta\rho} - 1\right)\tilde{y}_t - \frac{\beta}{1-\beta}\tilde{m} - \frac{\beta\rho}{1-\beta\rho}\tilde{l}\tilde{y}_t - \frac{\beta\mu}{1-\beta\mu}\tilde{e}\varepsilon_t \\ &= \frac{1}{1-\beta\rho}\left(\frac{\rho-1}{\sigma} - \beta\rho\tilde{l}\right)\tilde{y}_t - \frac{\beta}{1-\beta}\tilde{m} - \frac{\beta\mu}{1-\beta\mu}\tilde{e}\varepsilon_t \end{aligned}$$

The *T*-mapping is constructed by comparing this equation to the PLM (37)

$$(45)\quad T\begin{pmatrix} m^i \\ l^i \\ e^i \end{pmatrix} = \left(-\frac{\beta}{1-\beta}\tilde{m},\quad \frac{\rho-1}{\sigma(1-\beta\rho)} - \frac{\beta\rho}{1-\beta\rho}\tilde{l},\quad -\frac{\beta\mu}{1-\beta\mu}\tilde{e}\right)'$$

E-Stability for a single agent is determined by the ODEs

$$(46)\quad \frac{d}{d\tau}\left(m^i\right) = \left(-\beta(1-\beta)^{-1}\tilde{m} - m^i\right)'$$

$$(47)\quad \frac{d}{d\tau}\left(l^i\right) = \left(\frac{\rho-1}{\sigma(1-\beta\rho)} - \frac{\beta\rho}{1-\beta\rho}\tilde{l} - l^i\right)'$$

$$(48)\quad \frac{d}{d\tau}\left(e^i\right) = \left(-\beta\mu(1-\beta\mu)^{-1}\tilde{e} - e^i\right)'$$

Here, stability will be determined simultaneously for all agents. In an economy with n individual agents E-stability is determined by the stability of the ODEs

---

[4] See appendix A.2. It is assumed that the terms $\overline{Y}^i / \overline{C}^i$ and $\left(\overline{Y}^i - \overline{C}^i\right)/\overline{C}^i$ are not correlated to the log deviation of the agent's individual income or to his beliefs.

$$(49)\quad \frac{d}{d\tau}\begin{pmatrix} m^1 \\ m^2 \\ \vdots \\ m^n \end{pmatrix} = \left[-\underset{n\times n}{\boldsymbol{I}} - \beta(1-\beta)^{-1}\frac{1}{n}\underset{n\times n}{\boldsymbol{J}} \right] \begin{pmatrix} m^1 \\ m^2 \\ \vdots \\ m^n \end{pmatrix}$$

$$(50)\quad \frac{d}{d\tau}\begin{pmatrix} l^1 \\ l^2 \\ \vdots \\ l^n \end{pmatrix} = \frac{\rho-1}{\sigma(1-\beta\rho)}\underset{n\times 1}{\boldsymbol{J}} + \left[-\underset{n\times n}{\boldsymbol{I}} - \frac{\beta\rho}{1-\beta\rho}\frac{1}{n}\underset{n\times n}{\boldsymbol{J}} \right] \begin{pmatrix} l^1 \\ l^2 \\ \vdots \\ l^n \end{pmatrix}$$

$$(51)\quad \frac{d}{d\tau}\begin{pmatrix} e^1 \\ e^2 \\ \vdots \\ e^n \end{pmatrix} = \left[-\underset{n\times n}{\boldsymbol{I}} - \beta\mu(1-\beta\mu)^{-1}\frac{1}{n}\underset{n\times n}{\boldsymbol{J}} \right] \begin{pmatrix} e^1 \\ e^2 \\ \vdots \\ e^n \end{pmatrix}$$

where I and J are defined as in appendix A.4. The ODEs are asymptotically stable if all eigenvalues of their Jacobian matrices have negative real parts. Here, the Jacobi matrices are simply the matrix expressions in the square brackets. The Jacobi matrix for the ODE (49) has one distinct eigenvalue of $\lambda = -(1-\beta)^{-1} < 0$ and an $(n-1)$-fold eigenvalue of $\lambda = -1 < 0$, that for the ODE (50) has one distinct eigenvalue of $\lambda = -(1-\beta\rho)^{-1} < 0$ and an $(n-1)$-fold eigenvalue of $\lambda = -1 < 0$, and that for the ODE (51) has one distinct eigenvalue of $\lambda = -(1-\beta\mu)^{-1} < 0$ and an $(n-1)$-fold eigenvalue of $\lambda = -1 < 0$. Since all eigenvalues are negative, all ODEs are globally asymptotically stable and the system will converge under learning to the REE solution.

Consumption in the REE can be determined by inserting the equilibrium parameter values and the REE solution for the interest rate $r_t = \sigma^{-1}(\rho-1)\tilde{y}_t$ into equation (42). We get

$$(52)\quad \begin{aligned} c_t^i &= (1-\beta)b_t^i + \frac{\bar{Y}^i}{\bar{C}^i}\frac{1-\beta}{1-\beta\rho}y_t^i - \left(\sigma + \frac{\bar{Y}^i - \bar{C}^i}{\bar{C}^i}\right)\left(\frac{\beta^2\rho}{1-\beta\rho}\sigma^{-1}(\rho-1)\tilde{y}_t + \beta\sigma^{-1}(\rho-1)\tilde{y}_t\right) \\ &= (1-\beta)b_t^i + \frac{\bar{Y}^i}{\bar{C}^i}\frac{1-\beta}{1-\beta\rho}y_t^i - \left(1 + \frac{\bar{Y}^i - \bar{C}^i}{\bar{C}^i}\sigma^{-1}\right)\left(\frac{\beta\rho}{1-\beta\rho} + 1\right)\beta(\rho-1)\tilde{y}_t \end{aligned}$$

Defining $\lambda \equiv \frac{1-\beta}{1-\beta\rho}$ so that $1-\lambda = -\frac{\beta(\rho-1)}{1-\beta\rho}$, simplifying and rearranging yields the expression for the consumption in the REE solution used in the text

$$(53)\quad c_t^i = (1-\beta)b_t^i + \frac{\overline{Y}^i - \overline{C}^i}{\overline{C}^i}\left[\lambda y_t^i + (1-\lambda)\sigma^{-1}\tilde{y}_t\right] + \lambda\left(y_t^i - \tilde{y}_t\right) + \tilde{y}_t$$

where λ can alternatively be expressed as $\lambda = \frac{\overline{R}-1}{(\overline{R}-1)+(1-\rho)}$.

For an alternative way to determine consumption in the REE replace the expectations in (36) by the rational expectations

$$(54)\quad E_t^i y_T^i = \rho^{T-t} y_t^i$$

and

$$(55)\quad E_t^i r_T = \sigma^{-1}(\rho-1)\tilde{y}_T = \sigma^{-1}(\rho-1)\rho^{T-t}\tilde{y}_t$$

which yields

$$(56)\quad \begin{aligned} c_t^i &= (1-\beta)b_t^i + \sum_{T=t}^{\infty}(\beta\rho)^{T-t}\left[(1-\beta)\frac{\overline{Y}^i}{\overline{C}^i}y_t^i - \left(1+\frac{\overline{Y}^i-\overline{C}^i}{\overline{C}^i}\sigma^{-1}\right)\beta(\rho-1)\tilde{y}_t\right] \\ &= (1-\beta)b_t^i + \frac{1-\beta}{1-\beta\rho}\frac{\overline{Y}^i}{\overline{C}^i}y_t^i - \frac{\beta(\rho-1)}{1-\beta\rho}\left(1+\frac{\overline{Y}^i-\overline{C}^i}{\overline{C}^i}\sigma^{-1}\right)\tilde{y}_t \end{aligned}$$

Using the definition for λ as before and rearranging finally results again in (53).

A.6. EEE Learning

A.6.1. Detailed Derivation of the ALM

When using EEE learning, an agent determines his optimal current consumption according to the Euler equation for the current period

(57) $c_t = \hat{E}_t c_{t+1} - \sigma r_t$

His PLM of his consumption is linear and contains in addition to the variables of the REE solution also a constant and a sunspot variable ε_t

(58) $c_t^i = m^i + l^i y_{t-1}^i + k^i \tilde{y}_{t-1} + p^i b_{t-1}^i + e^i \varepsilon_{t-1} + \upsilon_t^i$

so his expectation of next periods consumption is

(59) $\hat{E}_t^i c_{t+1}^i = m^i + l^i y_t^i + k^i \tilde{y}_t + p^i b_t^i + e^i \varepsilon_t$

Inserting into (57) yields

(60) $c_t = m^i + l^i y_t^i + k^i \tilde{y}_t + p^i b_t^i + e^i \varepsilon - \sigma r_t$

and aggregating over all agents and solving for the equilibrium interest rate gives

(61) $r_t = \sigma^{-1} \left[\tilde{m} + n^{-1} \sum_{j=1}^{n} l^j y_t^j + \left(\tilde{k} - 1 \right) \tilde{y}_t + n^{-1} \sum_{j=1}^{n} p^j b_t^j + \tilde{e} \varepsilon_t \right]$

The ALM of the agent's consumption can be calculated by inserting this interest rate into (60):

(62) $c_t^i = \left(m^i - \tilde{m} \right) + l^i y_t^i - n^{-1} \sum_{j=1}^{n} l^j y_t^j + \left(k^i - \tilde{k} + 1 \right) \tilde{y}_t + p^i b_t^i - n^{-1} \sum_{j=1}^{n} p^j b_t^j + \left(e^i - \tilde{e} \right) \varepsilon_t$

In this equation, the current wealth is predetermined by the log-linearized flow budget constraint

(63) $b_t^i = \frac{\overline{B}^i}{\overline{C}^i} r_{t-1} + \beta^{-1} \frac{\overline{Y}^i}{\overline{C}^i} y_{t-1}^i + \beta^{-1} b_{t-1}^i - \beta^{-1} c_{t-1}^i$

where interest rate and consumption are given by the respective ALMs (61) and (62), each moved backwards by one period. Inserting yields

$$
\begin{aligned}
b_t^i &= \frac{\overline{B}^i}{\overline{C}^i}\sigma^{-1}\left[\tilde{m} + n^{-1}\sum_{j=1}^{n} l^j y_{t-1}^j + \left(\tilde{k}-1\right)\tilde{y}_{t-1} + n^{-1}\sum_{j=1}^{n} p^j b_{t-1}^j + \tilde{e}\,\varepsilon_{t-1}\right] \\
&+ \beta^{-1}\frac{\overline{Y}^i}{\overline{C}^i} y_{t-1}^i + \beta^{-1} b_{t-1}^i \\
&- \beta^{-1}\left[\left(m^i - \tilde{m}\right) + l^i y_{t-1}^i + \left(k^i - \tilde{k} + 1\right)\tilde{y}_{t-1} + p^i b_{t-1}^i + \left(e^i - \tilde{e}\right)\varepsilon_{t-1}\right] \\
&+ \beta^{-1}\left[n^{-1}\sum_{j=1}^{n} l^j y_{t-1}^j + n^{-1}\sum_{j=1}^{n} p^j b_{t-1}^j\right]
\end{aligned} \tag{64}
$$

or slightly rearranged

$$
\begin{aligned}
b_t^i &= \kappa^i\,\tilde{m} - \beta^{-1} m^i + \beta^{-1}\left(\frac{\overline{Y}^i}{\overline{C}^i} - l^i\right) y_{t-1}^i + \beta^{-1}\left(1 - p^i\right) b_{t-1}^i \\
&+ \left(\kappa^i\left(\tilde{k}-1\right) - \beta^{-1} k^i\right)\tilde{y}_{t-1} + \left(\kappa^i\,\tilde{e} - \beta^{-1} e^i\right)\varepsilon_{t-1} \\
&+ \kappa^i\, n^{-1}\sum_{j=1}^{n} l^j y_{t-1}^j + \kappa^i\, n^{-1}\sum_{j=1}^{n} p^j b_{t-1}^j
\end{aligned} \tag{65}
$$

with $\kappa^i = \frac{\overline{B}^i}{\overline{C}^i}\sigma^{-1} + \beta^{-1}$

For the individual income, the average income and the sunspot variable we have

$$
\begin{aligned}
\varepsilon_t &= \mu\,\varepsilon_{t-1} + \pi_t \\
y_t^i &= \rho\, y_{t-1}^i + \nu_t^i \\
\tilde{y}_t &= \rho\,\tilde{y}_{t-1} + \tilde{\nu}_t
\end{aligned} \tag{66}
$$

Inserting these expressions into (60) finally gives

$$
\begin{aligned}
c_t^i = {} & \left(1-\beta^{-1}p^i\right)m^i - \left(1-\kappa^i\, p^i\right)\tilde{m} \\
& + l^i\left(\rho\, y_{t-1}^i + \nu_t^i\right) + \beta^{-1}\left(\frac{\overline{Y}^i}{\overline{C}^i} - l^i\right)p^i y_{t-1}^i \\
& + \left(k^i - \tilde{k} + 1\right)\left(\rho\,\tilde{y}_{t-1} + \tilde{\nu}_t\right) + \left(\kappa^i\left(\tilde{k}-1\right) - \beta^{-1}k^i\right)p^i\,\tilde{y}_{t-1} \\
& + \left(e^i - \tilde{e}\right)\left(\mu\,\varepsilon_{t-1} + \pi_t\right) + \left(\kappa^i\,\tilde{e} - \beta^{-1}e^i\right)p^i\,\varepsilon_{t-1} \qquad (67) \\
& + \beta^{-1}p^i\left(1-p^i\right)b_{t-1}^i \\
& + \left(\kappa^i\, p^i - \rho\right)n^{-1}\sum_{j=1}^{n} l^j y_{t-1}^j - n^{-1}\sum_{j=1}^{n} l^j \nu_t^j \\
& + \kappa^i\, p^i n^{-1}\sum_{j=1}^{n} p^j b_{t-1}^j - n^{-1}\sum_{j=1}^{n} p^j b_t^j
\end{aligned}
$$

or, rearranged,

$$
\begin{aligned}
c_t^i = {} & \left(1-\beta^{-1}p^i\right)m^i - \left(1-\kappa^i\, p^i\right)\tilde{m} \\
& + \left[\left(\rho-\beta^{-1}p^i\right)l^i + \beta^{-1}p^i\,\frac{\overline{Y}^i}{\overline{C}^i}\right]y_{t-1}^i + l^i\nu_t^i \\
& + \left[\left(\rho-\beta^{-1}p^i\right)k^i - \left(\rho-\kappa^i\, p^i\right)\left(\tilde{k}-1\right)\right]\tilde{y}_{t-1} + \left(k^i - \tilde{k} + 1\right)\tilde{\nu}_t \\
& + \left[\left(\mu-\beta^{-1}p^i\right)e^i - \left(\mu-\kappa^i\, p^i\right)\tilde{e}\right]\varepsilon_{t-1} + \left(e^i - \tilde{e}\right)\pi_t \qquad (68) \\
& + \beta^{-1}p^i\left(1-p^i\right)b_{t-1}^i \\
& - \left(\rho-\kappa^i\, p^i\right)n^{-1}\sum_{j=1}^{n} l^j y_{t-1}^j - n^{-1}\sum_{j=1}^{n} l^j \nu_t^j \\
& + \kappa^i\, p^i n^{-1}\sum_{j=1}^{n} p^j b_{t-1}^j - n^{-1}\sum_{j=1}^{n} p^j b_t^j
\end{aligned}
$$

which describes the ALM for the current consumption as a function of exogenous variables from period $t-1$.[5]

A.6.2. Simultaneous Learning of the Interdependent Parameters

In this section it will be demonstrated how the equilibria and the stability of the learning process can be determined for the parameters m^i, k^i, and e^i. Since the learning process of

[5] Note that the parameters of the PLMs are noted without a date and collected regardless of their date to simplify the notation. See footnote 160 on p. 91 in the text.

these parameters is characterized by an interdependence between the individual agents, it must be considered simultaneously for all agents. As an alternative to the approach used in the text, which considered the parameters of a single agent and their averages, one can simultaneously analyze the learning process for all agents.

The T-map for the constant m^i for an individual agent is

$$(69)\quad T\left(m^i\right) = \left(1 - \frac{1-\beta}{\beta}\right)m^i - \left(1 - \frac{1-\beta}{\beta} - (1-\beta)\sigma^{-1}\frac{\overline{B}^i}{\overline{C}^i}\right)\tilde{m}$$

with a fixed point determined by the condition

$$(70)\quad m^i = \left(1 - \frac{1-\beta}{\beta}\right)m^i - \left(1 - \frac{1-\beta}{\beta} - (1-\beta)\sigma^{-1}\frac{\overline{B}^i}{\overline{C}^i}\right)n^{-1}\sum_{i=1}^{n} m^i$$

or

$$(71)\quad m^i = \left(1 - \frac{\beta}{1-\beta} + \beta\sigma^{-1}\frac{\overline{B}^i}{\overline{C}^i}\right)n^{-1}\sum_{i=1}^{n} m^i$$

Aggregating over all agents yields

$$(72)\quad 0 = \left(\frac{-\beta}{1-\beta} + \beta\sigma^{-1}n^{-1}\sum_{i=1}^{n}\frac{\overline{B}^i}{\overline{C}^i}\right)n^{-1}\sum_{i=1}^{n} m^i$$

where the sum $\sum_{i=1}^{n}\left(\overline{B}^i / \overline{C}^i\right)$ can be approximated by zero.[6] So one gets as the fixed point

$$(73)\quad 0 = n^{-1}\sum_{i=1}^{n} m^i$$

6 See the assumptions of section 3.1.2 and the discussion in section A.2 of the appendix.

and by (71) $m^i = 0$. E-stability of the fixed point can be determined by considering the ODE

$$(74)\quad \frac{d}{d\tau}m^i = -\frac{1-\beta}{\beta}m^i - \left(1-\frac{1-\beta}{\beta}-(1-\beta)\sigma^{-1}\frac{\overline{B}^i}{\overline{C}^i}\right)\tilde{m}$$

Stacking this ODE for all n agents in the economy yields the system

$$(75)\quad \frac{d}{d\tau}\begin{pmatrix} m^1 \\ m^2 \\ \vdots \\ m^n \end{pmatrix} = \left[-\frac{1-\beta}{\beta}\underset{n\times n}{I} + \left(\frac{1-\beta}{\beta}-1\right)\frac{1}{n}\underset{n\times n}{J} + (1-\beta)\sigma^{-1}\frac{1}{n}\begin{pmatrix} \frac{\overline{B}^1}{\overline{C}^1} & \frac{\overline{B}^1}{\overline{C}^1} & \cdots & \frac{\overline{B}^1}{\overline{C}^1} \\ \frac{\overline{B}^2}{\overline{C}^2} & \frac{\overline{B}^2}{\overline{C}^2} & \cdots & \frac{\overline{B}^2}{\overline{C}^2} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\overline{B}^n}{\overline{C}^n} & \frac{\overline{B}^n}{\overline{C}^n} & \cdots & \frac{\overline{B}^n}{\overline{C}^n} \end{pmatrix}\right]\begin{pmatrix} m^1 \\ m^2 \\ \vdots \\ m^n \end{pmatrix}$$

The Jacobi matrix for this ODE is the expression in square brackets. In the course of the calculation of its eigenvalues[7] the sum $\sum_{i=1}^{n}\left(\overline{B}^i/\overline{C}^i\right)$ has to be determined, which is approximated by zero as before. Then the Jacobi matrix has one distinct eigenvalue of $\lambda = -1 < 0$ and an $(n-1)$-fold eigenvalue of $\lambda = -\beta^{-1}(1-\beta) < 0$. Hence the system is globally asymptotically stable and the parameters converge to their equilibrium values under learning with probability 1.

The T-map for the parameter e^i which represents the influence of the sunspot variable for an individual agent is

$$(76)\quad T\left(e^i\right) = \left(\mu - \frac{1-\beta}{\beta}\right)e^i - \left(\mu - \frac{1-\beta}{\beta} - (1-\beta)\sigma^{-1}\frac{\overline{B}^i}{\overline{C}^i}\right)\tilde{e}$$

with a fixed point determined by the condition

[7] A useful rule for this kind of matrix is given in section A.7, paragraph c).

$$(77)\quad e^i = \left(\mu - \frac{1-\beta}{\beta}\right)e^i - \left(\mu - \frac{1-\beta}{\beta} - (1-\beta)\sigma^{-1}\frac{\overline{B}^i}{\overline{C}^i}\right)n^{-1}\sum_{i=1}^{n} e^i$$

or

$$(78)\quad e^i = \left(1 - \frac{\beta}{1-\beta\mu} + \frac{\beta(1-\beta)}{1-\beta\mu}\sigma^{-1}\frac{\overline{B}^i}{\overline{C}^i}\right)n^{-1}\sum_{i=1}^{n} e^i$$

for $\mu \neq \beta^{-1}$, which is assured by the condition $|\mu| \leq 1$ for the process of the sunspot variable. Aggregating over all agents yields

$$(79)\quad 0 = \left(\frac{-\beta}{1-\beta\mu} + \frac{\beta(1-\beta)}{1-\beta\mu}\sigma^{-1}n^{-1}\sum_{i=1}^{n}\frac{\overline{B}^i}{\overline{C}^i}\right)n^{-1}\sum_{i=1}^{n} e^i$$

where the sum $\sum_{i=1}^{n}\left(\overline{B}^i / \overline{C}^i\right)$ can again be approximated by zero. So one gets as the fixed point

$$(80)\quad 0 = n^{-1}\sum_{i=1}^{n} e^i$$

and by (78) $e^i = 0$. E-stability of the fixed point can be determined by considering the ODE

$$(81)\quad \frac{d}{d\tau}e^i = \left(\mu - \frac{1}{\beta}\right)e^i - \left(\mu - \frac{1-\beta}{\beta} - (1-\beta)\frac{\overline{B}^i}{\overline{C}^i}\sigma^{-1}\right)\tilde{e}$$

Stacking this ODE for all *n* agents in the economy yields the system

$$(82)\quad \frac{d}{d\tau}\begin{pmatrix} e^1 \\ e^2 \\ \vdots \\ e^n \end{pmatrix} = \left[\left(\mu - \frac{1}{\beta}\right)\underset{n\times n}{I} + \left(\frac{1-\beta}{\beta} - \mu\right)\frac{1}{n}\underset{n\times n}{J} + (1-\beta)\sigma^{-1}\frac{1}{n}\begin{pmatrix} \frac{\bar{B}^1}{\bar{C}^1} & \frac{\bar{B}^1}{\bar{C}^1} & \cdots & \frac{\bar{B}^1}{\bar{C}^1} \\ \frac{\bar{B}^2}{\bar{C}^2} & \frac{\bar{B}^2}{\bar{C}^2} & \cdots & \frac{\bar{B}^2}{\bar{C}^2} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\bar{B}^n}{\bar{C}^n} & \frac{\bar{B}^n}{\bar{C}^n} & \cdots & \frac{\bar{B}^n}{\bar{C}^n} \end{pmatrix}\right]\begin{pmatrix} e^1 \\ e^2 \\ \vdots \\ e^n \end{pmatrix}$$

The Jacobi matrix for this ODE is the expression in square brackets. With $\sum_{i=1}^{n}\left(\bar{B}^i/\bar{C}^i\right) \approx 0$ as before, the Jacobi matrix has one distinct eigenvalue of $\lambda = -1 < 0$ and an $(n-1)$-fold eigenvalue of $\lambda = -\beta^{-1}(1-\beta\mu) < 0$. Hence the system is globally asymptotically stable and the parameters converge to their equilibrium values under learning with probability 1.

The *T*-map for the parameter k^i which represents the influence of the average income is

$$(83)\quad T\left(k^i\right) = \left(\rho - \frac{1-\beta}{\beta}\right)k^i + \left(\rho - \frac{1-\beta}{\beta} - (1-\beta)\sigma^{-1}\frac{\bar{B}^i}{\bar{C}^i}\right)\left(1 - \lambda - \tilde{k}\right)$$

with a fixed point determined by the condition

$$(84)\quad k^i = \left(\rho - \frac{1-\beta}{\beta}\right)k^i + \left(\rho - \frac{1-\beta}{\beta} - (1-\beta)\sigma^{-1}\frac{\bar{B}^i}{\bar{C}^i}\right)\left(1 - \lambda - n^{-1}\sum_{i=1}^{n}k^i\right)$$

or

$$(85)\quad k^i = \left(\frac{1-\beta\rho-\beta}{1-\beta\rho} + \frac{\beta}{1-\beta\rho}(1-\beta)\sigma^{-1}\frac{\bar{B}^i}{\bar{C}^i}\right)\left(n^{-1}\sum_{i=1}^{n}k^i - 1 + \lambda\right)$$

Aggregating over all agents yields

$$(86)\quad k^i = \left(\frac{1-\beta\rho-\beta}{1-\beta\rho} + \frac{\beta}{1-\beta\rho}(1-\beta)\sigma^{-1}n^{-1}\sum_{i=1}^{n}\frac{\bar{B}^i}{\bar{C}^i}\right)\left(n^{-1}\sum_{i=1}^{n}k^i - 1 + \lambda\right)$$

where the sum $\sum_{i=1}^{n}\left(\overline{B}^i/\overline{C}^i\right)$ can again be approximated by zero. So one gets as the fixed point

$$\text{(87)}\quad n^{-1}\sum_{i=1}^{n}k^i = \frac{1-\beta\rho-\beta}{\beta}\left(-1+\lambda\right)=\rho-\lambda$$

and by (85) for the individual parameter

$$\text{(88)}\quad \begin{aligned} k^i &= \left(\frac{1-\beta\rho-\beta}{1-\beta\rho}+\frac{\beta}{1-\beta\rho}\left(1-\beta\right)\sigma^{-1}\frac{\overline{B}^i}{\overline{C}^i}\right)\left(\rho-1\right) \\ &= \rho-\lambda-\left(1-\lambda\right)\left(1-\beta\right)\sigma^{-1}\frac{\overline{B}^i}{\overline{C}^i} \end{aligned}$$

E-stability of the fixed point can be determined by considering the ODE

$$\text{(89)}\quad \frac{d}{d\tau}k^i = \left(\rho-\frac{1}{\beta}\right)k^i + \left(\rho-\frac{1-\beta}{\beta}-\left(1-\beta\right)\sigma^{-1}\frac{\overline{B}^i}{\overline{C}^i}\right)\left(1-\lambda-\tilde{k}\right)$$

Stacking this ODE for all n agents in the economy yields the system

$$\text{(90)}\quad \begin{aligned} \frac{d}{d\tau}\begin{pmatrix} k^1 \\ k^2 \\ \vdots \\ k^n \end{pmatrix} &= \left[\left(\rho-\frac{1}{\beta}\right)\underset{n\times n}{\boldsymbol{I}}+\left(\frac{1-\beta}{\beta}-\rho\right)\frac{1}{n}\underset{n\times n}{\boldsymbol{J}}+\left(1-\beta\right)\sigma^{-1}\frac{1}{n}\begin{pmatrix} \frac{\overline{B}^1}{\overline{C}^1} & \frac{\overline{B}^1}{\overline{C}^1} & \cdots & \frac{\overline{B}^1}{\overline{C}^1} \\ \frac{\overline{B}^2}{\overline{C}^2} & \frac{\overline{B}^2}{\overline{C}^2} & \cdots & \frac{\overline{B}^2}{\overline{C}^2} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\overline{B}^n}{\overline{C}^n} & \frac{\overline{B}^n}{\overline{C}^n} & \cdots & \frac{\overline{B}^n}{\overline{C}^n} \end{pmatrix}\right]\begin{pmatrix} k^1 \\ k^2 \\ \vdots \\ k^n \end{pmatrix} \\ &+\left(\rho-\frac{1-\beta}{\beta}-\left(1-\beta\right)\sigma^{-1}\frac{\overline{B}^i}{\overline{C}^i}\right)\left(1-\lambda\right)\underset{n\times 1}{\boldsymbol{J}} \end{aligned}$$

The Jacobi matrix for this ODE is the expression in square brackets. With $\sum_{i=1}^{n}\left(\overline{B}^i/\overline{C}^i\right)\approx 0$ as before, the Jacobi matrix has one distinct eigenvalue of $\lambda=-1<0$

and an $(n-1)$-fold eigenvalue of $\lambda = -\beta^{-1}\left(1-\beta\rho\right) < 0$. Hence the system is globally asymptotically stable and the parameters converge to their equilibrium values under learning with probability 1.

A.7. Helpful rules for the calculation of eigenvalues

For the calculation of the eigenvalues the following rules are used.

a) The $n \times n$ matrix J

(91) $$\underset{n\times n}{J} = \begin{pmatrix} 1 & 1 & \cdots & 1 \\ 1 & 1 & \cdots & 1 \\ \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & \cdots & 1 \end{pmatrix}$$

has one eigenvalue of $\lambda = n$ and $n-1$ eigenvalues of $\lambda = 0$.

b) Consider a matrix A with a set of eigenvalues $\{\lambda_i\}$. If A is multiplied by a scalar a, then the eigenvalues of the resulting matrix are $\{a\lambda_i\}$. The set of eigenvalues of a matrix B with

(92) $$B = aA - I$$

is $\{a\lambda_i - 1\}$. Thus, a matrix C with

(93) $$C = a\,n^{-1}J - bI$$

has one eigenvalue of $a-b$ and $n-1$ eigenvalues of $-b$.

c) Consider a vector $d = \begin{pmatrix} d^1 & d^2 & \cdots & d^n \end{pmatrix}'$ and a $n \times n$ matrix D with

(94) $$D = d \cdot \underset{1\times n}{J}$$

so that all columns of D are identical. Since D is of rank 1, there is exactly one eigenvalue λ with $\lambda \neq 0$. This eigenvalue must fulfil the requirement

$$(95)\quad D\,x = \lambda\,x$$

for some eigenvector $x = \begin{pmatrix} x^1 & x^2 & \cdots & x^n \end{pmatrix}'$. Thus, we have

$$(96)\quad \left(\sum_{i=1}^{n} x^i\right) d = \lambda\,x$$

which holds for $\lambda = \sum_{i=1}^{n} x^i$ and $d = x$. The matrix D has therefore one distinct eigenvalue $\lambda = \sum_{i=1}^{n} d^i$ with the eigenvector *d*, and an $(n-1)$-fold eigenvalue $\lambda = 0$.

■