

MERGING ECONOMIC FORECASTS⁺

by

Ehud Kalai*

and

Ehud Lehrer*

October 1990

Revised and Expanded February 1994

⁺The authors wish to thank, without implicating them, Sergiu Hart, John Hillas, Larry Jones, Jim Jordan, Morton Kamien, Andreu Mas-Colell, Jean-Francois Mertens, Stanley Reiter, Nancy Stokey, and Chuck Wilson. This research was partly supported by Grant No. SES-9022305 and SES-9011790 from the National Science Foundation, Economics.

*Department of Managerial Economics and Decision Sciences, J. L. Kellogg Graduate School of Management, Northwestern University, 2001 Sheridan Road, Evanston, Illinois 60208.

MERGING ECONOMIC FORECASTS

Abstract

Under the rational expectations assumption of Muth, economic agents use their perfect knowledge of the distribution of future prices to compute optimal current actions. In private forecasts equilibrium, introduced here, agents use subjective inaccurate forecasts about future prices to compute optimal current actions. The paper discusses truth-compatibility conditions guaranteeing convergence of the subjective forecasts to the correct one. Thus, under such conditions, every private forecast equilibrium must converge with time to a Muth-type rational expectations equilibrium.

1. Introduction

Under the rational expectations assumption of Muth (1961), economic agents anticipate correctly the distribution of future prices and use it to compute optimal current actions. It is obvious that future prices should be considered, but the assumption that agents know the correct distribution of future prices is questionable.

The purpose of this note is to exhibit conditions that lead the agents, after sufficiently long economic interaction, to an essentially correct forecast of future prices. In the model of a dynamic economy presented here, each individual agent starts with a truth-compatible private subjective forecast of future prices and uses it in choosing actions that maximize the present value of his long run expected utility. The two conditions, of truth-compatible initial forecasts and expected utility maximization, turn out to guarantee eventual learning of the correct forecast, and thus convergence to a Muth-type rational expectations equilibrium of the dynamic economy.

The more recent literature on rational expectations, as, for example, in Radner (1967), Lucas (1972) and later papers, have dealt with more sophisticated models of rational expectations. We refer the reader to the survey of Jordan and Radner (1982) for more elaborate discussions and further references. Much of this literature is concerned with the optimal use of private signals available to individual agents trying to anticipate the parameters of future economies.

Overall, learning papers in this literature can be roughly classified as having "irrational" or "rational" models. We refer the reader to Bray and Kreps (1987) and Blume and Easley (1992) for an elaborate discussion of these issues and references. The irrationality relates to the agents' assumption that what they attempt to learn is stationary. This assumption is violated when the process of learning is effective, leading the agents to change their behavior, and as a result violating the stationarity assumption. One needs to resort to additional restrictive assumptions or assume bounded rationality

to justify such models.

In rational models, on the other hand, agents are aware of the ongoing process of learning and its effects on changing behavior. Each agent is assumed to have a detailed and correct knowledge of the information structure of the other agents, their rules of learning, and correct prior probability distributions over unknown fundamental parameters. The assumption that the process of learning chosen by each such agent is optimal, given his information, yields a description of rational interactive learning behavior. One may think of this approach as imbedding the process of learning an equilibrium in a larger and more complex equilibrium model of learning under incomplete information. Thus, as pointed out by Bray and Kreps, the question of whether agents learn to follow an equilibrium is not really solved.

The current paper develops an open model which can be used to describe both rational and irrational agents. This means that agents are not assumed to model detailed components of the economy, other agents, other agents' information systems, and so forth. Nor are they assumed to assign prior probability distributions to such items. Instead, the agents' primitives of uncertainty are modeled directly on their payoff-relevant items, namely, future prices. Thus, we assume that each agent possesses a private forecast, i.e., a subjective distribution on future price paths, which guides him in determining an optimal dynamic economic plan. Irrationality of various degrees can come in since the private forecasts are left largely unspecified.

One immediate benefit of using an open--rather than closed--model is that it partly overcomes the Bray and Kreps criticism mentioned above. Namely, the economic system is not assumed to start at equilibrium in order to obtain convergence to equilibrium.

A second benefit of the open model is that it allows players to be less super-rational. In economies of even small size, the assumption that players model,

deterministically or stochastically, the entire working of the economy, including the functioning of other agents, seems highly unrealistic. It is important to note, however, that the open model does not preclude such super-rational agents, since a closed model can be thought of as a special case of the open one. This is the case since an equilibrium of a closed model induces individual forecasts to which the agents' actions are optimal. By taking these individually induced forecasts as the starting private forecasts in the open model, convergence to the correct forecast in the open model implies convergence also in the closed model. However, the convergence to the correct forecast demonstrated in this paper is obtained by considering beliefs placed directly on price forecasts, without the need to know where they came from.

It turns out that the proposed model cannot be completely open. As illustrated in the body of the paper, without some condition of compatibility between the subjective forecasts and the true forecast, convergence may not occur. Therefore, truth compatibility conditions, and resulting notions of convergence, become a major topic for this analysis. In an earlier paper (Kalai and Lehrer (1993)), the authors showed that absolute continuity of the truth relative to subjective beliefs guarantees convergence to a Nash equilibrium of a perfect-monitoring infinitely-repeated game. The first result described in the current paper is a rational expectations analog. Absolute continuity here means that every positive probability event, in the space of infinite future price paths, is not assigned probability zero according to an individual private forecast. This is a strong truth-compatibility condition since it also includes events in the infinite future, i.e., tail events. Not surprisingly, it yields strong convergence to rational expectations equilibrium in the sense that agents' forecasts of events, including tail events, will become accurate.

As pointed in Kalai and Lehrer (1994), for economic agents who discount future payoffs, assessing correctly the probability of tail events is unimportant. Assessing

correctly probabilities of finite horizon events is sufficient. Using more recent results of Lehrer and Smorodinsky (1993), we illustrate truth-compatibility conditions, weaker than absolute continuity, that yield convergence of this type. Under these conditions, agents may assign probability zero to positive probability events, but they never assign probability zero to a "neighborhoods of the truth." This means that if agents' forecasts are sufficiently diffused, and assign positive probability to every open set in a properly defined sense, then no matter what the real forecast is, convergence is assured. Such convergence guarantees that agents' forecasts are accurate for all finite horizon events at almost all sufficiently late periods. Under an additional condition, the forecasts become accurate at all sufficiently late periods.

Section 2 of the paper contains an example of a familiar storable commodities economy. The reader who is not interested in the details of the general model, truth-compatibility, and convergence notions, can obtain the general message of the paper by reading this example alone. However, the general results are applicable to a significantly broader set of economies.

Sections 3-5 contain the general model of the dynamic economy, the notions of rational expectations and private forecasts equilibria. It starts with a straightforward description of an infinite horizon dynamic economy and a Muth-type rational expectations equilibrium. What may be new there is that such an equilibrium is allowed to be stochastic, which means that the determination of each period's prices could be random (provided that it leads to optimal individual actions that clear the period's markets). What is more likely to be new in these sections is the concept of private forecasts equilibrium. At such an equilibrium, prices are determined in each period, so that agents' choice of subjectively optimal period's actions, relative to their updated private forecasts, clear the periods' markets.

Section 6 is devoted to expanded discussion of conditions of truth-compatibility

and merging of measures, and Section 7 is devoted to the statements of equilibrium convergence.

In Section 8 we return to a discussion of the applications to the storable economy example of Section 2.

2. An Example With Storable Commodities

We first illustrate the main message and concepts through a familiar simple example of an economy with two infinitely-divisible commodities, L and R. These commodities are perfect complements and can only be consumed together in equal amounts. All agents in the economy have identical linear function $u_i(c,c) = c$ describing their von Neumann-Morgenstern utility when consuming the vector (c,c) . The commodities are perfectly storable. Thus, any quantities owned and not consumed by an agent will be available to him, at zero storage cost, at the beginning of the next period.

The n agents in the economy can be classified into two types with different endowment mechanisms. Type 1 agents, labeled $i = 1, 2, \dots, m_1$, are each endowed at the beginning of each period $t = 1, 2, \dots$ with one unit of commodity L, i.e., $(1, 0)$. Type 2 agents, labeled $i = m_1 + 1, \dots, m_1 + m_2$, are each endowed at the beginning of each period one unit of R, i.e., $(0, 1)$.

It is useful to distinguish the case of $m_1 = m_2$, i.e., a balanced economy, from the case of inequality, say, $m_1 < m_2$, i.e., imbalanced economy. Due to the fact that L and R have to be consumed in equal quantities, the balanced economy must have equal quantities of L and R at the beginning of each period, while the imbalanced economy must have extra units of R at the beginning of each period.

In both economies the agents are told at the beginning of each period the period's prices of L and R and then they decide how much to trade, how much to

consume, and how much to store for the next period. Overall utility from an infinite consumption stream is evaluated according to the utility stream generated, discounted according to a common discount parameter.

One rational expectations equilibrium of the balanced economy has the auctioneer choosing the price path $(.5,.5), (.5,.5), \dots$. At the beginning of the first period, when an agent is told that the prices of L and R are the same, and knowing that under the correct forecast of future prices they will continue to be the same, his best response is to trade one-half of his endowed unit with one-half of the opposite unit to consume the vector $(.5,.5)$ immediately. This logic applies to all agents in all periods and as a result, one obtains even trades with markets clearing and full consumption (no storage) in every stage of the economy.

However, more complex rational expectations equilibria involving random choices of prices exist. Consider, for example, the balanced economy. Let μ denote a probability distribution on future price paths of L, p_L^1, p_L^2, \dots (with $p_R^t = 1 - p_L^t$). Periods' prices of L will be drawn according to μ , and constructing a rational expectations equilibrium, we assume that μ is known to the agents. We restrict our attention to distributions μ with the property that tomorrow's expected price of L equals today's realized price, i.e., $E_\mu(P_L^{t+1} | P_L^t) = P_L^t$. With such a property, when an agent is informed of today's realized price, he knows that storing any quantity cannot increase his expected utility from consumption tomorrow beyond what he can already obtain from it today. Hence, because of the discounting, he is better off trading optimally and consuming all he can at the current period. So if an auctioneer chooses prices randomly according to a distribution μ with the above property, all the agents in each period will trade at the period's realized price to have equal amounts of the two commodities, which they will consume immediately. The markets will clear and thus, any such distribution is a stochastic rational expectations equilibrium.

For the rational expectations equilibria of the imbalanced economy we obtain only the familiar extreme outcome. As we will show later, at any such equilibrium, deterministic or probabilistic, the auctioneer must assign probability one to the price path $1,1,\dots$ of the short commodity L (and $0,0,\dots$ to R). Every agent 1 through m_1 will be receiving, for free, a unit of R and consuming $(1,1)$ in each period. The other agents will hand over the needed units of R and will be consuming nothing.

We are interested, however, in an equilibrium which allows the agents to have subjective private forecasts, i.e., ones that do not necessarily coincide with the correct forecast, nor with each other's. At such an equilibrium each agent has a private forecast, $\tilde{\mu}_i$, about future prices which he uses and updates for computing his optimal trade-consumption-storage plan. The auctioneer still draws prices according to a distribution μ . But for the whole system, consisting of μ , the $\tilde{\mu}_i$'s and the agents' plans, to be at a private forecasts equilibrium it is required that: (i) the individual plans are optimal relative to the private forecasts, and (ii) that μ assigns positive probability only to prices that clear the markets when the agents use their subjectively optimal plans.

Clearly every rational expectations equilibrium can be viewed as a private forecasts equilibrium when the private forecasts coincide with the correct one used by the auctioneer, i.e., $\tilde{\mu}_i = \mu$ for all i . But private forecasts equilibrium allows for more. The agents may be unaware of the distribution of other agents, their preferences, and their initial endowments. For example, it is possible that agents 1 through m_1 think that the prices of L are likely to go up beyond the period's announced ones and as a result store their endowments for some time. For similar subjective reasons, agents $m_1 + 1$ through $m_1 + m_2$ may hold on to their endowment for some time. Thus, we may see periods with different prices where the markets clear since no trade is desired by any agent. As the above reasoning suggests, private forecasts equilibrium allows the description of speculative behavior.

The next two examples of private-forecasts equilibrium of a balanced economy distinguish the two situations studied in this paper. In the first example, subjective beliefs are not compatible with the truth and convergence does not occur. While, in the second example, convergence must occur due to such compatibility.

Example 2.1: All agents hold the following (identical) private forecast $\bar{\eta}$. Independently of past prices, $\bar{\eta}$ assigns probabilities of $1/3$ each to next period's prices being $(1,0)$, $(.5,.5)$ or $(0,1)$. Unlike this assessment, however, the real prices in each period are $(.5,.5)$ with certainty. If the agents are sufficiently patient, then whatever type of good they are endowed with, they will store it until they "get lucky" and have a price of 1 for their good. As a result, nobody would ever trade and the markets will clear in every period. Thus, the above is a private forecasts equilibrium.

At the above equilibrium, the private forecasts never converge to the correct one, no matter how long a history of constant equal prices was observed. This is due to the violation of the truth compatibility conditions discussed in the body of this paper. For example, the absolute continuity condition is violated since the event F , that prices will continue being $(.5,.5)$ forever, has real probability 1 but is assigned probability 0 by the individual agent's private forecasts described by $\bar{\eta}$.

It is true that such a violation will not become apparent to an agent in any finite time. If F^t is the event that up to time t all prices are $(.5,.5)$, then F^t has a real positive private probability and no clear-cut violation of subjective beliefs is observed. However, the agent should become unboundedly suspicious. F is the intersection of all the F^t 's and as the above private forecasts equilibrium unfolds, the agent sees the cumulative buildup of an infinite pattern of prices, $(.5,.5)$, that he considers impossible.

Example 2.2: All agents, of both types, hold an individual private forecast $\bar{\mu}$ determined

from a prior probability distribution on two forecasts described as follows. They assign probability $1 - \epsilon$ to the past-independent forecast $\bar{\eta}$ of the previous example, but also a probability $\epsilon > 0$ to the constant forecast of prices being $(.5,.5)$ forever. The real forecast, μ , assigns constant price path $(.5,.5)$ probability one.

As in the previous example, the agents' initial optimal behavior is to store their endowment with the hope of becoming lucky and trading their good at the price of 1. Nobody wishes to trade at $(.5,.5)$ during these initial periods and the markets clear. However, after sufficiently long history of $(.5,.5)$ prices, their posterior probability assigned to the random forecast, $\bar{\eta}$, will become sufficiently small, and from this time on they will trade at the prices $(.5,.5)$. Since we are in a balanced economy, each period's market from now on will clear and, in effect, they have converged to a rational expectations equilibrium.

The main message of this paper is that under compatibility assumptions of the beliefs with the truth, the individual forecasts of a private forecast equilibrium must eventually merge, and thus convergence to a rational expectations equilibrium is assured. So speculative behavior, as in Example 2.1, cannot persist. As an application for imbalanced economics, we will argue in the last section of the paper that at every private forecast equilibrium the price of the short commodity must eventually approach 1.

3. The Dynamic Economy

We let $N = \{1,2,\dots,n\}$ be a set of agents, each with a nonempty finite, or countable set, A_i , describing the conceivable actions that he may take in the intermediary stages of a discrete time infinite-horizon dynamic economy. A von Neumann-Morgenstern utility function describes the immediate payoff that agent i

receives, $u_i(a_i)$, when he takes an action $a_i \in A_i$. A common discount parameter, $0 < \lambda < 1$, will be used by all agents to evaluate utility payoff streams. The set $A = \times_i A_i$ will denote all conceivable action vectors.

In our storable goods example an action of player i is described by a vector $a_i = (b_1, b_2, c, s_1, s_2)$ describing respectively the quantities of L and R he buys, the joint quantity of L and R he consumes, and the quantities of L and R he stores. Actions $a_i \in A_i$ are restricted to consist of rational numbers with $c, s_1, s_2 \geq 0$. $u_i(a_i) = u_i(c, c) = c$.

We let P denote a nonempty countable set of (abstract) price vectors, or simply prices, that may emerge in the intermediary stages of the dynamic economy. Each agent i has a countable set of states, S_i , and a correspondence $F_i(s_i, p)$, describing the nonempty set of individually feasible actions he can take when he is at state s_i and the prices are p . A deterministic transition function describes a new state, $T_i(s_i, a_i)$, resulting from s_i when the action a_i is taken.

In the storable goods economy P consists of price vectors $p = (p_1, p_2)$ with p_1 and p_2 being rational nonnegative numbers summing to 1. An individual state there consists of a stage initial endowment of L and R vector $w_i = (w_1, w_2)$, consisting of nonnegative rational numbers. $F_i(w_i, p)$ consists of all the actions a_i satisfying the budget constraint, $bp \leq 0$, and the preservation of commodities, $w_i + b = s + (c, c)$. The transition rule there is $T_i(w_i, a_i) = (1, 0) + s$ for $i = 1, \dots, m_1$, and $T_i(w_i, a_i) = (0, 1) + s$ for $i = m_1 + 1, \dots, m_1 + m_2$.

A group feasible correspondence describes sets of action vectors, $G(s, p)$, that the group can take at every vector of states $s \in S = \times_i S_i$, when the prices are p . It is assumed that group feasible action-vectors exist and they must always consist of individually feasible actions, i.e., $\emptyset \neq G(s, p) \subseteq \times_i F_i(s_i, p)$ for every $s \in S$ and $p \in P$.

This is a market clearing condition and in the storable goods example $a = (a_1, \dots, a_n) \in G(w, p)$ if each $a_i \in F_i(w_i, p)$, and $\sum_{i \in N} b_i = 0$.

A dynamic economy in the environment described above is specified by a vector of initial states $s^1 = (s_1^1, \dots, s_n^1)$. At the first stage, prices $p^1 \in P$ will be announced, then the agents will choose a group feasible action vector $a^1 \in G(s^1, p^1)$ yielding themselves payoffs $u_i(a_i^1)$. For the second stage, the economy will transit to a new vector of states s^2 defined by $s_i^2 = T_i(S_i^1, a_i^1)$ in which similar choices of prices and feasible actions are made and payoffs result. Evolving in this fashion the dynamic economy generates infinite payoff streams which are evaluated by each agent using the common discount parameter λ .

In the storable goods example, the initial states are $w_i^1 = (1,0)$, for $i = 1, \dots, m_1$, and $w_i^1 = (0,1)$ for $i = m_1 + 1, \dots, m_1 + m_2$.

4. Equilibrium Concepts

The word path will denote a countable infinite sequence. For example, a price path is a sequence of price vectors p^1, p^2, \dots . We will use P^∞ to denote the set of price paths, A_i^∞ to denote the set of individual action paths, etc. A history will mean a finite initial segment of a path. For example, P^t , the t -times cartesian product of P with itself, denotes all price histories of length t . $\cup_{t \geq 1} P^t$ denotes the set of all price histories, etc.

A forecast is a probability distribution on the set of price paths, P^∞ . As usual, the σ -algebra used is the one generated by cylinder sets. A cylinder set here means a set of all price paths whose prefix coincides with a given price history. So, for a forecast μ and a price history $\bar{p}^t \in P^t$, we can think of \bar{p}^t as the event consisting of all price paths with initial segment \bar{p}^t and $\mu(\bar{p}^t)$ as the probability of this event. Following standard results in probability theory, in order to specify a forecast it suffices to specify the probabilities of all finite price histories.

An individual plan is a function $\pi_i: \cup_{t \geq 1} P^t \rightarrow A_i$ choosing an individual action

after every price history in an intertemporally feasible way. Formally, for every price path p^1, p^2, \dots , we define inductively the resulting history of individual state-action pairs $(s_i^1, a_i^1)(s_i^2, a_i^2) \dots$ starting with the initial state s_i^1 and the initial action $a_i^1 = \pi_i(p^1)$, and continuing inductively with $s_i^{t+1} = T_i(s_i^t, a_i^t)$ and $a_i^{t+1} = \pi_i(p^1, \dots, p^t)$. Individual intertemporal feasibility requires that $a_i^{t+1} \in F_i(s_i^{t+1}, p^{t+1})$.

For every forecast σ_i and feasible plan π_i we can compute the expected discounted payoff $u_i(\pi_i, \sigma_i)$ in the obvious way. We say that π_i is optimal relative to the forecast σ_i if for any feasible plan $\bar{\pi}_i$ $u_i(\bar{\pi}_i, \sigma_i) - u_i(\pi_i, \sigma_i) \leq 0$.

A private forecasts equilibrium is a triple $(\mu, \tilde{\mu}, \pi)$ consisting of a forecast μ , a vector of private forecasts $\tilde{\mu} = (\tilde{\mu}_1, \dots, \tilde{\mu}_n)$ and a vector of plans $\pi = (\pi_1, \dots, \pi_n)$ satisfying the following conditions. Each individual plan π_i is optimal relative to the individual forecast $\tilde{\mu}_i$ and with μ -probability one the individual plans are intertemporally group feasible, e.g., clear the periods markets. This means that for every μ -positive probability price history p^1, \dots, p^t, p^{t+1} and the resulting vector of individual states $s^{t+1} = (s_1^{t+1}, \dots, s_n^{t+1})$, the $t+1$ -st period vector of actions $\pi(p^1, \dots, p^t) \in G(s^{t+1}, p^{t+1})$.

The idea motivating the above definition is that, just as in a rational expectations equilibrium, the distribution by which period prices are determined, μ , is chosen to guarantee that markets clear. However, the individual agents choose their period's actions under the assumption that future prices will unfold according to their subjective forecasts, $\tilde{\mu}_i$, rather than μ_i , which they do not know.

For convenience of presentation, we will assume from now on that every μ -positive probability price history is assigned a positive probability by each $\tilde{\mu}_i$. This milder version of truth-compatibility will be implied by all the truth-compatibility conditions to be discussed later. It eliminates the need to deal with Bayesian updating after observing probability zero events.

A rational expectations equilibrium is defined formally to be a pair, (μ, π) ,

consisting of a forecast and a vector of individual plans for which $(\mu, (\mu, \mu, \dots, \mu), \pi)$ is a private forecasts equilibrium. In other words, it is the special case where the private forecasts happen to be accurate.

5. Induced Equilibria

Starting with a private forecasts equilibrium $e = (\mu, \tilde{\mu}, \pi)$ and a μ -positive-probability price history $\bar{p}^t \in P^t$ let $s (=s(\pi, \bar{p}^t))$ be the vector of states induced for stage $t + 1$ of the economy. We let $\mathcal{E}_{\bar{p}^t}$ denote the economy with this vector of initial states. Each forecast, the true one or the private subjective ones, of the original economy induces a forecast in the induced economy $\mathcal{E}_{\bar{p}^t}$. For example, the true forecast induces the forecast $\mu|_{\bar{p}^t}$ defined by $\mu|_{\bar{p}^t}(\bar{q}^r) = \mu(\bar{p}^t \bar{q}^r | \bar{p}^t)$. (By $\bar{p}^t \bar{q}^r$ we mean the price history of length $t + r$ which starts with the elements of \bar{p}^t and continues with the elements of \bar{q}^r .) We let $\tilde{\mu}|_{\bar{p}^t}$ be the vector of induced private forecasts.

The induced private plans are defined in the obvious way, $\pi_i|_{\bar{p}^t}(\bar{q}^r) = \pi_i(\bar{p}^t \bar{q}^r)$. It is easy to check that the private forecasts equilibrium e induces a private forecasts equilibrium $e|_{\bar{p}^t}$ in the economy $\mathcal{E}_{\bar{p}^t}$ defined by $(\mu|_{\bar{p}^t}, \tilde{\mu}|_{\bar{p}^t}, \pi|_{\bar{p}^t})$.

Our main results are that private forecasts equilibria induced after long price histories must converge to rational expectations equilibria, provided that the initial individual forecasts are compatible with the true forecast. The next section is therefore devoted to notions of compatibility-ensuring merging of forecasts.

6. Truth-Compatibility Leading to Merging of Probabilities

We let, as in the previous sections, $p = (p^1, p^2, \dots)$ denote a vector of future prices which we view as a sequence of random variables. Also, continuing with the previous notations, we let μ and $\tilde{\mu}_i$ be probability distributions on vectors p with the interpretation that μ is the "real" distribution and $\tilde{\mu}_i$ is a subjective one held by an

individual agent. With this in mind, all unspecified probabilities, in the statements to follow, refer to the true probability distribution. So, for example, if we say that $\tilde{\mu}$ merges to μ_i with probability one, or along almost all price paths, we refer to probability one according to μ .

Merging refers to the phenomenon that, after observing a long vector of past prices, $\bar{p}^t = (p^1, \dots, p^t)$, the subjective conditional probability on future prices, $\tilde{\mu}_i |_{\bar{p}^t}$, becomes close to the true one, $\mu |_{\bar{p}^t}$. However, if strong incompatibility exists between $\tilde{\mu}_i$ and μ , merging may not occur. For example, if μ assigns probability one on the single price path 1,1,1,... and $\tilde{\mu}_i$ places probability of $(1/2)^t$ on any length t path consisting of 0's and 1's, then, no matter how many 1's were observed, $\mu |_{\bar{p}^t}$ and $\tilde{\mu}_i |_{\bar{p}^t}$ will persistently stay the same (as they were originally) and no merging will occur.

We next describe conditions that rule out such serious inconsistencies.

6a. Absolute Continuity

We say that μ is absolutely continuous with respect to $\tilde{\mu}_i$, $\mu \ll \tilde{\mu}_i$, if $\tilde{\mu}_i(A) > 0$ for every event A with $\mu(A) > 0$. Absolute continuity means that any event that is possible is not ruled out by the agent. It implies, as we have assumed earlier, that the agent will never be "surprised." As prices of the economy unfold, he will never see a history of prices that he considered impossible. Technically, this condition will enable him to compute conditional probability distributions on future prices, since he would never have to condition on what he considers to be a zero probability event.

But absolute continuity is stronger than a "not be surprised" condition since the events A allowed in its definition can be in the infinite σ -algebra. Such events may never be determined, confirmed or contradicted, in any finite time. Thus, a more accurate verbal description is that the agent cannot become unboundedly suspicious, as was described in Example 2.1.

One instance where absolute continuity holds is when there are at most countably many possible forecasts, including the true one, and the subjective forecast is obtained by assigning a prior positive probability to each one of them. This yields a subjective forecast with a stronger property than absolute continuity, described as follows.

We say that $\tilde{\mu}_i$ has a grain of truth if $\tilde{\mu}_i$ can be expressed as $\alpha\mu + (1 - \alpha)\hat{\mu}$, where $0 < \alpha \leq 1$ and $\hat{\mu}$ is a probability measure. Thus, $\tilde{\mu}_i$ is a convex combination of truth, μ , and some residual probability distribution, $\hat{\mu}$.

6b. Diffused Beliefs

The following example captures the main idea behind this notion. Suppose that a coin with parameter $1/2$ is repeatedly and independently tossed to generate an infinite sequence of Heads and Tails. Let μ describe the distribution induced by this fair coin. An agent, not knowing the parameter of the coin, forms a subjective distribution $\tilde{\mu}_i$ using a uniform prior distribution on the interval of possible parameter values $(0,1)$. It is readily seen that the true distribution over infinite strings of Heads or Tails is not absolutely continuous with respect to the belief. For example, $\tilde{\mu}_i$ assigns probability zero to the event that the asymptotic frequency of Heads is precisely $1/2$, while the true probability is 1.

However, as tosses of the coin are observed, the posteriors over the possible parameters become concentrated around $1/2$. The point $1/2$, though, always attains zero probability and, therefore, the lack of absolute continuity remains in effect all the time.

Despite the lack of absolute continuity, after a sufficiently long time one can predict with high precision the distribution over the subsequent set of possible outcomes; it is approximately $1/2$ over Heads and Tails. Bayesian updating enables the observer to learn the correct distribution over short-run outcomes. But there is no way to learn to predict infinite horizon events, as in the case of absolute continuity. This,

however, should represent no drawback for agents who discount future utility.

What makes the learning possible in the coin example is that the belief is diffused around the right coin (with parameter 1/2); any neighborhood of 1/2 is assigned positive probability. There are two ways to formulate this idea. The first, which is more transparent, is a generalization of the coin setup to nonstationary distributions. The second is more general. For both we need the following definition in which, for any path $p = (p^1, p^2, \dots)$ of prices, we let \bar{p}^t be the t -prefix of p .

Definition 6.1: We say that distribution μ' is asymptotically ϵ -near μ if with μ -probability 1 there is a time T s.t. $t \geq T$ implies

$$|\mu(p^{t+1}|\bar{p}^t)/\mu'(p^{t+1}|\bar{p}^t) - 1| < \epsilon$$

(with $0/0$ defined to be 1).

Thus, if μ' is asymptotically ϵ -near μ , then from time T on the likelihood ratio of the two probabilities of the next upcoming short-run event, p^{t+1} , given \bar{p}^t is close to 1 up to an ϵ .

Suppose that there is a set of possible economies indexed by θ . The agent holds a prior probability distribution F_i over the set of indices Θ . Naturally, each economy \mathcal{E}_θ , $\theta \in \Theta$, induces a distribution μ_θ over paths of market prices. Moreover, the belief F_i defines an aggregate distribution denoted by $\tilde{\mu}_i$ over the same space. In other words, agent i assesses the evolution of market prices by the distribution $\tilde{\mu}_i$, which pertains to the average economy, as agent i perceives it.

In fact, the real economy, \mathcal{E} , is just one of the candidates. In a case where \mathcal{E} is assigned positive probability by F_i , the grain of truth condition is satisfied and we may resort to the previous section. The interesting case, though, is as in the coins example,

when \mathcal{E} is assigned zero probability. The chance of \mathcal{E} being the real economy is, according to agent i 's belief, zero. However, if neighborhoods of \mathcal{E} are ascribed positive probability, then learning will take place. To make it rigorous, define the ϵ -neighborhood of \mathcal{E} , C_ϵ , as the set of all indices θ such that μ_θ is asymptotically ϵ -near μ , where μ is as usual the real distribution (the one induced by \mathcal{E}). Formally,

$$C_\epsilon = \{\theta \in \Theta; \mu_\theta \text{ is asymptotically } \epsilon\text{-near } \mu\}.$$

We assume that C_ϵ is a measurable set in Θ .

Definition 6.2: We say that $\tilde{\mu}_i$ is diffused around μ if for every $\epsilon > 0$, $F(C_\epsilon) > 0$.

In words, $\tilde{\mu}_i$ is diffused around the real distribution if every neighborhood of μ (the one induced by the real economy) is assigned by F_i a positive prior probability. In the coins example, F_i is the uniform distribution over the set of all possible parameters, $\Theta = (0,1)$. Therefore, μ_θ is the distribution induced by an i.i.d. random variable that takes the value "Heads" with probability θ and the value "Tails" with probability $1 - \theta$. Here, $C_\epsilon = [1/2 - \epsilon, 1/2 + \epsilon]$ and it has positive weight according to F_i . Therefore, $\tilde{\mu}_i$ is diffused around $\mu = \mu_{1/2}$

There is a second way to phrase the diffusion relationship between μ and $\tilde{\mu}_i$ without having to pass through a parametric set-up.

Definition 6.3: $\tilde{\mu}_i$ is diffused around μ if for every $\epsilon > 0$ and for almost every path p there is a measure μ_ϵ s.t.

- (1) $\tilde{\mu}_i = \alpha\mu_\epsilon + (1 - \alpha)\hat{\mu}$, where $0 < \alpha \leq 1$ and $\hat{\mu}$ is a probability distribution.
- (2) There is a time T s.t. $t \geq T$ implies $|\mu(p^{t+1}|\bar{p}^t)/\mu_\epsilon(p^{t+1}|\bar{p}^t) - 1| < \epsilon$.

It turns out that the learning implied by the diffusion assumption alone is weaker than the one achieved in the coins example. There, coins in C_ϵ induce probabilities close to the real coin, but in addition, coins not in C_ϵ induce probabilities different from the real coin. Therefore, by observing enough outcomes, an agent can separate between close and remote distributions. We wish to adopt the same notion of separability for economies. That is, after observing a long history of market prices the agent should be able to rule out economies that are far from the correct one.

What needs to be added is a third condition to the former definition.

Definition 6.4: Suppose that $\tilde{\mu}_i$ is diffused around μ . We say that $\tilde{\mu}_i$ has the separation property if for every C_ϵ and $\theta \notin C_\epsilon$ with μ probability 1 there is a set of periods L having a positive lower density¹ and a number $d > 0$ s.t. for every $t \in L$ there exists a price p^{t+1} satisfying

$$|\mu(p^{t+1} | \bar{p}^t) - \mu_\theta(p^{t+1} | \bar{p}^t)| > d.$$

Notice that the set L can be a tail of periods, $\{T, T+1, \dots\}$. In this case, μ_θ and μ are different in assessing p_{t+1} given \bar{p}^t (one step ahead) by at least d . Notice also that d can be less than ϵ .

Back to the coin example: if $\theta \notin [1/2 - \epsilon, 1/2 + \epsilon]$, then, denoting by \bar{p}^t a history of t outcomes, $|\mu(\text{Heads} | \bar{p}^t) - \mu_\theta(\text{Heads} | \bar{p}^t)| > \epsilon$ for every t . Therefore, $\tilde{\mu}_i$ has the separation property.

¹The lower density of a set L is defined as $\liminf (1/n)|L \cap \{1, \dots, n\}|$.

7. Merging of Forecasts and Congvergence to Equilibrium

At a private forecasts equilibrium $e = (\mu, \tilde{\mu}, \pi)$, each agent's plan, π_i , is optimal relative to his private forecast, $\tilde{\mu}_i$. This implies that, after every $\tilde{\mu}_i$ -positive (and, hence, also for μ -positive) price history \bar{p}^t , $\pi_i|_{\bar{p}^t}$ must be optimal relative to $\tilde{\mu}_i|_{\bar{p}^t}$. So the induced vector, $e|_{\bar{p}^t} = (\mu|_{\bar{p}^t}, \tilde{\mu}|_{\bar{p}^t}, \pi|_{\bar{p}^t})$, must be a private forecast equilibrium in the induced economy $\mathcal{E}_{\bar{p}^t}$.

For convergence we will argue that such an $e|_{\bar{p}^t}$ can be made arbitrarily close to a rational expectations equilibrium for sufficiently large t 's. In general, closeness of a private forecast equilibrium, $e = (\mu, \tilde{\mu}, \pi)$, to a rational expectations equilibrium can be described in two ways. We may have each plan π_i be exactly optimal relative to the private forecast $\tilde{\mu}_i$ and $\tilde{\mu}_i$ be close to the true forecast μ . Alternatively, we may have each plan π_i be almost optimal relative to the true forecast μ , where almost optimal means that the difference between the expected payoff under π_i and the real optimal expected payoff differ by a small amount. We may think of the former as being close in forecasts accuracy to a rational expectations equilibrium and of the latter as being close in optimality. Under the notions of closeness of forecasts, discussed below, given that our agents discount future payoffs, closeness in forecasts accuracy implies closeness in optimization.

Following the above discussion, it suffices to argue that the individuals' forecasts become accurate with time, i.e., $\tilde{\mu}_i$ merges with μ . For that purpose, we first distinguish two notions of closeness of measures.

Definition 7.1: For $\epsilon > 0$, we will say that $\tilde{\mu}_i$ is ϵ -close to μ if there is an event Q with $\tilde{\mu}_i(Q), \mu(Q) > 1 - \epsilon$, and for every $A \subset Q$, $|\mu(A)/\tilde{\mu}_i(A) - 1| < \epsilon$ ($0/0$ is defined to be 1).

As was discussed earlier, this is a strong notion of closeness for two reasons. It means being accurate even on events in the infinite horizon of prices. Moreover, for all such events A , no matter how small, the ratios of the probabilities must be close to 1.

A weaker, yet sufficient, notion is the following.

Definition 7.2: $\tilde{\mu}_i$ is finite-horizon ϵ -close to μ if, after every $\tilde{\mu}_i$ -positive (and, hence, μ -positive) price history \bar{p}_t ,

$$|\mu(p^{t+1} | \bar{p}^t) - \tilde{\mu}_i(p^{t+1} | \bar{p}^t)| \leq \epsilon$$

for every price p^{t+1} .

Notice that while the definition above calls for having almost accurate conditional probabilities for one step future forecasts, it implies asymptomatic closeness for any given finite number of future steps. For example, for two steps $\mu(p^{t+2}, p^{t+1} | \bar{p}^t) = \mu(p^{t+2} | \bar{p}^t, p^{t+1})\mu(p^{t+1} | \bar{p}^t)$. Thus, by making each factor in the right side closely approximated, we can ensure obtaining a good two step approximation. Alternatively, one could describe finite-horizon ϵ -closeness by looking into the infinite horizon but discounting deviations that are t steps into the future by a factor of λ^t using any fixed discount factor $\lambda < 1$.

It is easy to verify that being ϵ -close is asymptotically stronger than being ϵ -close on finite horizons. In other words, we can make $\tilde{\mu}$ finite horizon ϵ -close to μ by making it δ -close to μ for sufficiently small δ .

7a. Under Absolute Continuity

Theorem 1 (Learning the Correct Forecast): If $\mu \ll \tilde{\mu}_i$ then for every $\epsilon > 0$ and μ -almost-every price path p^1, p^2, \dots there is a time T such that for every history $\bar{p}^t = (p^1, \dots, p^t)$ with $t \geq T$, $\tilde{\mu}_i |_{\bar{p}^t}$ is a ϵ -close to $\mu |_{\bar{p}^t}$.

Thus, under absolute continuity, Bayesian updating of forecasts leads eventually to ϵ -correct forecasts for arbitrarily small ϵ .

Theorem 1 is an immediate application of the version of Blackwell-Dubins (1962) result stated and proven in Kalai-Lehrer (1993). The statement there is identical to Theorem 1 here except that prices are replaced by general random variables.

Theorem 2 (Convergence to Rational Expectations Equilibrium): Let $e = (\mu, \tilde{\mu}, \pi)$ be a private forecasts equilibrium and $\epsilon > 0$. Assume that $\mu \ll \tilde{\mu}_i$ for every i . Then, for μ -almost-every price path p^1, p^2, \dots there is a time T such that for every price history $\bar{p}^t = (p^1, \dots, p^t)$ with $t \geq T$, $e |_{\bar{p}^t}$ is an ϵ -rational expectations equilibrium of $\mathcal{E}_{\bar{p}^t}$.

In the above theorem, the notion of ϵ -rational expectations equilibrium was not specified. This is so because all three interpretations are correct. The first interpretation is that each private forecast, $\tilde{\mu}_i |_{\bar{p}^t}$, to which π_i is an optimal plan, is almost accurate. That is, it is ϵ -close to the correct forecast $\mu |_{\bar{p}^t}$ on all events including tail events. Obviously, this is a direct consequence of Theorem 1. Such closeness implies, as we mentioned earlier, that private forecasts are ϵ -close to the true one on finite horizons. And being close on finite horizons implies asymptotic closeness in optimality. That is, each $\pi_i |_{\bar{p}^t}$ is ϵ -optimal relative to the true forecast $\mu |_{\bar{p}^t}$.

7b. Under Diffused Beliefs

The following results rely on generalizations of Theorem 1. Instead of absolute continuity, we assume that μ_i is diffused around μ for every i .

Remark: Theorem 1 implies that if $\mu \ll \tilde{\mu}_i$ then $\tilde{\mu}_i$ is diffused around μ and, moreover, satisfies the separation condition. Therefore, the assumption that $\tilde{\mu}_i$ is diffused around μ is indeed a generalization of the absolute continuity assumption.

The results in Lehrer and Smorodinsky (1993, 1994) imply the following.

Theorem 3 (convergence to rational expectations equilibrium): Let $e = (\mu, \tilde{\mu}, \pi)$ be a private forecasts equilibrium and $\epsilon > 0$.

- (a) If for every i $\tilde{\mu}_i$ is diffused around μ , then for μ -almost every price path p^1, p^2, \dots , there is a set of periods, L , having density 1,² such that for every time $\ell \in L$, $e|_{\bar{p}^\ell}$ is an ϵ -rational expectations equilibrium where $\bar{p}^\ell = (p^1, \dots, p^\ell)$.
- (b) If the condition of (a) holds and $\tilde{\mu}_i$ has the separation property for every i , then L is of the form $L = \{T, T + 1, T + 2, \dots\}$.

The distinction between the conclusions in (a) and (b) is that in (a) there may be a sparse set of periods when there is a discrepancy between some private assessments and the truth. Thus, under the condition in (a) there is no way to ensure that the forecasts and the truth will be eventually (from a certain period of time on) approximately the same. Under the hypothesis of (b), in contrast, there cannot be infinitely chaotic periods when there is a significant gap between the private forecasts

² L has density 1 if its lower density is 1.

and the truth; the various distributions will eventually coincide.

Just as in Theorem 2, the notion of ϵ -rational expectations used in Theorem 3 was not specified. Here, only two interpretations are correct. The first is that each private forecast, $\tilde{\mu}_i|_{\bar{p}^t}$, to which the plan $\pi_i|_{\bar{p}^t}$ is optimal, is finite horizons ϵ -close to the true forecast, $\mu|_{\bar{p}^t}$. With this interpretation, Theorem 3 is a direct consequence of the results proved by Lehrer and Smorodinsky. The second interpretation is obtained again by considering the plan's optimality in relation to the correct forecast. Each $\pi_i|_{\bar{p}^t}$ is ϵ -optimal relative to $\mu|_{\bar{p}^t}$.

8. Application

In this section we apply the main results to the imbalanced storable goods economy \mathcal{E} of Section 2. Recall that L is the short commodity there. The following corollary shows that while private forecast equilibrium allows speculative behavior it does not allow it in the long run. If we start with truth compatible beliefs, the price of the short commodity must converge to 1.

Corollary: Let μ be the correct forecast of a private forecast equilibrium of \mathcal{E} . Suppose the private forecasts are truth-compatible in the sense of absolute continuity or being diffused with the separation property. Then with μ -probability one $\lim_{t \rightarrow \infty} P_R^t = 0$.

Proof: Let $e = (\mu, \tilde{\mu}, \pi)$ be the private forecast equilibrium of the economy \mathcal{E} . By the main results, Theorems 2 and 3, for any $\epsilon > 0$, for μ almost every price path p_1, p_2, \dots there will be a time T such that, for all $t \geq T$, $e_{\bar{p}^t}$ is an ϵ -rational expectations equilibrium of $\mathcal{E}_{\bar{p}^t}$. (in the sense of ϵ -optimizing relative to the true forecast). Consider now the economy $\mathcal{E}_{\bar{p}^{t-1}}$ with the equilibrium of $e_a = (\mu|_{\bar{p}^{t-1}}, \tilde{\mu}|_{\bar{p}^{t-1}}, \pi|_{\bar{p}^{t-1}})$. This is the equilibrium where the first period prices were already drawn to be p^t , the agents

know this and have factored it into their posteriors, and it must be an ϵ -rational expectations equilibrium. Now we can truncate $\mathcal{E}_{\bar{p}^{t-1}}$ after $M (=M(t))$ additional periods so that the utility of any agents for consumption after period M is 0. For this truncated economy, $\mathcal{E}_{\bar{p}^{t-1}}, e_a$ is an ϵ -rational expectations δ -equilibrium where both ϵ and δ can be made arbitrarily small by choosing T and then M large enough. Now, consider $e_b = (\mu |_{\bar{p}}, \pi |_{\bar{p}^{t-1}})$. It must be a rational expectations η -equilibrium of $\mathcal{E}_{\bar{p}^{t-1}}$ where η can be made arbitrarily small by a choice of sufficiently large T and M .

Thus, to complete the proof it suffices to show that for a finite horizon economy as above, with known first period prices, at every rational expectations η -equilibrium the initial price of commodity L can be proven to be arbitrarily close to 1 by making η sufficiently small.

It is straightforward to verify that in the truncated economy, for a fixed forecast, the derived utility that an agent can obtain from a vector (w_L, w_R) of initial endowments is linear. This means that if μ is a forecast of a rational expectations η -equilibrium of $\mathcal{E}_{\bar{p}^{t-1}}$ it is a perfect forecast ξ -equilibrium of the 2-agent economy $\hat{\mathcal{E}}_{\bar{p}^{t-1}}$, which modifies $\mathcal{E}_{\bar{p}^{t-1}}$ as follows. Give agent $\hat{0}$ initially the excessive amount of the excessive commodity R , i.e., $\sum_i (w_{iR} - w_{iL})$, and nothing in all the following periods. Give agent $\hat{1}$ the rest of the endowments of all other agents combined at all times. Moreover, by making T and M sufficiently large we can make ξ sufficiently small. Notice that in the economy $\hat{\mathcal{E}}_{\bar{p}^{t-1}}$, agent $\hat{1}$ controls all the usable commodities and thus at an ξ -equilibrium he consumes everything except a quantity of utility at most ξ . Thus, player $\hat{0}$ can be made to consume arbitrarily small amounts. This would be impossible if the initial price p_R^t could not be made arbitrarily small.

References

- Blackwell, D. and L. Dubins (1962), "Merging of Opinions with Increasing Information," Annals of Statistics, 38, 882-886.
- Blume, L. and D. Easley (1992), "Rational Expectations and Rational Learning," mimeo, Cornell University.
- Bray, M. and D. Kreps (1987), "Rational Learning and Rational Expectations," in G. R. Feiwel (ed.), Arrow and the Ascent of Modern Economic Theory, pp. 597-625, New York: New York University Press.
- Jordan, J. and R. Radner, (1982), "Rational Expectations in Microeconomic Models: An Overview," Journal of Economic Theory, 26 (2), 201-223.
- Kalai, E. and E. Lehrer (1993), "Rational Learning Leads to Nash Equilibrium," Econometrica, 61, 1019-1045.
- Kalai, E. and E. Lehrer (1994), "Weak and Strong Merging of Opinions," Journal of Mathematical Economics, forthcoming.
- Lehrer, E. and R. Smorodinsky (1993), "Compatible Measures and Merging," mimeo.
- Lehrer, E. and R. Smorodinsky (1994), "Repeated Large Games with Incomplete Information," mimeo.
- Lucas, R. E. (1972), "Expectations and the Neutrality of Money," Journal of Economic Theory, 4, 103-124.
- Muth, J (1961), "Rational Expectations and the Theory of Price Movements," Econometrica, 29, 315-335.
- Radner, R. (1967), "Equilibre des Marches a Terme at au Comptout en Cas d'Incertitude," Cahiers d'Econometrie, C.N.R.S. Paris 4, 35-42.