

IHS Economics Series
Working Paper 186
February 2006

Functional Rational Expectations Equilibria in Market Games

Jamsheed Shorish





INSTITUT FÜR HÖHERE STUDIEN
INSTITUTE FOR ADVANCED STUDIES
Vienna

Impressum

Author(s):

Jamsheed Shorish

Title:

Functional Rational Expectations Equilibria in Market Games

ISSN: Unspecified

2006 Institut für Höhere Studien - Institute for Advanced Studies (IHS)

Josefstädter Straße 39, A-1080 Wien

E-Mail: office@ihs.ac.at

Web: www.ihs.ac.at

All IHS Working Papers are available online: http://irihs.ihs.ac.at/view/ihs_series/

This paper is available for download without charge at:

<https://irihs.ihs.ac.at/id/eprint/1690/>

186

Reihe Ökonomie
Economics Series

Functional Rational Expectations Equilibria in Market Games

Jamsheed Shorish

186

Reihe Ökonomie
Economics Series

Functional Rational Expectations Equilibria in Market Games

Jamsheed Shorish

February 2006

Institut für Höhere Studien (IHS), Wien
Institute for Advanced Studies, Vienna

Contact:

Jamsheed Shorish
Department of Economics and Finance
Institute for Advanced Studies
Stumpergasse 56
A-1060 Vienna, Austria
☎: +43/1/599 91-250
fax: +43/1/599 91-555
email: shorish@ihs.ac.at

Founded in 1963 by two prominent Austrians living in exile – the sociologist Paul F. Lazarsfeld and the economist Oskar Morgenstern – with the financial support from the Ford Foundation, the Austrian Federal Ministry of Education and the City of Vienna, the Institute for Advanced Studies (IHS) is the first institution for postgraduate education and research in economics and the social sciences in Austria. The **Economics Series** presents research done at the Department of Economics and Finance and aims to share “work in progress” in a timely way before formal publication. As usual, authors bear full responsibility for the content of their contributions.

Das Institut für Höhere Studien (IHS) wurde im Jahr 1963 von zwei prominenten Exilösterreichern – dem Soziologen Paul F. Lazarsfeld und dem Ökonomen Oskar Morgenstern – mit Hilfe der Ford-Stiftung, des Österreichischen Bundesministeriums für Unterricht und der Stadt Wien gegründet und ist somit die erste nachuniversitäre Lehr- und Forschungsstätte für die Sozial- und Wirtschaftswissenschaften in Österreich. Die **Reihe Ökonomie** bietet Einblick in die Forschungsarbeit der Abteilung für Ökonomie und Finanzwirtschaft und verfolgt das Ziel, abteilungsinterne Diskussionsbeiträge einer breiteren fachinternen Öffentlichkeit zugänglich zu machen. Die inhaltliche Verantwortung für die veröffentlichten Beiträge liegt bei den Autoren und Autorinnen.

Abstract

The rational expectations equilibrium has been criticized as an equilibrium concept in market game environments. Such an equilibrium may not exist generically, or it may introduce unrealistic assumptions about an economic agent's knowledge or computational ability. We define a rational expectations equilibrium as a probability measure over uncertain states of nature which exploits all available information in a market game, and which exists for almost all economies. Furthermore, if retrading is allowed, it is possible for agents to compute such a 'functional rational expectations equilibrium' using straightforward numerical fixed point algorithms.

Keywords

Market game, rational expectations equilibrium, bayesian updating, learning, computation

JEL Classification

G12, D83, C63

Comments

I would like to thank session participants of the SAET 2003 Conference in Rhodes, Greece, July 2003 for their comments and suggestions. All remaining errors are my own.

Contents

1	Introduction	1
2	The Market Game	2
2.1	The Agent's Problem	3
2.1.1	Expectations Formation	4
2.2	Market Clearing and Price Formation	4
2.3	Optimal Bids and Asks	5
3	The Functional REE Concept	6
4	Existence of an REE	8
5	Retrading and Functional REE	12
6	Conclusion	16
	References	17

1 Introduction

The application of the rational expectations equilibrium (REE) concept to market games has not been unambiguously supported by researchers seeking to better understand, and even define, an appropriate equilibrium selection criterion. Indeed, Dubey, Geanakoplos, and Shubik [1987] criticize the concept of REE in market games (and more generally, in a continuous, anonymous price formation mechanism). They focus instead upon equilibrium selection according to the Nash equilibria (NE) of the market game. While NE are available as solutions either generically (for a continuum of agents), or for an open set (for a finite number of agents), REE cannot be implemented by price formation mechanisms which generate these NE. Conversely, mechanisms which do admit REE (for example, if traders submit their entire demand function to an auctioneer) are argued to be too computationally complex to be realistic.

In addition, there is some evidence that NE alone may generate more realistic time series properties of observed prices. Jackson and Peck [1999] examine an asymmetric information, 2 period market game with an asset and a single consumption good. With an infinite number of agents and no noise traders, they demonstrate that information is not fully revealed in a (Bayesian) NE—instead, asset prices demonstrate *excess volatility* relative to the dividend process of the asset. By contrast, REE cannot be obtained by the price formation process.

The difficulties with using the REE concept in market games, as in the examples above, usually hinge upon three main criticisms. First, the definition of an REE is circular: the information content of prices is itself used to determine those prices. Second, as initially examined by Grossman and Stiglitz [1980] and subsequently by many others, REE may fail to exist when all agents are fully informed. In the absence of ‘noise’ or ‘liquidity’ traders, markets do not open. Finally, REE may be very difficult to find: if there is any weight to the notion of bounded rationality, it is likely that normal, ‘real-world’ economic participants will be unable to calculate any REE which might exist.

In this paper we extend the definition of an REE from the usual point representation

to one of a ‘functional REE’, in a simple market game environment. Starting from a similar model as Jackson and Peck [1999], but with a finite number of agents, a functional relation is defined between the expectations of the agents (here, defined by prior probability measures) and the ‘result’ of the economy, which is also given by a probability measure. When expectations match the result of the economy, i.e. when the relation possesses a functional fixed point, the economy is said to possess a functional REE: individual agent expectations about the economy’s law of motion are fulfilled, and all available information is being used.

In addition, if retrading of assets is allowed, then this functional REE can be learned—it is possible for agents to sample the functional relation by submitting many different priors, and observing the outcome. It is shown how to design an algorithm which can estimate the functional relation itself, and then find the REE probability measure by applying standard numerical fixed point algorithms.

Section 2 introduces the market game environment, while Section 3 introduces the functional REE as an equilibrium probability measure. Section 4 then demonstrates the existence of an REE probability measure, and Section 5 shows one possible method for numerically estimating such an REE. The final section concludes.

2 The Market Game

We consider a static game environment without retrading as in Jackson and Peck [1999] (retrading will be examined in Section 5). There exist $I < \infty$ agents who trade in a single consumption good, single asset economy. Each agent i possesses a preference relation over the consumption good, which is assumed to be representable by an individual utility function u_i . The utility function is further assumed to be continuous, strictly increasing and concave.

Each agent i also possesses an endowment e_i of the consumption good and the asset, defined as

$$e_i := (\bar{b}, \bar{q}), \quad 0 < \bar{b}, \bar{q} < \infty.$$

The asset pays a stochastic dividend d in the single good, drawn from a finite set of M possible dividends $\mathcal{D} := \{d^1, \dots, d^M\}$. Each agent i receives a random signal s_i about the value of the dividend, drawn from a finite set of N signals $\mathcal{S} := \{s^1, \dots, s^N\}$. Each signal s_i is drawn from \mathcal{S} according to an independently and identically distributed (i.i.d.) process $\mathbb{P}(s|d) : \mathcal{S} \times \mathcal{D} \rightarrow [0, 1]$, which is the probability of observing a private signal s conditional upon the dividend draw $d \in \mathcal{D}$. It is assumed that this distribution is common knowledge among all agents. The vector of all private signals drawn by the agents is denoted $s := (s_1, \dots, s_I)$. Occasionally, d is used to denote a ‘generic’ or representative member of \mathcal{D} , or will be used as an argument to emphasize a functional dependence (viz. the prior probability measure $\pi_i(d)$ defined below).

The consumption good and the asset are traded at a series of ‘trading posts’, which accept offers of the consumption good in return for the asset. These offers are denoted ‘bids’ b_i of an agent i for the asset. In addition, a trading post will also accept offers of the asset in return for the consumption good. For each agent i these offers are denoted ‘asks’ q_i . We let $b = (b_1, \dots, b_I)$ represent the vector of bids, and $q = (q_1, \dots, q_I)$ represent the vector of asks.

2.1 The Agent’s Problem

Each agent i seeks to maximize their expected utility of consumption, conditional upon their private signal:

$$\max_{\{b_i, q_i\}} \mathbb{E}[u_i(c_i)|s_i], \quad (2.1)$$

where feasible bids and asks obey $b_i \in [0, \bar{b}]$, $q_i \in [0, \bar{q}]$. Final consumption c_i depends upon both the realization of the dividend and the market clearing price (see below).

2.1.1 Expectations Formation

Agents are endowed with a set of priors over \mathcal{D} , denoted $\pi_i : \mathcal{D} \rightarrow [0, 1]$. It is assumed that these priors are non-zero everywhere, and

$$\sum_{j=1}^M \pi_i(d^j) = 1.$$

Since the conditional distribution of the signal s_i given any dividend value $d \in \mathcal{D}$ is known, after observing the signal an agent updates their prior using Bayes' Rule:

$$\pi_i(d|s_i) = \frac{\pi(s_i|d)\pi_i(d)}{\pi_i(s_i)}, \quad (2.2)$$

where

$$\pi_i(s) := \sum_{j=1}^M \pi(s|d^j)\pi_i(d^j) \quad (2.3)$$

is agent i 's unconditional probability measure over the signal space \mathcal{S} .

Using the posterior distribution $\pi_i(d|s_i)$, the agent's problem (2.1) may also be written as

$$\max_{\{b_i, q_i\}} \sum_{j=1}^M u_i(c_i(d^j))\pi_i(d^j|s_i). \quad (2.4)$$

2.2 Market Clearing and Price Formation

The market environment is a *market game*—agents are assumed to trade at the trading post in order to achieve their desired bundles. They submit bids and asks to the trading post, which calculates the ratio of bids to asks and set this as the market clearing price:

$$p = \frac{\sum_{i=1}^I b_i}{\sum_{i=1}^I q_i} = \frac{B}{Q}, \quad (2.5)$$

where $B = \sum_{i=1}^I b_i$ is the aggregate bid, and $Q = \sum_{i=1}^I q_i$ is the aggregate ask. Note that if B or Q is equal to 0 then $p = 0$.

After submitting their bids and asks each agent i receives a share of the consumption good equal to:

$$c_i = \bar{b} - b_i + \hat{d} \left(\frac{b_i}{p} + \bar{q} - q_i \right) + pq_i, \quad (2.6)$$

where the right hand side denotes (in turn) the net consumption good trade $\bar{b} - b_i$, the proceeds from the asset trading once the dividend value \hat{d} has been realized, $\hat{d} \left(\frac{b_i}{p} + \bar{q} - q_i \right)$, and finally the payoff from assets held, pq_i .

2.3 Optimal Bids and Asks

We proceed under the assumption that the individual consumers are intelligent enough to understand what their bids and asks will depend upon, but do not have enough information to completely specify this dependence.

From the consumer's optimization problem, we have the following first order conditions:

$$\mathbb{E} \left[\frac{\partial u_i}{\partial c_i} \left(\frac{d}{p} - 1 + b_i \left(-\frac{d}{p^2} \frac{\partial p}{\partial b_i} \right) \right) \right] = 0, \quad (2.7)$$

$$\mathbb{E} \left[\frac{\partial u_i}{\partial c_i} \left(p - d + q_i \frac{\partial p}{\partial q_i} \right) \right] = 0, \quad (2.8)$$

which may be simplified using the fact that $p = B/Q$ to

$$\mathbb{E} \left[\frac{\partial u_i}{\partial c_i} \left(\frac{d}{p} - 1 + b_i \left(-\frac{d}{p^2} \frac{1}{Q} \right) \right) \right] = 0, \quad (2.9)$$

$$\mathbb{E} \left[\frac{\partial u_i}{\partial c_i} \left(p - d - q_i \frac{p}{Q} \right) \right] = 0. \quad (2.10)$$

Solving these implicit equations for (b_i, q_i) yields the dependence of the individual's optimal strategy of bids and asks upon the posterior distribution $\pi_i(d|s_i)$ and the price level p :

$$b_i^* = b_i(\pi_i(d|s_i), p), q_i^* = q_i(\pi_i(d|s_i), p).$$

Using the fact that the price level will be determined from the bids and asks of all other agents, we can specify the *ex ante* dependence that the bid and ask functions will have upon other agents:

$$b_i^* = b_i(\pi(d|s)), q_i^* = q_i(\pi(d|s)),$$

where for simplicity of exposition we have defined

$$\pi(d|s) := (\pi_1(d|s_1), \dots, \pi_I(d|s_I)) \tag{2.11}$$

to be the vector of posterior probabilities over the dividend held by all consumers, conditional upon the vector of signals s received. Note that as both \mathcal{D} and \mathcal{S} are finite sets, $\pi(d|s)$ is a list of $I \times M \times N$ values.

For all but the entries corresponding to their own conditional probabilities and signal, the posterior probabilities $\pi(d|s)$ are unknown to each agent. Since the market game stipulates that each consumer submit one bid and ask, to close the model consumers must have some way of resolving this uncertainty about others' expectations and signals.

3 The Functional REE Concept

The information usable to each agent is summarized in the signal s_i , in the signal's conditional distribution $\pi(s_i|d)$, and in the prior probability distribution $\pi_i(d)$. As it is assumed that each agent may have a different prior distribution, there is no 'collapse' of the economy whereby prices are fully revealing and individuals either refuse to trade or are indifferent to trade (see e.g. Grossman and Stiglitz [1980]). Rather, the price will in general not be fully revealing (see Jackson and Peck [1999]), but will realistically reveal something about the underlying distribution of dividends prior to the actual dividend realization \hat{d} being announced.

In fact, this revelation is sufficient to define a rational expectations in the following way. Rather than focus solely upon the price outcome as being ‘fully revealing’ or not (which is unrealistic), an agent can have an *ex ante* probability measure $\pi(d)$ over the dividend space \mathcal{D} which turns out to be the ‘correct’ measure, *according to all information available to the agent*. That is, after submitting bids and asks and observing the market equilibrium price, the posterior probability over the dividend state after incorporating the equilibrium price as new information is the same as the prior distribution which was used when trades were submitted to the market. We call such a situation a ‘(functional) rational expectations equilibrium’ (REE), where the term ‘functional’ is indicative of the fact that the equilibrium is actually defined as a measure over all dividend states—here such a ‘function’ is merely a set of numbers as \mathcal{D} is discrete, but this may certainly be generalized if \mathcal{D} is continuous.

It is important that this equilibrium definition be consistent, i.e. that it be possible for agents to construct their bids and asks in a logical, or even ‘optimal’ fashion, according to their expectations. To this end, let us first define the available information that the agent may use after the market has cleared, and back out from this the optimal expectations which are to be used in the optimal bid and ask selection process.

Definition 1. *A temporary equilibrium under heterogeneous priors is a price p^* given by*

$$p^* := p(\pi(d|s)), \tag{3.1}$$

where $\pi(d|s)$ is as given in (2.11).

The appellation ‘temporary equilibrium’, which is used in dynamic models of learning and convergence to REE, is actually more appropriate for Section 5 when retrading is allowed. Here, of course, the model is static, so there is nothing ‘temporary’ about the equilibrium price p^* . But we introduce the term here because the REE concept really only has power in a dynamic environment where agents can potentially learn it—as the saying goes, an REE is the equilibrium which is learned *when all systematic learning errors have been eliminated*. This can only take place in a dynamic environment.

We can now begin to add structure to the definition of an REE. If the equilibrium is to contain all available information, and furthermore if it is to avoid systematic errors, then the prior probability measure over the dividend which agents infer *ex ante* from the economic system as a whole must equal the probability measure which the economic system signals through the temporary equilibrium price. In other words, an REE is precisely that prior probability measure $\pi^*(d)$ which ensures that when $\pi_i = \pi^* \forall i$, all agents are in fact observing this probability *ex post*, i.e. after updating their priors upon observing the price level p^* . Note that this probability measure need *not* equal the actual probability distribution over the dividend process $\pi(d)$. In fact, it is inconsistent to restrict the REE measure to the actual prior probability measure $\pi(d)$ —there may be many REE, depending upon how many prior probability measures are supported by the temporary equilibrium price p^* .

Thus, an REE probability measure is a measure $\pi^*(d)$ which obtains if all agents 1) use $\pi^*(d)$ as their prior measure for $d \in \mathcal{D}$, 2) observe their private signal s_i and use Bayes' Rule to update this prior to $\pi_i(d|s_i; \pi^*)$, 3) submit optimal bids and asks, and 4) condition upon the resulting temporary equilibrium price to perform Bayes' Rule once more, such that

$$\pi_i(d|p^*, s_i; \pi_i^*) = \pi^*(d) \forall i, \forall d \in \mathcal{D}.$$

4 Existence of an REE

The main result of the paper is that the set of REE probability measures $\pi^*(d)$ may be shown to exist, and may be found using relatively straightforward (but complex) computational techniques. We first show that there exists a set of REE probability measures for every ‘well-behaved’ market game without the common prior, common knowledge assumption. Second, using arguments developed for functional REE in Kelly and Shorish [2000], we demonstrate how to find these REE probability measures using the temporary equilibria and private signal data available to each agent.

To establish existence of an REE probability measure we return once more to the

optimal bid and ask functions for a consumer i from Section 2.3:

$$b_i^* = b_i(\pi(d|s)), q_i^* = q_i(\pi(d|s)).$$

We first consider the optimal bid function (treatment of the optimal ask function proceeds in the same fashion). The optimal bid function may be viewed as a function of the set of all possible signals received by all other agents, the set of all possible conditional probability measures, and the respective measures imposed upon these two sets.

Every agent i can form the probability that another agent (say agent j) observes a signal s_j , conditional upon observing their own signal s_i —this is because signals are iid draws with conditional distribution $\pi(s_i|d) \forall i$, and all agents know this:

$$\pi_i(s_j|s_i) := \sum_{k=1}^M \pi(s_j|d^k) \pi_i(d^k|s_i). \quad (4.1)$$

This probability measure defines the uncertainty a consumer has about the signals received by other agents. On the other hand, there also exists uncertainty about the conditional probability $\pi_j(d|s_j)$ held by agent j , which also forms part of the optimal bid function b_i^* .

Definition 2. *The conditional probability held by agent i about agent j 's unconditional probability distribution over the dividend space \mathcal{D} is ¹*

$${}^i\pi_j(d) := \sum_{k=1}^N \pi_j(d|s^k) \pi_i(s^k|s_i) \forall d \in \mathcal{D}. \quad (4.2)$$

Armed with this definition we can start to form agent i 's expectations about agent j 's conditional probability over the dividend space. First, using Bayes' Rule once more we can rewrite $\pi_j(d|s^k)$ in the above definition:

$$\pi_j(d|s^k) = \frac{\pi(s^k|d) \pi_j(d)}{\sum_{m=1}^M \pi(s^k|d^m) \pi_j(d^m)}. \quad (4.3)$$

¹Note the superscripts here—the summation in (4.2) is performed over all possible signals that agent j could have obtained, weighted by agent i 's conditional probability that agent j has received that signal.

We next make the following observations. Although agent i does not know the prior probability distribution $\pi_j(d)$ used by agent j , there is a best guess of this distribution given by ${}^i\pi_j(d)$ in (4.2). Hence, we can set $\pi_j(d) = {}^i\pi_j(d)$ in (4.3). In addition, agent i wishes to solve out for $\pi_j(d|s^k)$ in (4.2) and (4.3), i.e the agent wishes to find a number ${}^i\pi_j(d|s^k)$ such that $\forall k$,

$${}^i\pi_j(d|s^k) = \frac{\pi(s^k|d) \left(\sum_{n=1}^N {}^i\pi_j(d|s^n) \pi_i(s^n|s_i) \right)}{\sum_{m=1}^M \pi(s^k|d^m) \left(\sum_{n=1}^N {}^i\pi_j(d^m|s^n) \pi_i(s^n|s_i) \right)}, \quad \forall k = 1 \dots N, \forall d \in \mathcal{D}. \quad (4.4)$$

This is a set of $M \times N$ equations in ${}^i\pi_j(d|s_k)$, which are (for each agent i and j) simply a set of $M \times N$ unknowns. Since these equations are analytic in the unknowns, there exists a solution ${}^i\pi_j^*$ to (4.4). This solution is completely expressible in terms of the prior probability $\pi_i(d)$ held by agent i , and the commonly known conditional probability $\pi(s_i|d)$, through both ${}^i\pi_j(d)$ and $\pi_i(s|s_i)$.

From the above analysis we see that an agent can form consistent expectations about other agents' conditional probability distributions $\pi_j(d|s_j)$, using only the publicly available information $\pi(s|d)$ and the prior probability $\pi_i(d)$. This allows the agent to submit an optimal bid (or ask) by forming the proper conditional expectations given in the consumer's problem of Section 2.

After the bids and asks have been submitted, the market clears and the temporary equilibrium price p^* is found. This price depends both upon the prior probability measures $\pi_i(d)$ of all agents in the economy, and also upon all signals s_i received by agents and conditioned upon when trades were submitted. In what follows, we shall suppress the dependence of p^* upon the prior measures and concentrate upon the signals, and let $p^* = p(s)$ denote the temporary equilibrium price as a function of the vector of signals s .

Each agent now uses the temporary equilibrium price to update their probability measure $\pi_i(d|s_i)$, again using Bayes' Rule:

$$\pi_i(d|p^*, s_i) = \frac{\pi_i(p^*|d, s_i) \pi_i(d, s_i)}{\sum_{k=1}^M \pi_i(p^*|d^k, s_i) \pi_i(d^k, s_i)}. \quad (4.5)$$

It is straightforward to show that for a temporary equilibrium price p^* ,

$$\pi_i(p^* = p(s)|d, s_i) = \sum_{s \in p^{-1}(p^*)} \prod_{j=1}^I \pi(s_j|d), \quad (4.6)$$

where the summation is over all signal vectors $s \in \mathcal{S}$ which are compatible with the observed temporary equilibrium price p^* .

In addition, we also know that

$$\pi_i(d, s_i) = \pi(s_i|d)\pi_i(d). \quad (4.7)$$

Substitution of (4.6) and (4.7) into (4.5) yields a posterior distribution over the dividend space which is once more completely specified by $\pi(s_i|d)$ and $\pi_i(d)$.

Definition 3. *A rational expectations equilibrium (REE) probability measure is a measure $\pi^*(d)$ such that for every agent $i = 1 \dots I$,*

$$\begin{aligned} \pi_i(d) &= \pi^*(d), \\ \pi_i(d|p^*, s_i; \pi^*(d)) &= \pi^*(d), \end{aligned}$$

where we have emphasized the dependence of the posterior probability measure $\pi_i(d|p^*, s_i)$ on the prior probability measure $\pi^*(d)$.

The existence of an REE probability measure is relatively easy to verify.

Theorem 4.1. *For almost every economy, there exists a rational expectations equilibrium probability measure $(\pi_1, \dots, \pi_I) = (\pi^*, \dots, \pi^*)$ which is a fixed point of the mapping*

$$T \circ (\pi_1, \dots, \pi_I) := (\pi_1(d|p(s), s_1; \pi_1), \dots, \pi_I(d|p(s), s_I; \pi_I)). \quad (4.8)$$

The proof involves a straightforward application of the Knaster-Kuratowski-Mazurkiewicz lemma—see Border [1985], Corollary 5.7 p. 27. The only restrictions are 1) that for each prior probability measure $\pi_i(d)$, the set of posterior distributions be sufficiently rich to

encompass any convex combination of prior probability measures for any subset of agents, and 2) that there exists at least one possible prior probability measure $\hat{\pi}(d)$ and at least one agent i so that the set of possible posteriors $\pi_i(d|p(s), s_i; \hat{\pi}(d))$ is compact. Because the conditional expectations for the optimal bids and asks and the Bayesian updating of the distributions involve analytic functions, their power series representations are nearly always convergent and hence well-defined—the only possible problems occur when the normalization of the Bayesian updating fails, i.e. when the prior places zero weight on a possible dividend value, or when either bids or asks approach the lower boundary of the endowment set for all agents (in which case there is no informative price as $p^* = 0$).

5 Retrading and Functional REE

Having found that functional REE exist, i.e. that there almost always exist probability measures which incorporate all available information and support prior expectations of all agents, it remains to be shown that such equilibria can actually be learned. In order to do so, we open the market up to retrading, so that agents live in a dynamic environment which allows them to learn the mapping between their prior probability measure, and the temporary equilibrium posterior measure.

Retrading means that agents can resubmit bids and asks to the market after each trade has been completed. Once retrading is allowed, stationary equilibrium concepts have to be adjusted. For example, Dubey [1980] has shown that the NE from finite player market games are generically inefficient when retrading is allowed. This result was extended by Dubey and Rogowski [1990], showing that NE efficiency only occurs if and only if initial agent endowments are themselves efficient. Finally, if a continuum of agents is allowed, Dubey, Sahi, and Shubik [1993] show that the NE regains its preeminence under retrading prior to final consumption—in this case, NE are both efficient and Walrasian.

Although economies with a finite number of agents are inefficient, retrading does allow agents to improve the efficiency of their individual outcomes—Ghosal and Morelli [2004] show that agents have an incentive with retrading to recontract from the NE, in order to

move closer to the Pareto-frontier. After a finite number of periods the Pareto-frontier may not be reached, but any Pareto-efficient outcome may be approximated arbitrarily closely if the number of retrading periods is ‘large enough’.

At this stage cannot explicitly address these efficiency issues in what follows—rather, we consider the retrading environment as a ‘test ground’ for an agent to attempt to learn the mapping which takes their prior probability measure as an input, and returns (via the temporary equilibrium price) a posterior measure. The question of how convergence to (or toward) the Pareto frontier may be obtained under the learning paradigm introduced below is an interesting topic of future research.

We consider once more a finite agent market game, but now with a (finite or infinite) number of retrading periods $t = 1, 2, \dots$. For agents to attempt to learn the REE, which is a collection of functions representing the probability measure over the dividend space \mathcal{D} , we consider the following strategy adapted from Kelly and Shorish [2000]:

1. Agent i begins period t with a ‘prior’ probability measure π_{it} . If this is the first trading period then $\pi_{i1} = \pi_i(d|s_i)$, i.e. the updated prior probability measure $\pi_i(d)$ after the signal s_i is observed.
2. At each trading point t , agent i submits optimal bid and ask strategies as in Section 2, conditional upon the prior π_{it} .
3. Agents i observes the temporary equilibrium price p^* and uses it to update the prior as in Section 4, leading to the outcome $T \circ \pi_{it} = \pi_{it}^T := \pi_i(d|p^*, s_i; \pi_{it}(d))$.
4. Agent i sets a new prior for period $t + 1$: this may be accomplished with either a randomly selected probability measure, or else by setting $\pi_{it+1} = \pi_{it}^T$.
5. After τ periods, the agent estimates the operator T using the data set $(\pi_{it}, \pi_{it}^T)_{t=1 \dots \tau}$.
6. Finally, standard numerical methods allow agents to find the fixed point of the operator T —this is the (functional) REE of the economy, π^* .

Agents thus use retrading as a way to sample the space of possible posterior distributions, conditional upon the prior distributions they have submitted.

The main question to be answered is whether or not it is possible for agents to actually implement this algorithm in practice. There are two requirements to ensure that this algorithm is actually feasible:

Assumption 5.1. *For all π_{it} bounded (e.g. in the sup norm), π_{it}^T is also bounded—that is, the operator T sends bounded functions to bounded functions.*

Assumption 5.2. *Both π_{it} and π_{it}^T are Borel-measurable, that is, they possess at most a finite number of discontinuities.*

Assumption (5.1) simply states that the probability measures are well-defined. Note that in a continuous space of dividends this excludes using a Dirac-type of distribution, which is not implementable as a computable function except as a limit of integrable functions (which are in fact bounded). Here, though, the discreteness of \mathcal{D} ensures that this assumption is automatically verified. Assumption (5.2) is also satisfied by probability measures, as there are a finite number of dividends by assumption.

With these two assumptions in hand it is possible to implement the above algorithm, with the assistance of a class of functions which can serve as a ‘basis’ for the probability measures of interest. In Kelly and Shorish [2000], we explain in great detail how to accomplish this implementation, for a general class of functional REE. Briefly, the class of basis functions, known as ‘universal approximators’, allows an agent to use histogram data from the operator T to estimate the posterior probability measure π_{it}^T , for a given prior probability measure π_{it} , by fitting a parametric family of curves to the histogram.

By varying the prior probability measure used to compute bids and asks for the temporary equilibrium (which is also expressed using the basis functions), an entire map of (π_{it}, π_{it}^T) pairs can be generated over the retrading sample $t = 1 \dots \tau$. Using the same class of universal approximators, it is possible to find the mapping which turns the prior probability vector into the temporary equilibrium, posterior probability vector. This is a representation of the functional operator T (call it \hat{T}).

Lastly, an agent can use a numerical fixed point algorithm of choice to compute the fixed point (as a vector) of the representation \hat{T} . Using this vector fixed point, it is then possible to construct the functional fixed point π^* .

Notice that the agent's optimal decision-making is part of what builds the functional fixed point to be estimated—in particular, the bids and asks must be rationally generated from Bayes' Rule. But the strategy used may be generated from other optimizing criteria, or from strategic concerns (e.g. a Bayesian-Nash equilibrium; see Ghosal and Morelli [2004]).

In addition, it is important to note that along the equilibrium path, while the functional data set is generated, there is no point where trade ceases to exist. This is because the agent needs to sample 'enough' points in the function space to get a good fit for the functional mapping, and this will require submitting a (usually random) distribution of prior probability measures. So the prior probabilities may not even be rationally chosen—only the bids and asks, and the updating of the prior probability after the temporary equilibrium price is revealed, need be rationally generated. This is the content of step 4 above.

Of course, sampling this probability space is both computationally expensive and also expensive in terms of forgone gains to trade—the prior probability measures are not selected with any optimality criteria until the functional fixed point of T is found. Therefore, it may make sense for the agent to exploit the properties of the mapping T which may speed convergence—if, for example, the mapping has properties which are consistent with a contraction mapping, then it may make sense for the previously obtained posterior distribution π_{it}^T to become next period's prior π_{it+1} . This argument is also strengthened due to the way Bayes' Rule functions—the 'narrowing in on the truth' toward the true dividend distribution $\pi(d)$ is facilitated by adopting the previous period's posterior distribution as the following period's prior, and so convergence to the true dividend distribution (should it be an REE!) may obtain under learning.

6 Conclusion

For market games with a finite number of players and a finite number of uncertain states, there is in general a rational expectations equilibrium (REE) which is defined as a probability measure over the states. Moreover, this REE may be found using a numerical algorithm, which has the advantage of being stated in terms of computable components (like Fourier coefficients). This allows agents to harness the enormous computing power of the early 21st century to maximum advantage. Although it is an open question as to whether or not an REE is a desirable equilibrium concept, the notion that ‘expectations are fulfilled using all available information’ remains a compelling argument. In this environment, far from obtaining a no-trade REE, ‘most’ economies have a lively exchange of goods and assets as agents seek to both do the right thing (via optimization) and learn as much as possible on the way toward finding the REE, at which point learning ceases.

It would be very interesting to implement a market game in a computational environment, to demonstrate the algorithm in a similar way as in Kelly and Shorish [2000], and this is one topic of future research. For a continuum of agents, Shubik and Vriend [1999] have shown that with piecewise linear utility functions a simulated economy using genetic algorithms and classifier systems can numerically converge to stationary NE of market games. Designing a simulation with many agents estimating a functional REE would expand the applicability of simulation techniques to this finite agent environment.

Lastly, it is an open question as to whether or not it is ‘worth it’ to know the truth, i.e. to find an REE for any market where trading takes place over many periods. The success of ‘rule of thumb’ strategies implies that, on some level, higher welfare gains are achieved by using strategies which are not too complex. It may be that the REE computing algorithm designed here carries a higher welfare cost than an alternative algorithm which does not seek to find any equilibrium (e.g., a completely myopic, one-period-ahead strategy). Finding a good measure of the welfare costs of algorithms, and the associated efficiency of the REE itself, is another further direction of research.

References

- BORDER, K. (1985): *Fixed Point Theorems with Applications to Economics and Game Theory*. Cambridge University Press.
- DUBEY, P. (1980): “Nash Equilibria of Market Games: Finiteness and Inefficiency,” *Journal of Economic Theory*, 22(2), 363–376.
- DUBEY, P., J. GEANAKOPOLOS, AND M. SHUBIK (1987): “Revelation of Information in Strategic Market Games: A Critique of Rational Expectations,” *Journal of Mathematical Economics*, 16, 105–138.
- DUBEY, P., AND J. ROGOWSKI (1990): “Inefficiency of Smooth Market Mechanisms,” *Journal of Mathematical Economics*, 19, 285–304.
- DUBEY, P., S. SAHI, AND M. SHUBIK (1993): “Repeated trade and the Velocity of Money,” *Journal of Mathematical Economics*, 22, 125–137.
- GHOSAL, S., AND M. MORELLI (2004): “Retrading in Market Games,” *Journal of Economic Theory*, 114, 255–279.
- GROSSMAN, S. J., AND J. E. STIGLITZ (1980): “On the Impossibility of Informationally Efficient Markets,” *American Economic Review*, 70(3), 393–408.
- JACKSON, M. O., AND J. PECK (1999): “Asymmetric Information in a Competitive Market Game: Reexamining the Implications of Rational Expectations,” *Economic Theory*, 13(3), 603–628.
- KELLY, D. L., AND J. SHORISH (2000): “Stability of Functional Rational Expectations Equilibria,” *Journal of Economic Theory*, 95(2), 215–250.
- SHUBIK, M., AND N. VRIEND (1999): “A Behavioral Approach to a Strategic Market Game,” in *Computational Techniques for Modelling Learning in Economics*, ed. by T. Brenner. Kluwer Press.

Author: Jamsheed Shorish

Title: Functional Rational Expectations Equilibria in Market Games

Reihe Ökonomie / Economics Series 186

Editor: Robert M. Kunst (Econometrics)

Associate Editors: Walter Fisher (Macroeconomics), Klaus Ritzberger (Microeconomics)

ISSN: 1605-7996

© 2006 by the Department of Economics and Finance, Institute for Advanced Studies (IHS),
Stumpergasse 56, A-1060 Vienna • ☎ +43 1 59991-0 • Fax +43 1 59991-555 • <http://www.ihs.ac.at>
