

Rational Limits To Arbitrage

By

Jean-Pierre Zigrand

DISCUSSION PAPER 392

October 2001

FINANCIAL MARKETS GROUP
AN ESRC RESEARCH CENTRE

LONDON SCHOOL OF ECONOMICS



Any opinions expressed are those of the author and not necessarily those of the Financial Markets Group.

ISSN 0956-8549-392

Rational Limits to Arbitrage

Jean-Pierre Zigrand*

*Department of Accounting and Finance,
and Financial Markets Group
The London School of Economics
Houghton Street, London, WC2A 2AE
E-mail: j.p.z grand@lse.ac.uk*

Abstract

It is often argued that asset prices exhibit patterns incompatible with the behaviour of rational, optimizing agents. This paper proposes a rational framework which generates asset prices which appear irrational. This is accomplished by studying rational expectations equilibria in the presence of two realistic market frictions: immediacy risk (agents have to submit their demand functions before they know the equilibrium price) and asset-specific orders (investors have to submit one separate demand for each asset, which may not be contingent upon the prices of the other assets). We study some properties of such equilibria, in particular the prevalence of arbitrage and of informational inefficiencies.

JEL Classification Numbers: D50, D8, G12.

*I thank Pierre-André Chiappori, Roger Guesnerie, Erzo Luttmer, Heracles Polemarchakis, José Scheinkman and participants at the LSE, at the CEPR Summer Symposium in Gerzensee, the Workshop on Markets and Information in Venice, the University of Amsterdam, Brunel University, the Economic Theory Seminar at CORE and the University of Essex for helpful comments. All errors are mine.

1 Introduction

In this paper we investigate issues of asymmetric information and strategic intermediation in environments with realistic trading frictions.

Few practitioners would disagree with the statement that markets occasionally exhibit the following two imperfections: (i) they may be temporarily informationally inefficient, and (ii) they may temporarily be out of sync and exhibit mispricings and arbitrage opportunities.

As to (i), informational inefficiencies have traditionally been modelled by preventing revelation of private information by means of noise trading (see for instance (Grossman and Stiglitz 1980)) or of stochastic asset endowments (see for instance (Diamond and Verecchia 1981)). In this model we introduce one further channel leading to informational inefficiencies.

As to (ii), modelling mispricing situations has not been dealt with in an entirely satisfactory manner. In standard general equilibrium models, arbitrage is impossible since agents submit their Marshallian demand functions to the auctioneer, and any zero of the aggregate excess demand function is arbitrage-free by construction. In some sense, the individual demand functions, which map any conceivable price vector to a quantity vector, incorporate arbitraging strategies. So if the auctioneer was to fix a price allowing for a free lunch, the arbitrage actions coded into the Marshallian demand functions trigger, by price-taking, unbounded demands and supplies all on the same side of the market, proving that all equilibrium prices need to be free of arbitrage.

One obvious approach to generate arbitrage opportunities at an equilibrium is to introduce transaction costs or short-sale constraints to the model. This way of proceeding is a bit unsatisfactory, however, in accounting for mispricings since one effectively “explains” a \$1 mispricing gap by a transaction cost of \$1. It is one thing to introduce a device that generates arbitrage opportunities, but it is another thing to determine the extent of mispricing remaining in the market after trading. Furthermore, different investors face different transaction costs, and some investors face nearly zero marginal transaction costs. A more subtle approach to modelling equilibrium arbitrage opportunities introduces frictions via costly participation constraints and endogenous entry, see for instance (Fremault 1993), (Holden 1995) and (Zigrand 1999). Such frictions generate mispricings, but do not also impose exogenous limits to arbitrage, and the question as to the limits to arbitrage becomes interesting and relevant. For instance, in (Holden 1995) and (Zi-

grand 1999) the extent of arbitrage is limited by strategic considerations in a Cournot-Walras game, while in (Shleifer and Vishny 1997) the extent of arbitrage is limited by hidden information considerations.

The current paper can thus be viewed as contributing to the modelling of the two facets without relying on overly crude or counterfactual assumptions, but by modelling a plausible microstructure environment. First, we generate informational inefficiencies and arbitrage opportunities in a natural manner. Second, we show that the extent to which arbitrage opportunities survive at equilibrium is endogenous, even in a fully rational environment.

The framework we propose drops two simplifying (but arguably counterfactual) assumptions regularly made in financial economics:

Instantaneous Execution. Indeed, investors usually face immediacy risk or execution lags when trading assets, because they do not have instantaneous access to the floors and they cannot exactly forecast the prices at which their orders will be executed (or whether they will be executed in case of limit orders). In particular, this prevents investors from being able to execute riskless arbitrage strategies. This assumption is similar to (Kyle 1985), except that here (as in (Kyle 1989)) investors are free to submit limit orders (schedules), rather than limiting their trades to market orders (quantities) which would effectively prevent learning from prices. We thus reverse the standard timing, which assumes that investors first observe prices and then decide about their portfolio weights. The unpredictability we propose is common on financial markets and well-known to practitioners, as revealed by the universal use of limit orders on more volatile markets. Some of the underlying volatility may be due to new information about returns occurring between the instant the order was released and its execution, but a large part is due to portfolio rebalancing. So even an informational insider would face immediacy risk for he would be unlikely to be able to accurately forecast either portfolio rebalancing or the release of information to other informed participants. In terms of modelling, the unpredictability of non-informational portfolio rebalancing (“liquidity trades”) is due to the fact that the structure of the economy is not common-knowledge. This uncertainty about economic fundamentals is often simplified by relying on noise traders.

Marshallian Demand Schedules. In reality investors cannot submit an order for asset a contingent on prices of assets other than a . They have to submit schedules (composed of market and limit orders) to the auctioneers

(the limit-order book, market makers, specialists), one for each asset, that cannot depend on events other than the price of the asset. We refer to this restriction as asset-specificity. For instance, investors cannot submit a demand schedule that says that they want to purchase 1 unit of Apple at a price of \$x if the price of Microsoft is \$y. Effectively this also implies that we can think of auctioneers as being “local auctioneers,” one for each asset.

The Model. There is a unique perishable consumption commodity per date and state, chosen as the numeraire. All investors are price-takers. This is in contrast to (Klemperer and Meyer 1989) and (Kyle 1989) where oligopolists submit supply functions, realizing the impact on prices via a downward sloping demand schedule. We assume, in order to generate randomness, that the state of information that determines preferences and the endowment structure is not known initially. Each investor is then informed of his own preferences and endowments, but he gets no direct signal as to the structure and composition of the market, summarized by the preferences and endowments of the other agents.

Agents submit their asset demand schedules and the auctioneers compute aggregate demand schedules and clear markets. The resulting equilibrium prices may allow for arbitrage. This is not the end of the story, though. If investors have rational expectations, then they are able to refine their information by the information revealed by prices. Even though they cannot change their orders any more, they built in such contingencies into the demand schedules they originally handed over. Basically, we can view investor h 's optimization program as a team maximization problem with as many team members as there are different assets, $\{(h, a), a = 1, \dots, A\}$. Member a of investor h 's team (denoted by (h, a)) submits the schedule to the auctioneer of asset a , incorporating the following line of reasoning:

”I am rational in the sense that I know the equilibrium price function mapping shocks to prices. Unfortunately, I can only observe the price of asset a , and I cannot communicate with any of the other members (h, b) who observe q_b , $b \neq a$. But, given the structure of the economy, I can update my beliefs as to the composition of the whole market environment from any one price. The resulting price of asset a at which my order will be executed contains that information, allowing me to refine my information as to the prices of all the other assets.“

Related Literature. The investors' optimization problem has the same team flavor that was analyzed by (Marschak and Radner 1971) and later reviewed by (Radner 1986a) and (Radner 1986b). The two latter papers study related but simpler problems than the one in this paper, in particular as to the team-wide resource constraints, to the correlatedness of signals and to the rational information extraction from prices or actions. The REE aspects that occur here have been analyzed, in different forms, among others, by (Green 1977), (Grossman 1977), (Radner 1979) and (Allen 1981). The fact that each investor is unsure whom he is trading against is reminiscent of (Diamond and Verecchia 1981) where investors are endowed with random asset supplies. The only paper on arbitrage and information transmission we are aware of is (Fremault 1993). There is also a growing literature on the limits of arbitrage, especially intertemporal arbitrage, see for instance (Dow and Gorton 1994) or (Shleifer and Vishny 1997). The emphasis in this paper, by its static nature, captures relative mispricing of one portfolio of assets compared to another one.

Structure of the Paper. The structure of the economy is presented in Section 2. The two innovative market microstructure assumptions used in this paper are described in Sections 3 and 4 respectively. Section 5 introduces the investor's optimization problem, and Section 6 deals with the characterization of the rational expectations equilibrium. Section 7 concludes. All proofs are relegated to the appendix.

2 The Economy

As mentioned in the introduction, individuals face two layers of risk in this economy, first about the trading environment (the state of information) and then about the realization of endowments and asset payoffs (the state of nature). At time zero they face uncertainty about the market participants' endowments and preferences, which translates into uncertainty about prices. The asset demand functions they submit are not required to be either market or limit orders, but they are required to only depend on their own price, i.e. they need to be asset-specific. For instance, the demand schedule for asset a is a function of the price of asset a , q_a , only, as opposed to the whole price vector q . At time one this uncertainty is resolved, asset orders get executed and the

households consume. They are, by the very nature of immediacy risk, unable to retrade at the equilibrium price. The state of nature is realized at time two and final consumption occurs. It might be convenient to visualize investor h 's utility maximization program in the following fashion. Given that asset specific demand functions have to be handed over to the auctioneer before the state of information is realized, every investor's demand function for some asset a can be thought of as being represented by a team member (h, a) having the same preferences as the investor, and with the restriction that the A members cannot communicate, so that (h, a) 's private information (over and above the private information arising from h 's endowments and preferences) corresponds to the observation of the price of asset a , q_a .

2.1 General Setup of the Economy

We shall impose the following strong but standard assumptions on investors:

H1 H , the set (and the cardinality)¹ of households, is assumed to be finite, and there is only one consumption commodity. Trade occurs over three time periods: team members trade assets in period zero, they consume in period one, and finally they consume the proceeds (and are allowed to trade consumption) at time two.

H2 $(\mathcal{E}, \mathcal{F}, \mu)$ is the (complete) probability space of the parameters describing the states of information. (S, \mathcal{G}, p) with $S \equiv \{1, \dots, s, \dots, S\}$, $S < \infty$ is the probability space of states of nature.² They are assumed to be independent: knowledge of one's preference and endowments profile does not help one to predict the likelihood of states of nature, and thus of future asset payoffs. $\mathcal{E} \subset \mathbb{R}^E$ is a compact Borel set and contains an open subset of \mathbb{R}^E , with $E < \infty$. Assume that \mathcal{E} is a C^∞ manifold, possibly with boundary. μ is absolutely continuous with respect to E -dimensional Lebesgue measure λ^E on \mathbb{R}^E . For technical reasons, we assume that $\mu(\partial\mathcal{E}) = 0$ (i.e. \mathcal{E} is a contented set) and we can assume as well that \mathcal{E} is equal to the support of μ .

¹No confusion should arise by denoting the cardinality of a set by the same symbol as the set itself.

²So while the element $\epsilon \in \mathcal{E}$ determines, among others, the endowments and preferences state by state, the element $s \in S$ determines which element of the endowment and preference vector is actually realized, as well as the payoff of each asset.

H3 The information revelation unravels in three steps. Initially, investors have no information and observe their own preferences and endowments. We assume that the parameters determining the endowments $\{\omega^h\}_{h=1}^H \equiv \{\omega_0^h, \omega_1^h, \dots, \omega_S^h\}_{h=1}^H \in \mathbb{R}_{++}^{(S+1)H}$ and the utility functions $\{U^h\}_{h=1}^H \in \mathcal{U}^H$ (the set \mathcal{U} will be defined under H5) lie in the parameter space $\prod_{h=1}^H \mathcal{P} \subset \mathbb{R}^{HP}$ and depend for each investor h on the state of information $\epsilon \in \mathcal{E}$ via a measurable $g^h : (\mathcal{E}, \mathcal{F}) \rightarrow (\mathcal{P}, \mathcal{B}(\mathcal{P}))$, where \mathcal{B} denotes the Borel σ -algebra and \mathcal{P} is a smooth submanifold of \mathbb{R}^P , with $E \geq P + 2$. The mapping g^h is smooth on the interior of \mathcal{E} and is assumed to be a submersion. Thus we have

$$\epsilon \xrightarrow{g^h} e^h = g^h(\epsilon) \mapsto (\omega^h[e^h], U^h[e^h])$$

When ϵ is realized, each investor h knows the realization³ $e^h = g^h(\epsilon)$ (and thus knows $\{\omega^h, U^h\} \circ g^h \circ \epsilon$), but ignores the realization of $\{\omega^k, U^k\} \circ g^k \circ \epsilon$, $k \neq h$. $\mathcal{F}^h = \sigma(g^h)$ describes investor h 's information after learning their own parameters⁴. Since the uncertainty in our model is due to the composition of the market, it seems natural to assume that $\sigma(\mathcal{F}^1, \mathcal{F}^2, \dots) = \mathcal{F}$: prices cannot depend on more information than possessed by all the agents combined.

When trading subsequently the team member (h, b) learns the price of asset b , q_b , and refines his information to $\mathcal{F}^{h,b} = \mathcal{F}^h \vee \sigma(\phi_b)$.

H4 Assets pay off in the last period. The payoff matrix R is of $\dim S \times A$ and of full column rank, where A is the number of assets traded, $A \leq S < \infty$. Assets pay off in the numeraire, which is also the only consumption commodity. Row s of R is denoted by $d_s = (d_{1,s}, \dots, d_{A,s}) \in \mathbb{R}^A$. $d_{a,s}$ represents the numeraire payoff of asset a in the state of nature s . Also, assume there is $y \in \mathbb{R}^A$ such that $Ry \gg 0$. This assumption assures that not all prices are no-arbitrage prices, see (Geanakoplos

³In (Allen 1981) and (Allen 1985a), after the realization of the state of information, investors are assumed to know their endowments and prices but they find out about their preferences from prices. In her model, that's the only incentive to learn from prices. Here investors are assumed to know their endowments and preferences, and use that knowledge to predict prices.

⁴This is similar to the setup in (Green 1977). Notice that we chose this formulation over the formulation that agents refine their information from their endowments and preferences directly to keep the problem finite-dimensional. Noise trading is a special case of this setup.

and Polemarchakis 1986). For instance, it rules out $R = \begin{bmatrix} 1 \\ -1 \end{bmatrix}$ and $R = \begin{bmatrix} 1 & 1 \\ -1 & -1 \end{bmatrix}$, i.e. economies where no asset and no combination of assets can be used as a “safe asset” to carry purchasing power over to the next period.

H5 The consumption set is $X^h \equiv \mathbb{R}_{++}^{S+1}$. Von-Neumann–Morgenstern utility functions are smooth, strictly differentially concave, satisfy $\lim_{x \rightarrow 0} u_s^h(x) = +\infty$ for $s = 0, 1, \dots, S$, and are differentiable monotonic, by which we mean that $u_s^h(x) \geq k > 0$ for all $x > 0$, $s = 0, \dots, S$. Notice that the Von-Neumann-Morgenstern utility functions may contain discount factors. Investors are assumed to be expected utility maximizers as far as period-zero price uncertainty is concerned as well. We also define $U^h \equiv u_0^h + \sum_s p_s u_s^h$. \mathcal{U} is the set of such utility functions.

H6 Initial commodity endowments are $\omega^h \in \mathbb{R}_{++}^{S+1}$.

H7 The asset demand function for some asset a must be measurable w.r.t. the information of team member (h, a) , $\mathcal{F}^{h,a}$. It is denoted by $\theta_a^h = f_a^h(q_a, e^h)[\phi]$, where ϕ is the investor’s forecast function (to be defined in more detail below, basically it is a function mapping an $\epsilon \in \mathcal{E}$ to an equilibrium price vector $q = \phi(\epsilon)$).

Having reviewed the main ingredients, in the next two sections we elaborate upon the two assumptions that we consider innovative and that play a crucial role in the results derived.

Remark 1 *As alluded to before, our model encompasses a number of possible sources of immediacy risk. We mentioned preference and endowment risk. More restrictively, one may assume that preferences and endowments of agents $h = 1, \dots, H$ are common knowledge, but that agent $H+1$ is a representative noise trader. The noise trades of $H+1$ are randomly chosen in $\mathcal{E} \subset \mathbb{R}^A$, the set of noise trades on which the noise trader has no need to default and which still allow markets to clear, i.e. $\mathcal{E} = \{\epsilon \in \mathbb{R}^A : d_s \cdot \epsilon \in [-\omega_s^\epsilon, \sum_h \omega_s^h]\}$, all $s \in S\}$, where $\omega_s^\epsilon \in \mathbb{R}_{++}$ is the noise trader’s endowments in period 2 in state s .*

3 The Role of Immediacy Risk.

First, we would like to discuss two interpretations of Rational Expectations Equilibria, REE in short. In order to define by what we mean by REE, we need:

Definition 1 A price function or a forecast function ((Radner 1979)) ϕ is any $\mathcal{F}/\mathcal{B}^A$ -measurable mapping from the set of states of information \mathcal{E} to the space of asset price vectors, a subset of \mathbb{R}^A .

By a REE we then mean:

Definition 2 A Rational Expectations Equilibrium is a price function ϕ and asset demands $\{f_a^h\}_{a \in A, h \in H}$ solving the investors' problems, such that markets clear and forecasts are confirmed:

$$\sum_{h \in H} f_a^h(\phi(\epsilon), g^h(\epsilon))[\phi] = 0 \quad (\text{almost all } \epsilon \in \mathcal{E}), (\forall a \in A)$$

Commodities markets then clear by Walras' Law.

There are two interpretations of REE that are relevant in our context: First, in the market maker interpretation, the economy works as follows:

1. The auctioneer or the market maker knows everything about the deep parameters of the economy, summarized in the state of information ϵ . He makes the pricing function ϕ known to all participants.
2. Since he knows all the deep parameters ϵ , he quotes a price vector q that equals $\phi(\epsilon)$.
3. Investors observe q , update their information via the preimage $\phi^{-1}(q)$, and submit their market orders (i.e. quantities only).
4. By construction, all market orders exactly sum up to 0.

Alternatively, the algorithm approach seems closer to the actual workings of order-driven markets, where the auctioneer is a piece of software that takes the demand and supply schedules as inputs and matches them whenever possible.

1. The investors have the same forecast function ϕ , and we can think of the auctioneer as being a computer.

2. Investors submit (excess) demand functions. They are functions of prices q , since the investors do not observe market clearing prices at the time they submit the schedules. However, the demand schedules incorporate the updating $\phi^{-1}(q)$. In that respect they yield the equilibrium portfolio demand for each market clearing price.
3. The auctioneer sums up the demands and solves for q (if a solution exists). By investor rationality, $q = \phi(\epsilon)$ for each realization of ϵ (the forecast function coincides with the pricing function).

In this second interpretation, investors never explicitly observe prices and then update their beliefs and demands. The information refinement is directly integrated into the schedules. This introduces price uncertainty at the time the investors submit their orders. Investors know the price distribution, but they cannot be exactly sure at which prices their orders will be executed. This we call immediacy risk and it basically reverses the standard timing which assumes that investors first observe prices and then submit their orders. The traditional economy where agents have rational expectations and observe prices before choosing their optimal actions is called a Radner economy.

We can show the following equivalence relation:

Lemma 1 *If, in a Radner economy, we impose the additional constraint that agents have to submit their excess demand schedules before prices are observed, the Rational Expectations Equilibrium will be unaffected.*

In other words, both interpretations of REE coincide.

4 The Role of Asset-Specific Orders.

Let us depart from the Radner model by adding the unique constraint upon investors that the demand function for asset a may only be explicitly conditioned on its own price, rather than on the entire price vector. We call this restriction “asset-specificity.” Since we abstract from immediacy risk, the investor is able to observe the Radner price vector $q^R = \phi(\epsilon)$ before forming his demand, and submits the function $f_a^h(q_a; q_{-a}^R)$, a function of the first variable, q_a , only. The investor simply hides the second element q_{-a}^R from the auctioneer. The auctioneer solves $\sum_h f_a^h(q_a; q_{-a}^R) = 0$ for $q_a^R = \phi(\epsilon)$:

the standard equilibrium is trivially unaffected by this constraint. We have shown

Lemma 2 *Any REE in a Radner economy is also an equilibrium if asset demands are constrained to be asset-specific.*

Next, let us keep the asset-specificity constraint, but add the further constraint of immediacy risk. We now face a modelling choice as to whether \mathcal{E} , the support of μ , is a countable or uncountable set. It turns out that in order to generate nontrivial results, the underlying probability space has to be uncountable (as defined and assumed in H2). The reason is in fact very simple. If the state space is finite, we can show that the team member (h, a) can induce the payoff relevant joint signal from the price of asset a only (and the same must hold for all assets $a \in A$), i.e. typically any single price q_b is fully revealing. Indeed, an argument identical to Radner's ((Radner 1979)) shows that the following system of aggregate excess demands (a superscript "R" means the demand function is a standard REE demand function) has typically no solution (the two different realizations of the joint signals are e and e'):

$$\begin{aligned} \sum_h f_a^{R,h}(q, e^h) &= 0 & (\forall a \in A) \\ \sum_h f_a^{R,h}(q, e^{h'}) &= 0 & (\forall a \in A) \\ q_b &= q'_b \\ e &\neq e' \end{aligned}$$

Let us denote the Radner REE by $\hat{\phi}$ and assume asset-specificity away. Then every single price is fully revealing all the other prices via $q_{-a} = \hat{\phi}_{-a}(\hat{\phi}_a^{-1}(q_a))$. Denote such a price vector by q^R . Now do impose asset-specificity as well, and the investor hands over the schedule

$$f_a^h(q_a, e^h) = f_a^{R,h}(q_a, \hat{\phi}_{-a}(\hat{\phi}_a^{-1}(q_a)), e^h)$$

to the auctioneer clearing market a , and the full-communication equilibrium ensues, as $q^R = \hat{\phi}(\epsilon)$ (the Radner equilibrium price which solves $\sum_h f_a^{R,h}(q_a^R, q_{-a}^R, e^h) = 0$, all $a \in A$) solves this new system of equations as well: $\sum_h f_a^{R,h}(q_a^R, \hat{\phi}_{-a}(\hat{\phi}_a^{-1}(q_a^R)), e^h) = 0$, all a because $\hat{\phi}_{-a}(\hat{\phi}_a^{-1}(q_a^R)) = q_{-a}^R$.

Loosely speaking, however, investors have to be a bit 'more precise' because they have to be able to induce the joint signal from a single price, rather than from a price vector. We sum this discussion up in the following lemma.

Lemma 3 (Discrete States) *When the state of information is a discrete random variable, adding the combined assumption of asset-specific orders and of immediacy risk does not change the Radner REE.*

Typically, this is no longer true when an asset price does not reveal all payoff relevant information. For the Radner REE $\hat{\phi}$ to still be an equilibrium, it has to solve both $\sum_h f_a^{R,h}(\hat{\phi}_a(\epsilon), \hat{\phi}_{-a}(\epsilon), g^h(\epsilon))[\hat{\phi}] = 0$, the defining equilibrium condition, and $\sum_h f_a^h(\hat{\phi}_a(\epsilon), g^h(\epsilon))[\hat{\phi}] = 0$, $a = 1, \dots, A$. The function $f_a^h(\hat{\phi}_a(\epsilon), g^h(\epsilon))[\hat{\phi}]$ is the demand function for team member (h, a) , who decides about the holdings of asset a for agent h and whose information is given by $\mathcal{F}^{h,a}$. So unless the information gathered from $\hat{\phi}_{-a}$ is superfluous, it will typically be the case that $\sum_h f_a^h(\hat{\phi}_a, e^h)[\hat{\phi}] \neq 0$: the Radner REE pricing function $\hat{\phi}$ no longer clears markets, and the new one that does is denoted by ϕ . Consider the system of equations

$$\begin{aligned} 0 &= \sum_h f_a^{R,h}(q_a, q_{-a}, e^h)[\hat{\phi}] & , a = 1, \dots, A \\ 0 &= \sum_h f_1^h(q_1, e^h)[\hat{\phi}] \end{aligned}$$

To see intuitively why we cannot typically expect a solution to this system of $A + 1$ equations in A unknowns, assume to the contrary that there is a solution at some $\epsilon \in \mathcal{E}$. Perturb the endowments of agents 1 and 2 as $d\omega_s^1 = d_s^1$ and $d\omega_s^2 = -d_s^1$, $s = 1, \dots, S$, and $d\omega_0^1 = -q_1 = -d\omega_0^2$. Call this state of information ϵ' . It follows that in the Radner equilibrium prices are unaffected, $\hat{\phi}(\epsilon) = \hat{\phi}(\epsilon')$ and $df_1^{R,1} = -1 = -df_1^{R,2}$. However, Radner pricing $\hat{\phi}(\epsilon')$ is typically no longer an equilibrium in the immediacy risk economy. Agent 1 observes a new endowment vector and updates his information to \mathcal{F}^1 , but he does not observe $d\omega^2$, and similarly for agent 2. If $\hat{\phi}$ is still an equilibrium, the price is $q_1 = \hat{\phi}_1(\epsilon) = \hat{\phi}_1(\epsilon')$. But since for non-degenerate g functions the information of agents 1 and 2 are affected asymmetrically, we expect that $df_1^1 \neq -df_1^2$, so that $\sum_h f_1^h(\hat{\phi}_1(\epsilon'), g^h(\epsilon'))[\hat{\phi}] \neq 0$.

5 The Investor Problem

In this section we analyze the optimal investment decisions of a rational investor or household. Such an investor's views of the workings of an economy can be summarized by his forecast function, ϕ . Assume that $\phi : \mathcal{E} \rightarrow \mathbb{R}^A$ is $\mathcal{F}/\mathcal{B}^A$ measurable, where \mathcal{B}^A is the appropriate Borelian sigma-algebra.

The household problem consists in choosing $\mathcal{F}^{h,a}$ -measurable schedules⁵ $f_a^h(q_a, e^h)$ for every asset a , where $\mathcal{F}^{h,a} = \sigma(g^h, \phi_a)$ and where $e^h = g^h(\epsilon)$ and $q_a = \phi_a(\epsilon)$. A particular team member (h, b) tries to solve the interim problem [P:i]:

$$\begin{aligned} & \max_{\theta_b^h} E[U^h | \mathcal{F}^{h,b}] \\ & = \max_{\theta_b^h} E \left[u_0^h \left(\omega_0^h - \sum_{a \neq b} f_a^h(\phi_a, e^h) \phi_a - \theta_b^h q_b \right) \right. \\ & \quad \left. + \sum_s p_s u_s^h \left(\omega_s^h + \sum_{a \neq b} d_{a,s} f_a^h(\phi_a, e^h) + d_{b,s} \theta_b^h \right) \right] \Big| \mathcal{F}^{h,b} \end{aligned} \quad [\text{P:i}]$$

Assume that the differentiation and integration operators are exchangeable. Then the first order condition with respect to θ_b^h for an interior solution is

$$q_b = \sum_s d_{b,s} \left[\frac{E[p_s u_s^h (\omega_s^h + \sum_a f_a^h(\phi_a, e^h) d_{a,s}) | \mathcal{F}^{h,b}]}{E[u_0^h (\omega_0^h - \sum_a f_a^h(\phi_a, e^h) q_a) | \mathcal{F}^{h,b}]} \right]$$

The term in the outer brackets, call it $\lambda_s^{h,b}$, is the state-price that is specific to asset b , or maybe to a specific market-maker or exchange. The intertemporal marginal rate of substitution cannot, in contrast to standard models, play the role of a stochastic pricing kernel since the $\lambda^{h,b}$ used in the pricing of asset b does not typically price any other asset. This is the crucial

⁵Since each optimal portfolio θ_a^h will to be measurable w.r.t. $\sigma(g^h, \phi_a)$, there is a Baire function f_a^h satisfying $\theta_a^h = f_a^h(q_a, e^h)$ (refer for instance to (Krickeberg 1965), Theorem 2.5 on p.139). The optimal demand functions f_a^h also depend on the entire function ϕ (as opposed to a particular realization of the random variable ϕ), since the agent updates via $\phi_a^{-1}(q_a)$ and then forms his views about the other prices via the correspondence $\phi_{-a}(\phi_a^{-1}(q_a))$. We might thus write them as $\theta_a^h = f_a^h(q_a, e^h)[\phi]$. The dependence on e^h and ϕ (as a function) will often be left implicit in order to simplify notation.

asset-pricing implication of this model that will allow us to generate inefficient and mispriced equilibrium prices without recourse to ad-hoc assumptions. Informational innovations in one part of the markets are not immediately and simultaneously incorporated into all other prices, as standard REE models assume.

Given the demand functions for the other assets, $f_{-b}^h(\cdot)$, and the realization q_b , the FOC determines the optimal amount of asset b to buy, θ_b^h . Repeating that exercise for every value of q_b , we can trace out a demand function. However, we still have to guarantee that demand functions exist, i.e. that there are such functions solving the fixed-point problem.

Proposition 1 (Existence of demand functions) *Assume the investment opportunity set is truncated to a compact set. For given (e^h, ϕ) and under the standard assumptions, there exists a $\mathcal{F}^{h,a}$ -measurable demand function $f_a^h : q_a \mapsto \theta_a^h = f_a^h(q_a, e^h)[\phi]$ for any asset $a \in A$.*

The investment opportunity set is truncated to a compact set for we want to show the existence of demand functions for any pricing function ϕ and any realization $\epsilon \in \mathcal{E}$, some of which may allow for an arbitrage.

6 Equilibrium

The existence of REE with truly partially revealing prices is in general problematic, c.f. (Green 1977), (Kreps 1977) and (Jordan and Radner 1982) for instances of non-existence, and (Allen 1985a), (Allen 1985b), (Ausubel 1990), (Heifetz and Polemarchakis 1998) and the family of CARA-normal finance models for different special contexts in which equilibria do exist. Over and above the problems typically encountered in this strand of literature, the set of admissible equilibrium price functions is a delicate issue here. It is well-known that in standard economies the image of ϕ is a subset of \mathcal{Q} , the set of no-arbitrage prices defined as follows:

Definition 3 *A price vector q admits No Arbitrage (NA) if there is no trade $y \in \mathbb{R}^A$ such that*

$$\begin{aligned} & \text{either } y'q \leq 0 \quad \text{and} \quad Ry > 0, \\ & \text{or } y'q < 0 \quad \text{and} \quad Ry \geq 0 \end{aligned}$$

The set of such vectors is denoted by \mathcal{Q} .

This is overly restrictive in our setup: the fact that $\phi(\epsilon)$ allows for arbitrage does not prevent it from being an equilibrium, since no investor may realize that there is an arbitrage (as defined in Definition 3, which makes no reference to information sets). Further, assume that ϕ is such that at $\phi_a(\epsilon)$ team member (h, a) knows for sure that there is an arbitrage strategy involving asset a . This does not necessarily cause the price system to be nonviable, since the other legs of the operation may not realize there is an arbitrage. Worse, suppose all legs of the arbitrage trade know there is an arbitrage at $\phi(\epsilon)$, but that one or more of them do not know that the others know it. In short, the team members will not be able to fully take advantage of this opportunity unless it is common knowledge (CK) among the relevant subset of team members. They will only take advantage of what in effect amounts to a very good deal. But due to the possibility that the submitted demands do not form an arbitrage portfolio, their decision problem is well-defined and a solution exists. The section below contains a number of illustrative examples that show that the set of admissible pricing functions ϕ is much larger than the set in standard RE economies. Even price realizations that are far inside the set of arbitrage prices may be compatible with a REE.

6.1 Illustrations

Assume that the first asset pays off 1 in the first state only, while the second asset is a riskless bond, $R = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$. The absence of arbitrage corresponds to the requirement that $q_2 > q_1 > 0$. As a convention, on the figures below the price realization is represented by a bold point, and the set of possible realizations of the price vector by the respective team member is represented by a bold line. Example 1 shows a ϕ that is not compatible with equilibrium because it is CK that there is an arbitrage.

Insert
figure 1
here

In Figure 2, team member $(h, 1)$ knows there is an arbitrage, but $(h, 2)$ does not. The bold line represents the set of all possible prices according to member $(h, 1)$'s information set, and it is included in the set of arbitrage prices.

Insert
figure 2
here

The next example (Figure 3) shows a ϕ and a price realization at which both team members know there is an arbitrage, but no-one knows that the

other knows.

Insert
figure 3
here

An example where ϕ and the price realization are such that both know there is an arbitrage, member 2 knows that 1 knows it, but 1 does not know that 2 knows is illustrated in Figure 4.

Insert
figure 4
here

The final example (Figure 5) illustrates how much larger the set of admissible price functions can be, compared to the standard model. Even at a price realization like the one given by the bold dot it is easy to verify that it is not common knowledge that there is an arbitrage (i.e. 1 knows, 2 knows, 1 knows that 2 knows, 2 knows that 1 knows, 1 knows that 2 knows that 1 knows etc, but not ad infinitum). Thus there is some probability of miscoordination which guarantees that their portfolio problem is well-defined.

Insert
figure 5
here

6.2 Domain of the Pricing Function

In order to summarize the previous discussion a bit more formally, we define \mathcal{G} to be the powerset of A , the set of all subsets of A . Also, define the set of realizations ϵ that yield an arbitrage opportunity that can be taken advantage of by some set of traders $G \in \mathcal{G}$ having the information of investor h by

$$E^{h,G}(\phi) \equiv \left\{ \epsilon \in \mathcal{E} : \text{it is } \{\sigma(g^h(\epsilon), \phi_a(\epsilon))\}_{a \in G} \text{ - common-knowledge that } \right. \\ \left. \exists \lambda \in \mathbb{R}_{++}^S \text{ such that } \phi_a(\epsilon) = \sum_{s \in S} d_{a,s} \lambda_s \text{ for all } a \in G \right\}$$

By an event being $\{\sigma(g^h(\epsilon), \phi_a(\epsilon))\}_{a \in G}$ -CK we mean that the event is CK for the subset of investor h 's team members $a \in G$, given that the sigma-algebra of (h, a) is given by $\sigma(g^h(\epsilon), \phi_a(\epsilon))$, all $a \in G$.

A pricing function then needs to belong to

$$\Phi \equiv \left\{ \phi : \mathcal{E} \rightarrow \mathbb{R}^A, \mathcal{F}/\mathcal{B}^A\text{-measurable, such that } \mu \left(\bigcup_{h \in H, G \in \mathcal{G}} E^{h,G}(\phi) \right) = 0 \right\}$$

Notice that this formulation also eliminates single-agent arbitrages in which an agent (h, a) knows there is an arbitrage that he can exploit on his

own. Indeed, if $G = \{a\}$, then realizations ϵ such that there is no $\lambda \in \mathbb{R}_{++}^S$ for which $\phi_a(\epsilon) = \sum_s \lambda_s d_{a,s}$ have zero measure if $\phi \in \Phi$. For instance, a bond with a sure payoff of 1 in each state cannot have a non-positive price.

The following proposition is then an easy consequence of Proposition 1:

Proposition 2 *For $\phi \in \Phi$, demand functions are well-defined.*

There are no general equilibrium existence results for standard rational expectations economies, and there are a number of counterexamples to existence (e.g. (Green 1977)). Our aim is not to extend on the literature on existence. Still, we would briefly like to point out that, just as in (Allen 1985a), equilibria exist in our model if we allow for what may be called a fuzzy price economy, where the actual transacting price is a properly perturbed version of the “true equilibrium price.” In general, however, the frictions we introduce in this paper exacerbate the difficulties of establishing the behaviour of demand functions, the domain of viable prices and the existence of equilibria. Therefore, in the remainder of this paper, we shall assume that an equilibrium exists and study some of its qualitative properties.

6.3 Arbitrage

The joint assumption of immediacy risk and of asset-specificity may generate arbitrage opportunities with positive probability, depending on (\mathcal{E}, ϕ, R) .⁶ Since in our economy prices are determined locally, $q_a = d_a \cdot \lambda^a$, $a \in A$, there is no guarantee that there is a $\lambda \gg 0$ for which $q_a = \phi_a(\epsilon) = d_a \cdot \lambda$, $\forall a \in A$. As a function of ϵ , prices of different assets are determined by different state-prices and depend on local information, and there is no reason why the resulting price vector q should not admit free lunches. This is clearest if we assume that the state of information ϵ_a only affects demand for asset a , in which case $q_a = d_a \cdot \lambda^a(\epsilon_a)$. As long as the fact that q represents an arbitrage is not common knowledge for any subset of agents (h, A') , $A' \subset A$, there is no force driving prices into the set of no-arbitrage prices.

The ingredients needed to induce $\phi(\mathcal{E}) \cap \mathcal{Q} \neq \emptyset$ are threefold. First, ϕ must be well-behaved. In particular, ϕ must not map all points $\epsilon \in \mathcal{E}$ into “too small a set.” Second, it is obvious that \mathcal{E} needs to be “large enough,”

⁶Of course, in the fuzzy price economy, the probability of arbitrage depends on the exogenous specification of the support of the fuzzy prices. If the support is all of \mathbb{R}^A , arbitrage price vectors trivially occur in a fuzzy price equilibrium.

else the equilibria will be too close to the Radner REE (if ϕ is well-behaved), which cannot allow for any arbitrage. Lastly, we have to allow for asset payoffs that are general enough so that an arbitrage can be constructed at all. This latter requirement is guaranteed by H4 which warrants that there is a portfolio y such that $Ry \gg 0$.

If these conditions are satisfied, it becomes apparent that arbitrage is possible, for we have enough degrees of freedom to choose ϵ for which $\{\lambda \in \mathbb{R}^S : R'\lambda = \phi(\epsilon)\} \cap \mathbb{R}_{++}^S = \emptyset$. A simple example (taken from (Zigrand 2001)) may be instructive.

Example 1 *Consider the CAPM economy (quadratic Von Neumann - Morgenstern utility functions over time two consumption and linear time zero utility function over time one consumption). Assume also that $\epsilon \in \mathbb{R}^A$ represents a stochastic asset supply, assumed to be multivariate normal $N(0, \Sigma_\epsilon)$.*

Because the support of ϵ is all of the real space, $\mathcal{E} = \mathbb{R}^A$, this specification simplifies the equilibrium existence problem. There cannot be any situation where it is common knowledge to some subset of traders that there is an arbitrage opportunity, $\mu(\cup_{h \in H, G \in \mathcal{G}} E^{h,G}(\phi)) = 0$, and the domain for ϕ is the set of all $\mathcal{F}/\mathcal{B}^A$ -measurable functions.

Then it can be shown that the REE pricing function ϕ is an isomorphism, $\phi(\epsilon) = F + G\epsilon$, with parameters $F \in \mathbb{R}^A$ and G a diagonal and invertible $A \times A$ matrix. Arbitrage may therefore occur with strictly positive μ -probability (a simple proof is in the appendix).

Interestingly, if ϕ is bijective as in this example, the equilibrium price vector is fully revealing the state of information. In that case the equilibrium is fully efficient and Pareto optimal in standard economies. Nevertheless, even in that benchmark economy, equilibria may exhibit arbitrage opportunities and allocational inefficiencies under the maintained plausible assumptions of asset-specificity and immediacy risk. The next section shows that equilibria are characterized by informational inefficiency as well.

6.4 Informational Efficiency

Here we analyze the information transmission across exchanges or across pits at these equilibria. If each pit is hit by an independent demand or supply shock and if no investors can trade on more than one pit, no information can be transmitted. In our setup, however, prices are correlated so that each

observed asset price reveals information about all the states of information. Still, at a given typical realization (q_b, e^h) , markets remain inefficient:

Proposition 3 (Degree of Information Revelation) *Assume that the equilibrium price function ϕ is smooth. Define (h, b) 's observation function $i^{h,b} : \mathcal{E} \rightarrow \mathcal{D} \equiv \mathbb{R} \times \mathcal{P}$ by $i^{h,b} : \epsilon \mapsto (\phi_b, g^h)(\epsilon)$. It follows from H3 that $i^{h,b}$ is a smooth mapping between smooth manifolds. Then for a dense set of prices and individual characteristics (q_b, e^h) , denoted by Z , whose complement is of Lebesgue measure zero in \mathbb{R}^{P+1} , the dimension of residual uncertainty facing team member (h, b) is $E - P - 1$.*

This result can be strengthened when more is known on the observation function:

Example 1 (contd) *In this economy the observation functions $i^{h,b} = \phi_b$ are submersions. Since submersions do not admit any critical points, the previous proposition is actually true for any state of information. In other words, the dimension of residual uncertainty facing team member (h, b) is $E - 1$ with μ -probability one.*

7 Conclusion

This paper proposes a plausible framework in which to address general equilibrium pricing implications.

First, we show that under standard assumptions, some micro market structure aspects (if taken one at a time) are not crucial because equilibria coincide with standard REE.

Second, we analyze a class of models in which the marketstructure does matter. We do this by combining two plausible and appealing conditions on how orders are computed, submitted and executed. The implication is that a piece of information on one market needs not be immediately reflected in all other prices. These marketstructure assumptions dispense with noise trading as a means to avoid complete information revelation and thus help avoid the Grossman-Stiglitz paradox (see (Grossman and Stiglitz 1980)). While agents would have been perfectly informed from prices in a standard economy, calling for a blurring device in order to generate imperfect information and an incentive to gather and pay for informative signals, no such device is necessary here.

Third, we show that given the imperfect knowledge of any one trader about the other traders' deep parameters, equilibria may exhibit mispricings, and therefore arbitrage opportunities.

The proposed marketstructure assumptions break perfect risk-sharing, enable investors to gain from better information and allow us to view markets much more as the decentralized places they are, as opposed to the idealized frictionless simultaneous worldwide auction envisaged by Walras. We believe that the plausible framework introduced in this paper therefore lends itself to the study of a number of more applied financial problems. Some applications of the setup proposed in this paper to asset pricing can be found in (Zigrand 2001).

One final application of the framework is to view it as a mechanism to generate arbitrage opportunities, and to explicitly account for arbitrageurs. Along the line of (Zigrand 1999), arbitrageurs can be fruitfully modelled as strategic Cournot competitors, taking as given the Walrasian correspondence generated in this model. In other words, by paying fixed entry fees, arbitrageurs acquired a technology that allows them to trade without immediacy risk. With costly entry into the arbitraging business, the Cournot-Walras equilibrium characterizes the number of active players, the degree of integration as well as the degree of informational efficiency since arbitrageurs effectively make markets more revealing.

Appendix: Proofs

Proof of Lemma 1 At a standard REE, investors observe the price realization $q = \phi(\epsilon)$ and solve their asset allocation problem

$$\theta^{h,*} = \arg \max_{\theta \in \mathbb{R}^A} E[U^h \mid \sigma(g^h, \phi)] \quad [R]$$

Since each optimal portfolio $\theta^{h,*}$ is measurable w.r.t. $\sigma(g^h, \phi)$, this generates a Baire function (Borel-measurable function $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$) $f^{h,*}$ satisfying $\theta^{h,*} = f^{h,*}(q, e^h)$ (refer for instance to (Krickeberg 1965), Theorem 2.5 on p.139).

On the other hand, assume the same standard investor optimization problem, but with the added feature of price uncertainty at the time zero optimization stage. Recall that investors have common prior beliefs (summarized by the measure μ) over states of information ϵ . They choose, given their information $\mathcal{F}^h = \sigma(g^h)$, an (version of the) optimal Baire function $f^h : \mathbb{R}^A \rightarrow \mathbb{R}^A$ s.t.

$$f^h = \arg \max_{F: \mathbb{R}^A \rightarrow \mathbb{R}^A} E \left[u_0^h (\omega_0^h(\epsilon) - \phi(\epsilon) \cdot F(\phi(\epsilon)); \epsilon) + \sum_s p_s u_s^h (d_s \cdot F(\phi(\epsilon)) + \omega_s^h(\epsilon); \epsilon) \mid \sigma(g^h) \right] \quad [IR]$$

First, we show that the argmax f^h of the immediacy risk problem [IR] also solves the Radner problem [R]. Assume to the contrary that, for $q \in W$ where the set W is measurable with respect to $\sigma(g^h)$ with positive measure (conditional on $\sigma(g^h)$), $f^h(q) \neq f^{h,*}(q)$. Then on W we have that

$$E[U^h(f^{h,*}) \mid \sigma(g^h, \phi)] > E[U^h(f^h) \mid \sigma(g^h, \phi)]$$

The function F defined as $F = f^h$ if $q \notin W$ and $F = f^{h,*}(q)$ if $q \in W$ then satisfies

$$E[U^h(F) \mid \sigma(g^h)] > E[U^h(f^h) \mid \sigma(g^h)]$$

contradicting the maximality of f^h .

Second, we show that the function $f^{h,*}$ also solves [IR]. Since there is no ex-ante link between the events represented by the sets in $\sigma(g^h, \phi)$, the investor can optimize event by event, and the derived demand functions are the standard Marshallian demand functions. In other words, since

$E[U^h|\sigma(g^h)] = E[E[U^h|\sigma(g^h, \phi)]|\sigma(g^h)]$, maximizing $E[U^h|\sigma(g^h, \phi)]$ automatically maximizes the entire expression.

It follows that the resulting equilibrium price vector will indeed be in the set of no-arbitrage prices, which we denote by \mathcal{Q} (defined in Section 6), and coincides with the REE equilibrium of the economy without immediacy risk (ignoring issues arising from multiple and sunspot equilibria). In other words, both interpretations of REE that we advanced above are formally equivalent. ■

Proof of Proposition 1 (Existence of Demand Functions) The proof goes along similar lines as the proof of 1. The fixed point problem can be transformed into an easier problem as follows. Rather than maximize the interim objective, given $\mathcal{F}^{h,a}$, we maximize the ex-ante objective [P.ea.], i.e. given \mathcal{F}^h , by choosing (for a given e^h) Baire functions $f_a^h : \mathbb{R} \rightarrow \mathbb{R}$ for all $a \in A$ (which we could again denote by $f_a^h(q_a, e^h)[\phi]$):

$$V^h(g^h)[\phi] = \max_{\{f_a^h: \mathbb{R} \rightarrow \mathbb{R}\}_{a=1}^A} E[U^h(x^h(f^h))|\mathcal{F}^h] \quad [\text{P.ea.}]$$

where the budget constraints are used in the functions

$$(x^h \circ f^h) \circ q = \begin{bmatrix} (x_0^h \circ f^h) \circ q \\ \vdots \\ (x_s^h \circ f^h) \circ q \\ \vdots \end{bmatrix} = \begin{bmatrix} \omega_0^h - \sum_a f_a^h(q_a)q_a \\ \vdots \\ \omega_s^h + \sum_a f_a^h(q_a)d_{a,s} \\ \vdots \end{bmatrix}$$

Measurability is then guaranteed: for two realizations q and q' of the random variable ϕ , $f_a^h(q_a) = f_a^h(q'_a)$ if $q_a = q'_a$. If this standard problem admits a solution, this same solution also has to be a solution to each team member's problem. Indeed, let f^* represent the solution to the ex-ante problem, and let $\theta_b^h(q_b) \in \mathbb{R}$ represent the solution to the interim (given q_b) problem of

$$\max_{\theta_b^h \in \mathbb{R}} E[U^h|\mathcal{F}^{h,b}]$$

Keeping in mind that $E[U^h|\mathcal{F}^h] = E[E[U^h|\mathcal{F}^{h,b}]|\mathcal{F}^h]$, assume that $f_b^*(q_b)$ does not maximize the interim problem (for a nonnull set of prices $q_b \in A$).

Then setting $f_b(q_b) = \theta_b^h(q_b)$ for $q_b \in A$ and $f_b(q_b) = f_b^*(q_b)$ for $q_b \in \mathbb{R} \setminus A$ yields a strictly higher ex-ante utility, contradicting the optimality of f^* .

Assume that the space of demand functions $f : \mathbb{R}^A \rightarrow \mathbb{R}^A$ (but where each of the A functions only maps \mathbb{R} to \mathbb{R}) is the normed vector space \mathcal{L}_1^A of measurable functions with the norm

$$\|f\| \equiv E[|f| \|\mathcal{F}^h\|] = E \left[\left(\sum_{a=1}^A f_a(q_a)^2 \right)^{1/2} \|\mathcal{F}^h\| \right]$$

Obviously, for certain pricing functions, the optimization problem may not admit a solution unless we restrict the domain of choice. We are thus led to truncate the investment opportunity set by insisting that investors choose $(f, x) \in K \cap \mathcal{B} \equiv (K' \times K'') \cap \mathcal{B}$, K' and K'' compact rectangles in \mathcal{L}_1^A and \mathcal{L}_1^{S+1} respectively, with center at the origin, and \mathcal{B} the budget set. Here the object of choice x is also a Baire function, $x : \mathbb{R}^A \rightarrow \mathbb{R}^{S+1}$, for given e^h and ϕ . The constraint set is compact and convex. The demand functions solve the truncated problem:

$$\max_{(f,x) \in K \cap \mathcal{B}} W(x) \equiv E[U^h(x) \|\mathcal{F}^h\|]$$

The objective function W is continuous and strictly concave. To see continuity, $\lim_{n \rightarrow \infty} E[U^h(x_n) \|\mathcal{F}^h\|] = E[U^h(\lim_{n \rightarrow \infty} x_n) \|\mathcal{F}^h\|]$, notice that $|U^h(x_n)| \leq w \equiv \sup_{\tilde{x} \in K''} U^h(\tilde{x}) \leq M_U < \infty$, and that w is evidently integrable. So by the Lebesgue convergence theorem,

$$\lim_{n \rightarrow \infty} E[U^h(x_n) \|\mathcal{F}^h\|] = E[\lim_{n \rightarrow \infty} U^h(x_n) \|\mathcal{F}^h\|] = E[U^h(\lim_{n \rightarrow \infty} x_n) \|\mathcal{F}^h\|]$$

where the last equality follows from the assumed continuity of $U^h(\cdot)$. Strict concavity is evident: $W(\lambda x' + (1 - \lambda)x'') = E[U^h(\lambda x' + (1 - \lambda)x'') \|\mathcal{F}^h\|] < E[\lambda U^h(x') + (1 - \lambda)U^h(x'') \|\mathcal{F}^h\|] = \lambda W(x') + (1 - \lambda)W(x'')$.

Now since K is compact and W continuous, as functions of (e^h, ϕ) , (f, x) is nonempty and convex-valued, and by the strict concavity of W and the convexity of the constraint set, (f, x) is single-valued. ■

Proof that Arbitrage can Occur with Positive Probability in Example 1 Consider the set $Y \equiv \{y \in \mathbb{R}^A : Ry \gg 0\}$, nonempty by H3. It is sufficient to show that the set $P^* \equiv \{q \in \mathbb{R}^A : \exists y \in Y : q'y \leq 0\}$ has

strictly positive derived measure over asset prices, $\eta \equiv \mu \circ \phi^{-1}$. Define the set $P \equiv -Y$. We show first that $P \subset \text{int}(P^*)$, and second that $\eta(P) > 0$.

P only contains arbitrage prices: take any $q \in P$ and let the asset demand be $y = -q \in Y$. Then since $y \in Y$, $Ry \gg 0$, and since $y = -q$, $q'y = -q'q < 0$: $P \subset \text{int}(P^*)$.

Upon inspection, it is clear that the linear mapping defined by R is transversal to \mathbb{R}_{++}^S . It follows that Y is a manifold of dimension A and has positive A -Lebesgue measure, and so does P . Notice that $\phi(\mathcal{E}) \cap P \neq \emptyset$. Since the pricing function ϕ is an isomorphism, the preimage of P in \mathcal{E} via ϕ has positive μ -measure as well. ■

Proof of Proposition 3. Similar to (Allen 1985b) and (Heifetz and Polemarchakis 1998). By Sard's Theorem ((Guillemin and Pollack 1974)), the set of critical values of $i^{h,b}$ is of Lebesgue measure zero and its complement Z is dense. By the regular value theorem, $i^{h,b^{-1}}(q_b, e^h)$ is an $E - P - 1$ dimensional smooth manifold for every $(q_b, e^h) \in Z$. ■

References

- Allen, B. (1981). “Generic Existence of Completely Revealing Equilibria for Economies with Uncertainty when Prices Convey Information”. *Econometrica* 49, 1173–1199.
- Allen, B. (1985a). “The Existence of Fully Rational Expectations Approximate Equilibria with Noisy Price Expectations”. *Journal of Economic Theory* 37, 213–253.
- Allen, B. (1985b). “The Existence of Rational Expectations Equilibria in a Large Economy with Noisy Price Expectations”. *Journal of Mathematical Economics* 24, 67–103.
- Ausubel, L. (1990). “Partially Revealing Rational Expectations Equilibrium in a Competitive Economy”. *Journal of Economic Theory* 50, 93–126.
- Cotter, K. (1985). “Similarity of Information and Behavior with a Pointwise Convergence Topology”. *Journal of Mathematical Economics* 15, 25–38.
- Diamond, D. and R. Verecchia (1981). “Information Aggregation in a Noisy Rational Expectations Economy”. *Journal of Financial Economics* 9, 221–235.
- Dow, J. and G. Gorton (1994). “Arbitrage Chains”. *Journal of Finance* 49(3), 819–849.
- Fremault, A. (1993). “Execution Lags and Imperfect Arbitrage: the Case of Stock Index Arbitrage”. *Economics Letters* 43, 103–109.
- Geanakoplos, J. and H. Polemarchakis (1986). “Existence, Regularity, and Constrained Suboptimality of Competitive Allocations When the Asset Market is Incomplete”. In R. S. Walter Heller and D. Starrett (Eds.), *Uncertainty, Information and Communication. Essays in Honor of Kenneth J. Arrow, Volume III*. Cambridge University Press.
- Green, J. (1977). “The Non-Existence of Informational Equilibria”. *Review of Economic Studies* 44, 451–463.
- Grossman, S. (1977). “The Existence of Futures Markets, Noisy Rational Expectations and Informational Externalities”. *Review of Economic Studies* 44(3), 431–449.

- Grossman, S. and J. Stiglitz (1980). “On the Impossibility of Informationally Efficient Markets”. *American Economic Review* 70, 393–408.
- Guillemin, V. and A. Pollack (1974). *Differential Topology*. Prentice Hall.
- Heifetz, A. and H. Polemarchakis (1998). “Partial Revelation with Rational Expectations”. *Journal of Economic Theory* 80, 171–181.
- Holden, C. (1995). “Index Arbitrage as Cross-Sectional Market Making”. *Journal of Futures Markets* 15(4), 423–455.
- Jordan, J. and R. Radner (1982). “Rational Expectations in Microeconomic Models: An Overview”. *Journal of Economic Theory* 26, 201–223.
- Klemperer, P. and M. Meyer (1989). “Supply Function Equilibria in Oligopoly Under Uncertainty”. *Econometrica* 57, 1243–1277.
- Kreps, D. (1977). “A Note on ‘Fulfilled Expectations’ Equilibria”. *Journal of Economic Theory* 14, 32–43.
- Krickeberg, K. (1965). *Probability Theory*. Addison-Wesley.
- Kyle, A. (1985). “Continuous Auctions and Insider Trading”. *Econometrica* 53, 1315–1336.
- Kyle, A. (1989). “Informed Speculation With Imperfect Competition”. *Review of Economic Studies* 56, 317–356.
- Marschak, J. and R. Radner (1971). *The Economic Theory of Teams*. Yale University Press.
- Radner, R. (1979). “Rational Expectations Equilibrium: Generic Existence and the Information Revealed by Price”. *Econometrica* 47, 655–678.
- Radner, R. (1986a). “Allocation of a Scarce Resource under Uncertainty: An Example of a Team”. In C. B. McGuire and R. Radner (Eds.), *Decision and Organization. Essays in Honor of Jacob Marschak* pp. 217–236. University of Minnesota Press.
- Radner, R. (1986b). “Teams”. In C. B. McGuire and R. Radner (Eds.), *Decision and Organization. Essays in Honor of Jacob Marschak* pp. 189–215. University of Minnesota Press.
- Shleifer, A. and R. Vishny (1997). “The Limits of Arbitrage”. *Journal of Finance* 52, 35–55.

- Zigrand, J.-P. (1999). “Arbitrage and Endogenous Market Integration”. Technical Report 319, FMG Discussion Paper.
- Zigrand, J.-P. (2001). “Rational Asset Pricing Implications from Realistic Trading Frictions”. Mimeo, The London School of Economics.




Figure 1: ϕ is not admissible. It is common knowledge that there is an arbitrage opportunity.




Figure 2: ϕ is admissible. 1 knows there is an arbitrage, but 2 doesn't.




Figure 3: ϕ is admissible. Both know there is an arbitrage, but no agent knows that the other agent knows.




Figure 4: ϕ is admissible. Both know there is an arbitrage, 2 knows that 1 knows, but 1 does not know whether 2 knows.




Figure 5: ϕ is admissible. Both know there is an arbitrage, and each one knows that the other one knows, and 1 knows that 2 knows that 1 knows etc., but not ad infinitum.