Learning and stability of the Bayesian–Walrasian equilibrium

Marialaura Pesce [a], Nicholas C. Yannelis [b,c,*]

[a] *Dipartimento di Matematica e Statistica, Universitá di Napoli Federico II, Napoli 81026, Italy*
[b] *Economics - School of Social Sciences, The University of Manchester, Oxford Road, Manchester M13 9PL, UK*
[c] *Department of Economics, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA*

ARTICLE INFO

ABSTRACT

Under the Bayesian–Walrasian Equilibrium (BWE) (see Balder and Yannelis, 2009), agents form price estimates based on their own private information, and in terms of those prices they can formulate estimated budget sets. Then, based on his/her own private information, each agent maximizes interim expected utility subject to his/her own estimated budget set. From the imprecision due to the price estimation it follows that the resulting equilibrium allocation may not clear the markets for every state of nature, i.e., exact feasibility of allocations may not occur. This paper shows that if the economy is repeated from period to period and agents refine their private information by observing the past BWE, then in the limit all agents will obtain the same information and market clearing will be reached. The converse is also true. The analysis provides a new way of looking at the asymmetric equilibrium which has a statistical foundation.

© 2010 Elsevier B.V. All rights reserved.

1. Introduction

It is by now well known that a Rational Expectations Equilibrium (REE), as introduced in Radner (1979) may not exist. It only exists in a generic sense and not universal. Moreover, it fails to be fully Pareto optimal and incentive compatible and it is not implementable as a perfect Bayesian equilibrium of an extensive form game (Glycopantis et al., 2005, 2009). The difficulty with the rational expectations equilibrium is that agents are maximizing interim expected utility conditioned on their own private information and on the information that the equilibrium prices have generated, i.e., agents are acting as knowing the equilibrium prices.

Balder and Yannelis (2009) introduced a new concept called Bayesian–Walrasian Equilibrium (BWE) which has Bayesian features. In particular, agents form price estimates based on their own private information,[1] and in terms of those prices they can formulate estimated budget sets. Then, based on his/her own private information, each agent maximizes interim expected utility subject to his/her own estimated budget set. From the imprecision due to the price estimation it follows that the resulting equilibrium allocation may not clear the markets for every state of nature, i.e., exact feasibility of allocations may not occur. This new concept exists under the standard continuity and concavity assumptions on the utility function and furthermore, it was shown in Balder and Yannelis (2009) that in the example of the non-existence of a REE of Kreps (1977), a BWE exists, but markets do not clear exactly. The non-market clearing of the BWE, is of course not desirable as in equilibrium there may be excess supply or demand. However, Balder and Yannelis (2009) do show that in the case of perfect

---

foresight (correct price estimates) and also in the case of symmetric information, a BWE exists and the exact feasibility of allocations holds.

The main purpose of this paper is to provide a dynamic rationalization of the BWE and show that exact market clearing will always hold in the limit. In particular, we show that if the economy is repeated from period to period and agents refine their private information by observing the past BWE, then in the limit all agents will obtain the same information and market clearing will be reached.[2]

In other words, in the limit there is no asymmetric information among agents, and all markets are cleared for every state of nature. Furthermore, we show that the limit BWE always exists, it is Pareto optimal and obviously incentive compatible. Hence, the limit BWE has desirable properties, i.e., it exists under the standard assumptions and it is efficient. This is quite interesting because, even if in each period we have an excess demand or supply, repeating the asymmetric information economy, agents eventually will learn all the information (i.e., become symmetrically informed) and in that limit economy the BWE satisfies exact feasibility.

It is rather surprising that the converse is also true, i.e., starting from a BWE in the limit symmetric economy we can construct a sequence of approximate BWE or ϵ-BWE which converges to it. This means that even if agents make mistakes along the sequence by optimizing approximately their conditional expected utility subject to the estimated budget set, and even if market clearing does not hold in the sequence of the ϵ-BWE, still in the limit agents will reach the exact BWE, where exact optimization and exact feasibility hold. This result can be viewed as a stability property of the BWE.

The above results have several implications.

First, the BWE concept where agents are symmetrically informed and obviously all markets are cleared can be viewed as the limit of an asymmetric BWE where the asymmetrically informed agents maximize interim utility based on their own information subject to an estimated budget constraint and market clearing does not occur. Agents by observing past BWE learn the information of others and the asymmetries on the private information disappear in the limit.

Second, the above results shed light on some important issues regarding the effect of information. Indeed, as the learning process unravels near the limit the additional information acquired does not change drastically equilibrium outcomes, that is small perturbations of the information structure do not have drastic effects on the equilibrium outcomes.

Third, since the BWE is based on estimated prices and thus on estimated budget constraints, our new modeling provides a statistical foundation of our equilibrium concept which seems to be suitable to econometric applications.

Finally, we have provided a new way to examine the Walrasian equilibrium concept under uncertainty, which is not subjected to the criticism of the REE. Our new concept universally exists and it is efficient.

The paper is organized as follows: Section 2 introduces the reader to the BWE and its relationship to the REE. Section 3 discusses the dynamics and the idea of learning. Also in the same section we prove results on the convergence of private information. Section 4 introduces formally the limit symmetric BWE and Section 5 provides the conditions under which the BWE converges to the symmetric BWE. Section 6 shows that any symmetric BWE can be reached as the limit of an approximate sequence of BWE. Some concluding remarks are collected in Section 7.

2. Bayesian–Walrasian equilibria

We consider an **exchange economy with asymmetric information**

$$\mathcal{E} = \left\{ (\Omega, \mathcal{F}, \mathbb{P}); \quad (X_i, \mathcal{F}_i, u_i, e_i)_{i \in I = \{1, \dots, n\}} \right\},$$

where,

1. $(\Omega, \mathcal{F}, \mathbb{P})$ is a probability measure space describing the **exogenous uncertainty**, i.e., Ω is the countable set of all possible states of nature and $\mathcal{F}$ is the σ-field that represents the set of all events.
2. $(X_i, \mathcal{F}_i, u_i, e_i)_{i \in I}$ is the set of **agent's characteristics**. Each economic agent $i \in I$ is characterized by:
   - $X_i : \Omega \to 2^{\mathbb{R}^\ell_+}$ is agent i'**random consumption set**.
   - $\mathcal{F}_i$ is a measurable partition[3] of $(\Omega, \mathcal{F})$ denoting **the private information** of agent i. The interpretation is as usual: if $\omega \in \Omega$ is the state of nature that is going to be realized, agent i observes $E^{\mathcal{F}_i}(\omega)$ the element of $\mathcal{F}_i$ which contains ω.
   - a **random utility function** representing his preferences:
     $$u_i : \Omega \times \mathbb{R}^\ell_+ \to \mathbb{R}$$
     $$(\omega, x) \to u_i(\omega, x).$$

Through the paper, we will assume that for all $i \in I$ and $\omega \in \Omega$, $u_i(\omega, \cdot)$ is strongly monotone.

---

[2] Such a conjecture was made in Balder–Yannelis (2009, p. 396) and it was left as an open question if it is true. Theorem 6.2 in Section 6, provides an affirmative answer to the Balder–Yannelis conjecture.

[3] By an abuse of notation we will still denote by $\mathcal{F}_i$ the σ-algebra that the partition $\mathcal{F}_i$ generates.

- a **random initial endowment** of physical resources represented by the function

$$e_i : \Omega \to \mathbb{R}_+^\ell.$$

We assume that the function e_i is $\mathcal{F}_i$-measurable and $e_i(\omega) \in X_i(\omega)$ for all $\omega \in \Omega$.
- $\mathbb{P}$ is **the common prior**.

Suppose that for all $\omega \in \Omega$, $\mathbb{P}(\omega) > 0$. Upon the realization of $\omega \in \Omega$, which agent i perceives through $E^{\mathcal{F}_i}(\omega)$, he/she forms the **conditional probability** given by[4]

$$\mathbb{P}_i(\omega'|\omega) = \begin{cases} \dfrac{\mathbb{P}(\omega')}{\mathbb{P}(E^{\mathcal{F}_i}(\omega))} & \text{if } \omega' \in E^{\mathcal{F}_i}(\omega), \\ 0 & \text{otherwise.} \end{cases}$$

Denote by

$$\bar{L}_{X_i} = \{x_i : \Omega \to \mathbb{R}_+^\ell \quad \text{such that} \quad x_i(\omega) \in X_i(\omega) \text{ for all} \quad \omega \in \Omega\}$$
$$L_{X_i} = \{x_i \in \bar{L}_{X_i} \quad \text{such that} \quad x_i(\cdot) \text{ is } \mathcal{F}_i - \text{measurable}\}.$$

Let $\bar{L}_X = \prod_{i=1}^n \bar{L}_{X_i}$ and $L_X = \prod_{i=1}^n L_{X_i}$.

An **allocation** x for an economy $\mathcal{E}$ is a function

$$x : I \times \Omega \to \mathbb{R}_+^\ell$$
$$(i, \omega) \to x_i(\omega), \quad \text{such that} \quad x \in L_X.$$

An allocation x is said to be **feasible** if for all $\omega \in \Omega$,

$$\sum_{i \in I} x_i(\omega) = \sum_{i \in I} e_i(\omega).$$

Agent i's **interim expected utility** is the function $v_i : \Omega \times L_{X_i} \to \mathbb{R}$, defined by

$$v_i(x_i|\mathcal{F}_i)(\omega) = \sum_{\omega' \in \Omega} u_i(\omega', x_i(\omega')) \mathbb{P}_i(\omega'|\omega).$$

We can also express the interim expected utility using the conditional probability as

$$v_i(x_i|\mathcal{F}_i)(\omega) = \sum_{\omega' \in E^{\mathcal{F}_i}(\omega)} u_i(\omega', x_i(\omega')) \frac{\mathbb{P}(\omega')}{\mathbb{P}(E^{\mathcal{F}_i}(\omega))}.$$

A **random price vector** is a $\mathcal{F}$-measurable, non-zero function $p : \Omega \to \Delta$, where $\Delta = \{p \in \mathbb{R}_+^\ell : \sum_{j=1}^\ell p^j = 1\}$. Given a random price vector p, every agent i adopts the following conditional expectation:

$$\hat{p}_i(\omega) = \sum_{\omega' \in \Omega} p(\omega') \mathbb{P}_i(\omega'|\omega) = \sum_{\omega' \in E^{\mathcal{F}_i}(\omega)} p(\omega') \frac{\mathbb{P}(\omega')}{\mathbb{P}(E^{\mathcal{F}_i}(\omega))}. \tag{1}$$

In other words, $\hat{p}_i(\omega)$ is agent i's **Bayesian price estimate** of the random price vector p, given that the state ω has been realized.[5]

---

[4] Notice that $\sum_{\omega' \in E^{\mathcal{F}_i}(\omega)} \frac{\mathbb{P}(\omega')}{\mathbb{E}^{\mathcal{F}_i}(\omega)} = 1$.

[5] Since we consider countably many states of nature, then for each $i \in I$ and each good $h \in \{1, \ldots, \ell\}$, x_i^h and e_i^h are real sequences (i.e., $(x_i^h(\omega_1), x_i^h(\omega_2), \ldots)$ and $(e_i^h(\omega_1), e_i^h(\omega_2), \ldots)$). In the same spirit of Peleg and Yaari (1970), we do not restrict the choice of equilibrium prices to the dual space, but we look at them as functions defined from Ω to the $\mathbb{R}^\ell$-simplex, and hence for each good h, p^h is a real sequence as well as x_i^h and e_i^h are. Therefore, our approach is not exactly the infinite one, but it is closer to the finite-dimensional case. The reason is that we consider price-quantity inner products point wise $\hat{p}_i(\omega) \cdot x_i(\omega)$, where for each ω and each i, $\hat{p}_i(\omega)$ and $x_i(\omega)$ are elements of $\mathbb{R}_+^\ell$ (see for example the budget set below). We consider the topology of point wise convergence, that is the sequence x^t converges to x^* if and only if $\lim_{t \to \infty} x^t(\omega) = x^*(\omega)$ for all $\omega \in \Omega$. Similarly, a sequence of price p^t converges to p^* if and only if $\lim_{t \to \infty} p^t(\omega) = p^*(\omega)$ for all $\omega \in \Omega$. Moreover, notice that even if (1) is a countable sum, since for each ω, $p(\omega) \in \Delta$, then for all $i \in I$ and each $\omega \in \Omega$, $\hat{p}_i(\omega) \le 1$. We do not consider uncountably many states since we need a weaker topology and the convergence arguments may not hold. We leave it as an open question. However, it should be pointed out that recent results by Podczeck et al. (2010) show that with infinitely many states asymmetric information equilibrium need not exist.

Using this natural estimate for the price, given his/her information σ-algebra, agent i forms the following **estimated budget set**

$$\hat{B}_i(\omega, p) = \{x_i \in X_i(\omega) : \hat{p}_i(\omega) \cdot x_i \leq \hat{p}_i(\omega) \cdot e_i(\omega)\} = \left\{ x_i \in X_i(\omega) : \sum_{\omega' \in \Omega} p(\omega') \mathbb{P}_i(\omega'|\omega) \cdot x_i \leq \sum_{\omega' \in \Omega} p(\omega') \mathbb{P}_i(\omega'|\omega) \cdot e_i(\omega) \right\}.$$

We now extend definition of BWE, introduced by Balder and Yannelis (2009), to differential information economies with a countable set of states of nature.

Definition 2.1. An **Asymmetric Bayesian–Walrasian Equilibrium (BWE)** of the differential information exchange economy $\mathcal{E}$ is a pair (p, x) such that

(*i*) $p : \Omega \to \Delta$ is a random price vector,

(*ii*) $x = (x_i)_{i \in I} \in L_X$ is an allocation,

(*iii*) $x_i(\omega) \in \text{argmax}_{y_i \in \hat{B}_i(\omega, p)} v_i(y_i|\mathcal{F}_i)(\omega)$ for all $\omega \in \Omega$ and every $i \in I$,

(*iv*) $p(\omega) \cdot \sum_{i \in I} (x_i(\omega) - e_i(\omega)) = \max_{1 \leq h \leq \ell} \sum_{i \in I} (x_i(\omega) - e_i(\omega))_h$ for all $\omega \in \Omega$.

Condition (*iii*) states that for all state of nature agent i maximizes his/her interim expected utility based on his/her own private information subject to his/her estimated budget set. Condition (*iv*) is an unusual substitute for the classical feasibility property of Walrasian equilibria. As noted by Balder and Yannelis (see Proposition 2.1, p. 389), condition (*iv*) must be considered the price that has to be paid when one wishes to avoid the unrealistic hypothesis surrounding the REE. However, we obtain the exact feasibility, i.e.,

(*iv*) $\sum_{i \in I} x_i(\omega) = \sum_{i \in I} e_i(\omega)$ for all $\omega \in \Omega$,

provided that

(a) $\mathcal{F}_i = \mathcal{F}_j$ for all $i, j \in I$ (symmetric information)
(b) $\hat{p}_i = p$ for all $i \in I$ (correct price estimate or perfect forecasting of price p),

as the following proposition shows (see proposition 2.1 in Balder and Yannelis (2009) for the case of finite states of nature).

Proposition 2.2.   *If (p, x) is a BWE as in Definition 2.1, then free-disposal feasibility*

$$\sum_{i \in I} x_i(\omega) \leq \sum_{i \in I} e_i(\omega) \quad \text{for all } \omega \in \Omega$$

holds in each of the following cases:

(*a*) $\mathcal{F}_i = \mathcal{F}_j$ *for all $i, j \in I$*   (*symmetric information*)

(*b*) $\hat{p}_i = p$ *for all $i \in I$*   (*correct price estimate or perfect forecasting of price p*).

Moreover, if in additional $u_i(\omega, \cdot)$ is strongly monotonic on $X_i(\omega)$ for all $i \in I$ and for all $\omega \in \Omega$, then the above feasibility condition sharpens into

$$\sum_{i \in I} x_i(\omega) = \sum_{i \in I} e_i(\omega) \quad \text{for all} \quad \omega \in \Omega.$$

Proof.   Similar arguments to the ones used in Balder–Yannelis (2009, Proposition 2.1) can be adopted to complete the proof. $\square$

Thus, we are interested in finding a framework which guarantees that either condition (*a*) or condition (*b*) are satisfied, and hence in the BWE exact feasibility of allocations is obtained, as the definition below shows.

Definition 2.3. A pair (p^*, x^*) is said to be a **symmetric Bayesian–Walrasian equilibrium** if

(i^*) $p^* : \Omega \to \Delta$ is a random price vector,

(ii^*) $x^* = (x_i^*)_{i \in I} \in \bar{L}_X$ is an allocation $\mathcal{F}^*$ – measurable,

   where $\mathcal{F}^*$ is the symmetric information (i.e., for all $i \in I$, $\mathcal{F}_i = \mathcal{F}^*$)

(iii^*) $x_i^*(\omega) \in argmax_{y_i \in \hat{B}_i(\omega, p^*)} v_i(y_i | \mathcal{F}^*)(\omega)$ for all $\omega \in \Omega$ and every $i \in I$,

(iv^*) $\sum_{i \in I} x_i^*(\omega) = \sum_{i \in I} e_i(\omega)$ for all $\omega \in \Omega$.

We will consider a dynamic framework, in which agents learn by observing the past BWE and refining their private information. This allows us to get in the limit a symmetric information economy and under certain additional assumptions a full information economy (i.e., all agents know everything, i.e., $\mathcal{F}_i = \mathcal{F}$ for all $i \in I$). Therefore, we will prove not only that in the limit a BWE satisfies the exact feasibility, but also that it exists, it is (interim) Pareto optimal and obviously incentive compatible.

Let $\sigma(p)$ be the smallest sub-σ-algebra of $\mathcal{F}$ for which p is measurable and let $\mathcal{G}_i = \sigma(p) \vee \mathcal{F}_i$ denote the smallest σ-algebra[6] containing both $\sigma(p)$ and $\mathcal{F}_i$. The notion below is due to Radner (1979), (see also Allen (1981)).

Definition 2.4. A pair (p, x), where p is a price system and $x = (x_1, \ldots, x_n) \in \bar{L}_X$ is an allocation, is a **rational expectations equilibrium** (REE) for the economy $\mathcal{E}$, denoted by REE($\mathcal{E}$) if

 (i) for all i the consumption function $x_i(\cdot)$ is $\mathcal{G}_i$-measurable.
 (ii) for all i and for all ω the consumption function maximizes

$$v_i(x_i | \mathcal{G}_i)(\omega) = \sum_{\omega' \in E^{\mathcal{G}_i}(\omega)} u_i(\omega', x_i(\omega')) \frac{\mathbb{P}(\omega')}{\mathbb{P}(E^{\mathcal{G}_i}(\omega))},$$

   (where $E^{\mathcal{G}_i}(\omega)$ is the event in $\mathcal{G}_i$ which contains ω and $\mathbb{P}(E^{\mathcal{G}_i}(\omega)) > 0$) subject to

$$p(\omega) \cdot x_i(\omega) \le p(\omega) \cdot e_i(\omega)$$

   i.e. the budget set at state ω, and
(iii) $\sum_{i=1}^n x_i(\omega) = \sum_{i=1}^n e_i(\omega)$ for all ω.

The REE and BWE are different concepts. For example, as it is shown in Balder–Yannelis (2009, p. 390), in the Kreps example the REE does not exist, but the BWE exists. However, in the case that the private information of each agent in a BWE is singletons and if the REE is full revealing,[7] then both notions become an ex-post Walrasian equilibrium.

3. Learning and convergence of private information

Let $T = \{1, 2, \ldots\}$ be the discrete set of time horizon and $\mathcal{E}^t$ be a differential information economy with initial endowment e^t for each $t \in T$. Denote by $\sigma(X_i, u_i, e_i^1)$ the σ-algebra generated by random utility function, the random initial endowment, and the random strategy set of agent i at beginning, that is $\mathcal{F}_i^1 = \sigma(X_i, u_i, e_i^1)$. At time t, there is additional information available to agent i acquired by observing past BWE prices and allocations. We can express each agent's private information recursively by

$$\mathcal{F}_i^{t+1} = \mathcal{F}_i^t \vee \sigma(p^t, x^t),$$

where $\sigma(p^t, x^t)$ is the information that the BWE generates at period t, i.e., the smallest σ-algebra for which the Bayesian–Walrasian equilibrium in period t is measurable, and $\mathcal{F}_i^t \vee \sigma(p^t, x^t)$ is the join (i.e., the smallest σ-algebra containing

---

[6] The field of events discernable by every player is the **"coarse"** σ-field $\bigwedge_{i \in I} \mathcal{F}_i$, which is the largest σ-algebra contained in each $\mathcal{F}_i$. While, agents by pooling their information, discern the events in the **"fine"** σ-field $\bigvee_{i \in I} \mathcal{F}_i$, which denotes the smallest σ-algebra containing all $\mathcal{F}_i$.

[7] A REE price p is said to be full revealing if the information it generates is singletons, that is $\sigma(p) = \mathcal{F}$, and hence for each agent i, $\mathcal{G}_i = \mathcal{F}$.

$\mathcal{F}_i^t$ and $\sigma(p^t, x^t)$). Thus, for all $i \in I$,

$$
\begin{aligned}
\text{if } t = 1 \quad \text{then,} \quad \mathcal{F}_i^1 &= \sigma(X_i, u_i, e_i^1) \\
\text{if } t = 2 \quad \text{then,} \quad \mathcal{F}_i^2 &= \mathcal{F}_i^1 \vee \sigma(x^1, p^1) \\
\text{if } t = 3 \quad \text{then,} \quad \mathcal{F}_i^3 &= \mathcal{F}_i^2 \vee \sigma(x^2, p^2) \\
&= \mathcal{F}_i^1 \vee \sigma(x^1, p^1) \vee \sigma(x^2, p^2) \\
&\vdots
\end{aligned}
$$

Therefore, the private information of agent i at time t is given by

$$
\mathcal{F}_i^t = \sigma(X_i, u_i, e_i^1, (p^{t-1}, x^{t-1}), (p^{t-2}, x^{t-2}), \dots, (p^{t-(t-1)}, x^{t-(t-1)})) = \mathcal{F}_i^1 \vee \bigvee_{k=1}^{t-1} \sigma(p^k, x^k),
$$

where $(p^{t-1}, x^{t-1}), (p^{t-2}, x^{t-2}), \dots, (p^{t-(t-1)}, x^{t-(t-1)})$ are past Bayesian–Walrasian equilibria. Clearly, for each agent i and each time period t we have

$$
\mathcal{F}_i^t \subseteq \mathcal{F}_i^{t+1} \subseteq \mathcal{F}_i^{t+2} \subseteq \cdots
$$

The above expression represents a **learning process** for player i and it generates a sequence of exchange economies $\mathcal{E}^t = \{(X_i, \mathcal{F}_i^t, u_i, e_i^t) : i = 1, \dots, n\}$.

Notice that our framework is not a dynamic model since we repeat the economy over time allowing each agent to refine his/her own private information. Thus, the utility function u_i as well as the random consumption set X_i do not change. However, since agents update their information, the σ-algebra $\mathcal{F}_i^t$ changes from period to period, and consequently, the interim expected utility function changes as well.

We suppose that at beginning (i.e., $t = 1$) there are at least two agents who are asymmetrically informed, where for each agent i, the private information $\mathcal{F}_i^1$ is the information generated by the initial endowment, the utility function and the consumption set, that is $\mathcal{F}_i^1 = \sigma(X_i, u_i, e_i^1)$. We call this assumption a **non-trivial differential information economy**.

Notice that agents do not form expectations over the entire horizon but only for the current period, i.e., each agent's interim expected utility is based on his/her current period private information. Obviously, since the private information set of each agent becomes finer over time, the interim expected utility of each agent is changing as well. The information gathered at a given time t will affect the outcome in periods $t + 1, t + 2, \dots$. Furthermore notice that for each agent $i \in I$ and each time $t \in T$,

$$
\mathcal{F}_i^t = \mathcal{F}_i^1 \vee \left(\bigvee_{k=1}^{t-1} \sigma(p^k, x^k) \right),
$$

that is for all i and for all t, $\mathcal{F}_i^t = \mathcal{F}_i^1 \vee G^t$, where for all $t \in T$, $G^t = \bigvee_{k=1}^{t-1} \sigma(x^k, p^k)$ is the information generated by all past Bayesian–Walrasian equilibria until time t.

Thus, the information of each player i, after t periods, has two components. The former is given by his/her initial private information $\mathcal{F}_i^1$ which contributes to the asymmetric part and the latter is given by the information generated by the past period BWE allocations which creates the common part of information since all agents see the same BWE outcome in each period. In the limit economy, the private information of each agent i, $\mathcal{F}_i^*$, is the initial private information of i, $\mathcal{F}_i^1$, together with the information generated by all past periods BWE, that is

$$
\mathcal{F}_i^* = \mathcal{F}_i^1 \vee \bigvee_{k=1}^{\infty} \sigma(x^k, p^k) = \mathcal{F}_i^1 \vee G^{\infty}.
$$

Notice that even in the limit economy, agents may have different private information. Consider, for example, the case in which in each period $t \in T$, the information generated by the equilibria is coarser than the information of each trader, that is $\sigma(x^t, p^t) \subseteq \mathcal{F}_i^t$ for all $i \in I$. This means that no agent learns anything, that is

$$
\mathcal{F}_i^{t+1} = \mathcal{F}_i^t \qquad \text{for all } i \in I \text{ and for all } t \in T.
$$

Since in each period agents learn nothing, even in the limit economy the information of each agent i is the initial one. Hence, it might be that in each period the same economy is replicated and therefore the equilibria remains the same. This makes the convergence results trivial, because a constant sequence of identical equilibria converges always to itself and vice versa. Clearly, we want to avoid this trivial case. In other words, we want the learning process to make sense. Formally, we need that at least one agent learns something in at least one period. This is guaranteed by the following assumption which we call **"non-trivial learning"**.

(A.0) $\quad \mathcal{F}_i^1 \subseteq G^{\infty} = \bigvee_{k=1}^{\infty} \sigma(x^k, p^k) \quad$ for all $i \in I$.

Assumption $(A.0)$ states that for each agent i, the pooled information generated by all past equilibria is at least as fine as agent i's initial information.

Remark 3.1. Notice that the "no-trivial learning assumption" does not mean that all agents will learn something, it only implies that at least one agent learns something in some period. Indeed, if nobody can learn, that is

$$\text{for all } i \in I \text{ and for all } t \in T \quad \mathcal{F}_i^t = \mathcal{F}_i^{t+1},$$

then for all $i \in I$, $\mathcal{F}_i^1 = \mathcal{F}_i^* = \mathcal{F}_i^1 \vee G^\infty$, and therefore by $(A.0)$, $\mathcal{F}_i^1 = G^\infty$. This means that already at the beginning, (i.e. $t = 1$) all agents have the same information, which is a contradiction to the assumption of the non-trivial differential information economy.

Remark 3.2. We notice that, under the non-trivial learning assumption $(A.0)$, as time goes to infinity each agent gets the same information, given by the information generated by all the past periods Bayesian–Walrasian equilibria, i.e., $G^\infty = \bigvee_{k=1}^\infty \sigma(p^k, x^k)$. In other words, the private information $\mathcal{F}_i^t$ of each player i converges to G^∞. This follows from the fact that in the limit, the information of each agent is given by

$$\mathcal{F}_i^* = \mathcal{F}_i^1 \vee G^\infty,$$

which is finer than G^∞. On the other hand, by the non-trivial learning assumption $(A.0)$, for all $i \in I$, $\mathcal{F}_i^1 \subseteq G^\infty$. Hence, for each agent i, $\mathcal{F}_i^*$ coincides with G^∞. This means that as time goes on, the common part of information becomes prevalent and in the limit all agents will have the same information generated by all the past periods BWE. Therefore, in the limit we reach an economy with symmetric information in which, if utility functions are strongly monotone, a BWE satisfies the exact feasibility as indicated in Section 2, condition (a). Notice that $(A.0)$ is also a necessary condition for $\mathcal{F}_i^* = G^\infty$ for all $i \in I$. In fact, if for all $i \in I$, $\mathcal{F}_i^* = G^\infty$, then $(A.0)$ holds.

Remark 3.3. If Ω is finite, since agents can not learn forever, from Remark 3.2 it follows that there exists a period $s \in T$ from which all agents will have the same information G^∞, that is, there exists $s \in T$ such that for all $i \in I$,

$$\mathcal{F}_i^t = G^\infty \quad \text{for all } t > s.$$

Remark 3.4. Notice that we have stated that in the limit we reach an economy in which all agents have the same information. However, this does not mean that in the limit agents will know everything (i.e., G^∞ may not be singletons). Indeed, it might be the case that $G^\infty \subset \mathcal{F}$. However, if there exists a period t in which the pooled information of all agents is the full information $\mathcal{F}$, then after that period, agents will learn everything. Formally,

$$\text{if there exists } k \in T \text{ such that} \quad \bigvee_{i \in I} \mathcal{F}_i^k = \mathcal{F}, \quad \text{then } G^\infty = \mathcal{F}. \tag{2}$$

Indeed, for all $i \in I$, $\mathcal{F}_i^k \subseteq G^\infty$ (by assumption $(A.0)$), thus

$$\mathcal{F} = \bigvee_{i \in I} \mathcal{F}_i^k \subseteq G^\infty \subseteq \mathcal{F},$$

hence $G^\infty = \mathcal{F}$.

Remark 3.5. Typically, it is assumed that in the starting period $t = 1$, $\bigvee_{i \in I} \mathcal{F}_i^1 = \mathcal{F}$. As noted above, this condition is a particular case of (2) (i.e., $k = 1$). This means that in the limit the information of each agent will be generated by the partition of only singletons. Therefore, in the limit the measurability assumption will play no role, since the full information BWE will be an ex-post Walrasian equilibrium. Our more general framework includes the above case, where $\bigvee_{i \in I} \mathcal{F}_i^1 = \mathcal{F}$, as a special case.[8]

4. Limit symmetric Bayesian–Walrasian equilibrium

We define the **limit symmetric information economy**

$$\mathcal{E}^* = \{(\Omega, \mathcal{F}, \mathbb{P}), (X_i, \mathcal{F}_i^*, u_i, e_i^*)_{i \in I}\}$$

as follows: $(\Omega, \mathcal{F}, \mathbb{P})$, X_i and u_i are defined as in Section 2. While,

$$\mathcal{F}_i^* = G^\infty = \bigvee_{k=1}^\infty \sigma(x^k, p^k) \qquad \text{for all } i = 1, \ldots, n$$

and e_i^* is the initial endowment of agent i which is assumed to be $\mathcal{F}_i^*$-measurable and such that $e_i^*(\omega) \in X_i(\omega)$ for all $\omega \in \Omega$. Obviously, all agents in this economy have the same private information generated by the Bayesian–Walrasian equilibria in all periods (i.e., $\mathcal{F}_i^* = G^\infty = \bigvee_{k=1}^\infty \sigma(x^k, p^k)$).

---

[8] As it was pointed out in Section 2, in the special case of the full revealing REE, the BWE and REE coincide as they lead to an ex-post Walrasian equilibrium.

Remark 4.1. Since in the limit all agents will get the same information (i.e., $\mathcal{F}_i^* = G^\infty$ for all $i \in I$), then they will have also the same Bayesian estimate price (i.e., $\hat{p}_i(\omega) = \hat{p}(\omega)$ for all $\omega \in \Omega$), and therefore, if the utility functions are strongly monotone, then the limit symmetric BWE satisfies the exact feasibility.

For each i, define the set[9]

$$L_{X_i}^\infty = \{x_i \in \bar{L}_{X_i} : \quad x_i(\cdot) \text{ is } G^\infty - \text{measurable}\}.$$

Let $L_X^\infty = \prod_{i=1}^n L_{X_i}^\infty$.
The BWE($\mathcal{E}^*$) is defined for the economy $\mathcal{E}^*$ analogously with the economy $\mathcal{E}^t$.

Definition 4.2. A pair (p^*, x^*) is said to be a **limit symmetric Bayesian–Walrasian equilibrium** of the limit symmetric information economy $\mathcal{E}^*$, denoted by BWE($\mathcal{E}^*$), if

(i^*) $p^* : \Omega \to \Delta$ is a random price vector,

(ii^*) $x^* \in L_X^\infty$,

(iii^*) $x_i^*(\omega) \in argmax_{y_i \in \hat{B}_i(\omega, p^*)} v_i(y_i | G^\infty)(\omega)$ for all i and for all ω.

(iv^*) $\sum_{i \in I} x_i^*(\omega) = \sum_{i \in I} e_i^*(\omega)$ for all $\omega \in \Omega$.

Condition (iii^*) indicates that agents maximize interim expected utility subject to a point wise budget set conditioned on an event which is an element of the pooled information that all past BWE prices and allocations have generated. Notice that this event is the same for each agent, as the private information of each agent $\mathcal{F}_i^*$ is the same for each individual, i.e., $\mathcal{F}_i^* = G^\infty$ for all $i \in I$. Condition (iv^*) is the standard feasibility condition.

Using Kreps's example, Balder and Yannelis have underlined that the BWE and REE are distinct concepts. It is also true that in general, the limit of the BWE equilibria may be far from the REE because the learning process might be "trapped" at a point, where not all information has been revealed by prices. However, in the case of "non-trivial learning" in the limit the learned information becomes prevalent. In particular, in this case, the equilibrium price p^* of the limit symmetric BWE generates a σ-algebra $\sigma(p^*)$ which is coarser than the common information G^∞, i.e., $\sigma(p^*) \subseteq G^\infty$. This allows us to observe that for all $\omega \in \Omega$ and $i \in I$,

$$\left[\sum_{\omega' \in E^{G^\infty}(\omega)} p^*(\omega') \frac{\mathbb{P}(\omega')}{\mathbb{P}\left(E^{G^\infty}(\omega)\right)} \right] \cdot [x_i^*(\omega) - e_i^*(\omega)] = p^*(\omega) \cdot [x_i^*(\omega) - e_i^*(\omega)].$$

Therefore in the case of non-trivial learning the limit symmetric BWE can be viewed as a REE. The convergence property of the REE is another issue and we plan to work on this problem in a future paper.

Although, in this economy all agents have the same information, it may be still the case that the pooled information over the entire horizon may not be singletons. Therefore, the G^∞ measurability condition of allocations makes sense. However, if the pooled information over the entire horizon is singletons, Definition 2.3 reduces to an ex-post Walrasian equilibrium.

5. Convergence of the Bayesian–Walrasian equilibria to the symmetric one

In the following section we will provide conditions under which any sequence $(p^t, x^t) \in BWE(\mathcal{E}^t)$ will have a subsequence that converges to $(p^*, x^*) \in BWE(\mathcal{E}^*)$ in the limit symmetric information economy $\mathcal{E}^*$. This result on convergence to a limit symmetric BWE suggests that, as the learning process unravels near the limit, the additional information acquired does not change drastically equilibrium outcomes. In order to appreciate the value of the next theorem, one should contemplate the consequence of its failure: it would mean that small perturbation of the information structure would have drastic effects on the equilibrium outcome, which would have implications on the robustness of the equilibrium concept.

The following assumptions will be needed for our results.

(A.1) For every $i \in I$, $X_i : \Omega \to 2^{\mathbb{R}_+^\ell}$ is a convex and closed-valued correspondence.[10]
(A.2) $u_i(\omega, \cdot)$ is continuous, strongly monotone and bounded for all $\omega \in \Omega$ and for each $i \in I$.
(A.3) For every $i \in I$, $t \in T$ and for all $\omega \in \Omega$, the initial endowment $e_i^t(\omega)$ is an interior point of $X_i(\omega)$.
(A.4) $\{e_i^t, \mathcal{F}_i^t\}_{t \in T}$ is a martingale.
(A.5) $\{\sum_{i=1}^n e_i^t, \wedge_{i=1}^n \mathcal{F}_i^t\}_{t \in T}$ is a martingale.

---

[9] Notice that since for all $i \in I$, $\mathcal{F}_i^t \subseteq \mathcal{F}_i^{t+1}$ for all t, then $L_{X^t} \subseteq L_X^\infty$ for all t.
[10] Notice that X_i is also a non-empty correspondence, since we assume that it contains the initial endowment, i.e., $e_i(\omega) \in X_i(\omega)$ for all $\omega \in \Omega$.

Remark 5.1. Endowments of agents are allowed to vary over time according to measurability requirements as the private information of each agent become finer and finer due to the learning process. However, this is just a generalization since in the easy case of constant initial endowments (i.e., $e_i^1 = e_i^t$ for all i and t) hypotheses (A.4) and (A.5) are satisfied. Indeed, for all i, $e_i^1(\cdot)$ is $\mathcal{F}_i^1$-measurable and since $\mathcal{F}_i^1 \subseteq \mathcal{F}_i^t$ then for all i and t $e_i^1(\cdot)$ is $\mathcal{F}_i^t$-measurable. Moreover, for all $i \in I$, $t \in T$ and $s = 1, 2, \ldots$

$$E\left[e_i^{t+s}|\mathcal{F}_i^t\right] = E\left[e_i^1|\mathcal{F}_i^t\right] = e_i^1 = e_i^t$$

$$E\left[\sum_{i \in I} e_i^{t+s}|\bigwedge_{i \in I}\mathcal{F}_i^t\right] = E\left[\sum_{i \in I}e_i^1|\bigwedge_{i \in I}\mathcal{F}_i^t\right] = \sum_{i \in I}e_i^1 = \sum_{i \in I}e_i^t.$$

This means that $\{e_i^1, \mathcal{F}_i^t\}_{t \in T}$ and $\{\sum_{i=1}^n e_i^1, \wedge_{i=1}^n \mathcal{F}_i^t\}_{t \in T}$ are martingales.

Remark 5.2. Notice that the above assumptions may not guarantee the existence of a BWE when Ω is countable. On the other hand, if Ω is finite and $u_i(\omega, \cdot)$ is concave for all i and ω, from assumptions (A.1), (A.2) and (A.3) it follows that in each period a BWE exists (see Theorem 4.1 in Balder and Yannelis (2009)).

Remark 5.3. From assumption (A.4) it follows that $e_i^t(\omega)$ converges to e_i^* (recall the Martingale convergence Theorem), and (A.5) implies that the total initial endowment also converges.[11]

It is easy to see that the following lemma, needed for the next theorem, holds true (for sake of completeness see the proof in Pesce and Yannelis (2009)).

Lemma 5.4. *Let $a_n^1, a_n^2, \ldots, a_n^d$ be d sequences that converge respectively to $a^1, a^2, \ldots, a^d$. Then*

$$\lim_{n \to +\infty} \max_{1 \le j \le d} a_n^j = \max_{1 \le j \le d} \lim_{n \to +\infty} a_n^j.$$

We are now ready to state the following Theorem.

Theorem 5.5. *Let $\{\mathcal{E}^t : t \in T\}$ be a sequence of differential information economies satisfying assumptions $(A.0) - (A.4)$ and let $(p^t, x^t) \in \mathrm{BWE}(\mathcal{E}^t)$. Then, there exists a subsequence (p^{t_n}, x^{t_n}) which converges to $(p^*, x^*) \in \mathrm{BWE}(\mathcal{E}^*)$.*

Proof. Let (p^t, x^t) be a sequence of BWE, we need to find a subsequence (p^{t_n}, x^{t_n}) that converges to (p^*, x^*). Let $Z = \{(p^t, x^t) : t \in T\}$ where $(p^t, x^t) \in \mathrm{BWE}(\mathcal{E}^t)$ for each t. Let L_Δ be the space of all $\mathcal{F}$-measurable functions $f : \Omega \to \Delta$. Define $L_{P,X} = L_\Delta \times L_X$. We may take the random consumption set of each agent to be $X_i(\omega) = [0, \sup_t \sum_{i=1}^n e_i^t(\omega)]$. Such a set is clearly compact, convex and nonempty, therefore L_X is compact. Since L_Δ is defined on a compact set Δ, it is also compact. Therefore by Tychonoff's theorem, their product $L_{P,X}$ is compact. Thus, the closure of the set Z, $\bar{Z}$, is compact, and hence there exists a subsequence $\{(p^{t_n}, x^{t_n}) \in L_{P,X} : n = 1, 2, \ldots\}$ which converges to (p^*, x^*).

We must show that (p^*, x^*) is a Bayesian–Walrasian equilibrium for the limit symmetric information economy $\mathcal{E}^*$, i.e., $(p^*, x^*) \in \mathrm{BWE}(\mathcal{E}^*)$. It is clear that p^* is a price system. Since for all t and for all i, $\mathcal{F}_i^t \subseteq G^\infty$ and x_i^t is $\mathcal{F}_i^t$-measurable, then x_i^t is G^∞-measurable for all t and i (recall (A.0)). Therefore, $x^* \in L_X^\infty$. Thus the conditions (i^*) and (ii^*) hold.

We want now to prove that condition (iii^*) also holds. Suppose otherwise that x^* violates (iii^*), then there exist an agent i, a state ω^*, and an allocation $y \in L_X^\infty$ such that

$$v_i(y_i|G^\infty)(\omega^*) > v_i(x_i^*|G^\infty)(\omega^*) \tag{3}$$

and

$$\hat{p}^*(\omega^*) \cdot y_i(\omega^*) \le \hat{p}^*(\omega^*) \cdot e_i^*(\omega^*).$$

Notice that if ω^* is such that $\hat{p}^*(\omega^*) \cdot y_i(\omega^*) = \hat{p}^*(\omega^*) \cdot e_i^*(\omega^*)$, then since the conditional expected utility is continuous, from assumption $(A.3)$, we can find $\tilde{y}_i$ in the neighborhood of y_i such that $v_i(\tilde{y}_i|G^\infty)(\omega^*) > v_i(x_i^*|G^\infty)(\omega^*)$ and $\hat{p}^*(\omega^*) \cdot \tilde{y}_i(\omega^*) < \hat{p}^*(\omega^*) \cdot e_i^*(\omega^*)$. Without loss of generality, let $\tilde{y}_i$ equal to y_i, then it follows from above that

$$\hat{p}^*(\omega^*) \cdot y_i(\omega^*) < \hat{p}^*(\omega^*) \cdot e_i^*(\omega^*). \tag{4}$$

For any i, $\mathcal{F}_i^{t_n}$ is an increasing sequence of private information, then for any time t, it is coarser than its limit G^∞, and thus $E^{G^\infty}(\omega^*) \subseteq E^{\mathcal{F}_i^{t_n}}(\omega^*)$. Let us define $y_i^{t_n}$ by

$$y_i^{t_n}(\omega) = \begin{cases} y_i(\omega) & \text{for} \quad \omega \in E^{G^\infty}(\omega^*) \\ x_i^{t_n}(\omega) & \text{for} \quad \omega \notin E^{G^\infty}(\omega^*). \end{cases}$$

---

[11] Similar assumptions were made in Koutsougeras and Yannelis (1999).

Then,

$$v_i(y_i^{t_n}|\mathcal{F}_i^{t_n})(\omega^*) = \sum_{\omega \in E^{\mathcal{F}_i^{t_n}}(\omega^*)} u_i(\omega, y_i^{t_n}(\omega)) \frac{\mathbb{P}(\omega)}{\mathbb{P}(E^{\mathcal{F}_i^{t_n}}(\omega^*))}$$

$$= \sum_{\omega \in E^{G^\infty}(\omega^*)} u_i(\omega, y_i^{t_n}(\omega)) \frac{\mathbb{P}(\omega)}{\mathbb{P}(E^{\mathcal{F}_i^{t_n}}(\omega^*))} + \sum_{\omega \in E^{\mathcal{F}_i^{t_n}}(\omega^*) \setminus E^{G^\infty}(\omega^*)} u_i(\omega, y_i^{t_n}(\omega)) \frac{\mathbb{P}(\omega)}{\mathbb{P}(E^{\mathcal{F}_i^{t_n}}(\omega^*))}$$

$$= \sum_{\omega \in E^{G^\infty}(\omega^*)} u_i(\omega, y_i(\omega)) \frac{\mathbb{P}(\omega)}{\mathbb{P}(E^{\mathcal{F}_i^{t_n}}(\omega^*))} + \sum_{\omega \in E^{\mathcal{F}_i^{t_n}}(\omega^*) \setminus E^{G^\infty}(\omega^*)} u_i(\omega, x_i^{t_n}(\omega)) \frac{\mathbb{P}(\omega)}{\mathbb{P}(E^{\mathcal{F}_i^{t_n}}(\omega^*))}$$

$$= v_i(x_i^{t_n}|\mathcal{F}_i^{t_n})(\omega^*) + \left[v_i(y_i|G^\infty)(\omega^*) - v_i(x_i^{t_n}|G^\infty)(\omega^*) \right] \frac{\mathbb{P}(E^{G^\infty}(\omega^*))}{\mathbb{P}(E^{\mathcal{F}_i^{t_n}}(\omega^*))}.$$

Since $v_i(y_i|G^\infty)(\omega^*) > v_i(x_i^*|G^\infty)(\omega^*)$ (see (3)) and $x_i^{t_n}$ converges to x_i^*, by the continuity of the conditional expected utility $v_i(\ \cdot\ |G^\infty)(\omega^*)$, there exists large enough t such that $v_i(y_i|G^\infty)(\omega^*) > v_i(x_i^{t_n}|G^\infty)(\omega^*)$. Therefore, we have

$$v_i(y_i^{t_n}|\mathcal{F}_i^{t_n})(\omega^*) > v_i(x_i^{t_n}|\mathcal{F}_i^{t_n})(\omega^*).$$

Then,

$$\hat{p}_i^{t_n}(\omega^*) \cdot y_i^{t_n}(\omega^*) > \hat{p}_i^{t_n}(\omega^*) \cdot e_i^{t_n}(\omega^*)$$

or equivalently by the definition of $y_i^{t_n}$ in the state ω^*,

$$\hat{p}_i^{t_n}(\omega^*) \cdot y_i(\omega^*) > \hat{p}_i^{t_n}(\omega^*) \cdot e_i^{t_n}(\omega^*). \tag{5}$$

Notice that, since $E^{\mathcal{F}_i^t}(\omega^*) \supseteq E^{\mathcal{F}_i^{t+1}}(\omega^*) \supseteq E^{G^\infty}(\omega^*)$ for all i and t, $\mathbb{P}(E^{\mathcal{F}_i^t}(\omega^*))$ converges to $\mathbb{P}(E^{G^\infty}(\omega^*))$, then $\mathbb{P}(\omega^*|\mathcal{F}_i^t)$ converges to $\mathbb{P}(\omega^*|G^\infty)$. Moreover, since p^t converges to p^*, then $\hat{p}_i^t(\omega^*)$ converges to $\hat{p}^*(\omega^*)$ for all $i \in I$. Therefore, assumption $(A.4)$ and the condition (5) imply that

$$\hat{p}^*(\omega^*) \cdot y_i(\omega^*) \geq \hat{p}^*(\omega^*) \cdot e_i^*(\omega^*),$$

which contradicts (4). Hence, also the third condition (iii^*) holds. To complete the proof, we must show that condition (iv^*) holds as well.

By the definition of BWE($\mathcal{E}^{t_n}$), for all $\omega \in \Omega$,

$$p^{t_n}(\omega) \cdot \sum_{i \in I} (x_i^{t_n}(\omega) - e_i^{t_n}(\omega)) = \max_{1 \leq h \leq \ell} \sum_{i \in I} (x_i^{t_n}(\omega) - e_i^{t_n}(\omega))_h,$$

by the continuity of the inner product and by Lemma 5.4 we have that

$$p^*(\omega) \cdot \sum_{i \in I} (x_i^*(\omega) - e_i^*(\omega)) = \max_{1 \leq h \leq \ell} \sum_{i \in I} (x_i^*(\omega) - e_i^*(\omega))_h.$$

Therefore, (p^*, x^*) must be a Bayesian–Walrasian equilibrium of the limit symmetric information economy and hence by Proposition 2.2, it follows that $\sum_{i=1}^n x_i^*(\omega) = \sum_{i=1}^n e_i^*(\omega)$ for all $\omega \in \Omega$. This means that $(x^*, p^*) \in BWE(\mathcal{E}^*)$.   □

Corollary 5.6. *If Ω is finite and $u_i(\omega, \cdot)$ is concave for all i and ω, then under assumptions* (A.0)–(A.4), *there exists a limit symmetric Bayesian–Walrasian equilibrium (p^*, x^*) of the limit economy $\mathcal{E}^*$.*

Proof. Theorem 4.1 in Balder and Yannelis (2009) guarantees in each period t the existence of a Bayesian–Walrasian equilibrium (x^t, p^t) for the economy $\mathcal{E}^t$. From Theorem 5.5 it follows that a limit symmetric Bayesian–Walrasian equilibrium exists.   □

The pair (x^*, p^*) is trivially incentive compatible because in the economy $\mathcal{E}^*$ all agents have the same information. We now recall the definition of incentive compatibility (see Krasa and Yannelis (1994) or Koutsougeras and Yannelis (1993)).

Definition 5.7. An allocation x fulfills coalitional incentive compatibility if and only if the following does not hold: there exist a coalition $S \subsetneq I$ and states $a, b, a \neq b$ which members of $I \setminus S$ are unable to distinguish and such that members of S are strictly better off by announcing b whenever a has actually occurred. Formally x is coalitional incentive compatible if and only if there do not exist a coalition $S \subsetneq I$ and states $a, b, a \neq b$, such that

    (1) $b \in E^{\mathcal{F}_i}(a)$ for every $i \notin S$ and

    (2) $u_i(a, e_i(a) + x_i(b) - e_i(b)) > u_i(a, x_i(a))$ for every $i \in S$.

It is still an open question whereas in each period $t \in T$ the BWE satisfies the incentive compatibility since the feasibility requirement may not hold. On the other hand, in the limit symmetric information economy, incentives to misreport the realized state of nature do not exist. Therefore the limit symmetric information BWE is coalitional incentive compatible, and fortiori individual incentive compatible.

Moreover a limit symmetric Bayesian–Walrasian equilibrium of the economy $\mathcal{E}^*$ is also (interim) Pareto optimal as the next proposition indicates. We first recall the notion of an (interim) Pareto optimal allocation.

Definition 5.8. An allocation $x \in L_X^\infty$ is said to be interim Pareto optimal if the following is not true: there exist $y \in L_X^\infty$ and a state $\omega \in \Omega$, such that

(i) $\quad v_i(y_i|G^\infty)(\omega) > v_i(x_i|G^\infty)(\omega) \quad$ for all $i \in I$

(ii) $\quad \displaystyle\sum_{i \in I} y_i(\omega) = \sum_{i \in I} e_i^*(\omega).$

The proof of the Proposition below is standard.

Proposition 5.9. *Any limit symmetric BWE of the limit symmetric information economy $\mathcal{E}^*$ is (interim) Pareto optimal.*

6. Stability of the symmetric Bayesian–Walrasian equilibrium

We now show the converse of Theorem 5.5, i.e., any limit symmetric information BWE can be reached by a sequence of approximate BWE outcomes. We can view this result as a stability property of Bayesian–Walrasian equilibrium, in the sense that we can always construct a route to reach the limit symmetric information BWE.

Before we state the result, we first define the notion of an approximate (or ε-) BWE where each agent optimizes the conditional expected utility within a small error $\varepsilon > 0$ in each state. Since we allow the possibility of making errors in maximization, agents are considered to be bounded rational.

Definition 6.1. A pair (p, x) is said to be an ε-**Bayesian–Walrasian Equilibrium** (ε-BWE) of the differential information exchange economy $\mathcal{E}$ if

(i) $\quad p : \Omega \to \Delta$ is a random price vector,

(ii) $\quad x \in L_X$ is an allocation,

(iii) $\quad$ for all i and for all ω the consumption function

$x_i \quad$ maximizes the conditional expected utility within $\varepsilon > 0$, i.e.,

$v_i(x_i|\mathcal{F}_i)(\omega) \geq max\{v_i(y_i|\mathcal{F}_i)(\omega) : \ y_i(\omega) \in \hat{B}_i(\omega, p)\} - \varepsilon$

and $\quad x_i(\omega) \in \hat{B}_i(\omega, p),$

(iv) $\quad p(\omega) \cdot \sum_{i \in I}(x_i(\omega) - e_i(\omega)) = \max_{1 \leq h \leq \ell} \sum_{i \in I}(x_i(\omega) - e_i(\omega))_h$ for all $\omega \in \Omega$.

We now show that a BWE in the limit can be reached by a sequence of approximate BWE allocations. In other words, given a price-consumption pair which is a BWE for the limit symmetric economy, and obviously exact feasibility holds, we can construct a sequence of ε-BWE that converges to it. Notice that in the sequence, agents maximize approximately their interim utility subject to their estimated budget set and also market clearing does not hold. Therefore, we can conclude that the BWE is stable.

The following theorem on the approximation of BWE equilibria has an important interpretation as Theorem 5.5. Here again in order to appreciate the value of this result one should contemplate of its failure: if some equilibria in the limit economy failed to be approximated even by approximate BWE, that would mean that those equilibria are artifacts of the definition and they would be irrelevant for analytical purposes. That would cast doubts on the value of the concept of BWE itself.

Theorem 6.2. *Let $\{\mathcal{E}^t : t \in T\}$ be a sequence of differential information exchange economies satisfying assumptions* (A.0)–(A.5) *and let $(p^*, x^*) \in \text{BWE}(\mathcal{E}^*)$. Then, for any $\varepsilon > 0$, there exists a sequence $(p^t, x^t) \in \text{BWE}_\varepsilon(\mathcal{E}^t)$ such that (p^t, x^t) converges to (p^*, x^*).*

Proof. Let $x_i^t = \mathbf{E}[x_i^*|\wedge_{j=1}^n \mathcal{F}_j^t]$ and $p^t = \mathbf{E}[p^*|\wedge_{j=1}^n \mathcal{F}_j^t]$. Then, since $\{\wedge_{j=1}^n \mathcal{F}_j^t\}$ is monotone increasing, we have

$x_i^t = \mathbf{E}[x_i^*|\wedge_{j=1}^n \mathcal{F}_j^t] = \mathbf{E}[\mathbf{E}[x_i^*|\wedge_{j=1}^n \mathcal{F}_j^{t+1}]|\wedge_{j=1}^n \mathcal{F}_j^t] = \mathbf{E}[x_i^{t+1}|\wedge_{j=1}^n \mathcal{F}_j^t],$

$p^t = \mathbf{E}[p^*|\wedge_{j=1}^n \mathcal{F}_j^t] = \mathbf{E}[\mathbf{E}[p^*|\wedge_{j=1}^n \mathcal{F}_j^{t+1}]|\wedge_{j=1}^n \mathcal{F}_j^t] = \mathbf{E}[p^{t+1}|\wedge_{j=1}^n \mathcal{F}_j^t].$

Hence, by construction $\{(p^t, x^t), \wedge_{j=1}^n \mathcal{F}_j^t\}$ is a martingale and by the martingale convergence theorem (p^t, x^t) converges to $(p^*, x^*) \in \text{BWE}(\mathcal{E}^*).$

To complete the proof, we must show that (p^t, x^t) lies in $\mathrm{BWE}_\varepsilon(\mathcal{E}^t)$ for all but finite t's. That is, the set

$$K = \{t \in T \,:\, (p^t, x^t) \notin BWE_\varepsilon(\mathcal{E}^t)\}$$

should be finite. Suppose, by way of contradiction, $(p^t, x^t) \notin \mathrm{BWE}_\varepsilon(\mathcal{E}^t)$ for infinitely many t's.
By using the feasibility (iv^*) of x^* and assumption (A.5), we have

$$
\begin{aligned}
\sum_{i=1}^n x_i^t &= \sum_{i=1}^n \mathbf{E}[x_i^* | \wedge_{j=1}^n \mathcal{F}_j^t] = \mathbf{E}\left[\sum_{i=1}^n x_i^* | \wedge_{j=1}^n \mathcal{F}_j^t\right] \\
&= \mathbf{E}\left[\sum_{i=1}^n e_i^* | \wedge_{j=1}^n \mathcal{F}_j^t\right] = \sum_{i=1}^n \mathbf{E}[e_i^* | \wedge_{j=1}^n \mathcal{F}_j^t] \\
&= \sum_{i=1}^n e_i^t \qquad \text{by} \quad (A.5).
\end{aligned}
$$

Thus, for all $\omega \in \Omega$ condition (iv) is satisfied. In fact,

$$p^t(\omega) \cdot \sum_{i \in I}(x_i^t(\omega) - e_i^t(\omega)) = 0 = \max_{1 \le h \le \ell} \sum_{i \in I}(x_i^t(\omega) - e_i^t(\omega))_h.$$

Since x_i^t is defined as the conditional expectation on $\wedge_{j=1}^n \mathcal{F}_j^t$, it is $\mathcal{F}_i^t$-measurable, and hence the conditions (i) and (ii) are satisfied.

Then, from the above assumption, there exist an agent i and a state $\omega^* \in \Omega$ such that, for infinitely many t's,

$$x_i^t(\omega^*) \notin \mathrm{argmax}_{y_i^t \in \hat{B}_i(\omega^*, p^t)} v_i(y_i^t | \mathcal{F}_i^t)(\omega^*) - \varepsilon. \tag{6}$$

Therefore, there exists a sequence $y_i^t \in L_{X_i^t}$ such that

$$v_i(y_i^t | \mathcal{F}_i^t)(\omega^*) > v_i(x_i^t | \mathcal{F}_i^t)(\omega^*) + \varepsilon, \text{ and}$$

$$\hat{p}_i^t(\omega^*) \cdot y_i^t(\omega^*) \le \hat{p}_i^t(\omega^*) \cdot e_i^t(\omega^*). \tag{7}$$

Since $y_i^t \in L_{X_i^t}$, which is a compact set as noted in the proof of Theorem 5.5, then we can find a subsequence $\{y_i^{t_n}\}$ of $\{y_i^t\}$ which converges to y_i.

Since for all[12] i, $\hat{p}_i^{t_n}$ converges to $\hat{p}^*$, $y_i^{t_n}$ to y_i, and $e_i^{t_n}$ to e_i^* (recall assumption (A.4)), then from (7) it follows that

$$\hat{p}^*(\omega^*) \cdot y_i(\omega^*) \le \hat{p}^*(\omega^*) \cdot e_i^*(\omega^*).$$

Moreover taking limits in the expression below

$$v_i(y_i^{t_n} | \mathcal{F}_i^{t_n})(\omega^*) > v_i(x_i^{t_n} | \mathcal{F}_i^{t_n})(\omega^*) + \varepsilon,$$

it follows that

$$v_i(y_i | G^\infty)(\omega^*) \ge v_i(x_i^* | G^\infty)(\omega^*) + \varepsilon > v_i(x_i^* | G^\infty)(\omega^*),$$

which is a contradiction to the fact that $(p^*, x^*) \in \mathrm{BWE}(\mathcal{E}^*)$. Thus K must be a finite set, i.e., $(p^t, x^t) \notin BWE_\epsilon(\mathcal{E}^t)$ for only finite number of t's. Therefore, there exists a sequence $(p^t, x^t) \in BWE_\epsilon(\mathcal{E}^t)$ that converges to $(p^*, x^*) \in BWE(\mathcal{E}^*)$, as it was to be shown. $\quad\square$

7. Conclusions

We introduced the concept of a symmetric BWE. This concept is free of the undesirable property of the asymmetric BWE which does not allow for market clearing. Most importantly, the symmetric BWE can provide a dynamic foundation for the asymmetric one since by repetition, agents observe the equilibrium outcomes, learn the information of the others and in the limit symmetric information prevails. Thus, despite the fact that from period to period the agent's price estimates are not accurate and market clearing does not hold, nonetheless, in the limit the correct estimates will prevail and market clearing will hold. We also showed that the converse is true, i.e., given a symmetric BWE, we can always construct a sequence of approximate BWE which converges to the symmetric BWE.

---

[12] See the proof of Theorem 5.5.

The concept of the symmetric BWE is Pareto optimal, coalitional incentive compatible and it exists universally (compare with the REE which exists only generically). This new way has a statistical foundation as it is based on price estimates and on an estimated budget set, which makes our modeling attractive to empirical analysis.

Acknowledgements

References

Allen, B., 1981. Generic existence of completely revealing equilibria with uncertainty, when prices convey information. Econometrica 49, 1173–1199.
Balder, E.J., Yannelis, N.C., 2009. Bayesian–Walrasian equilibria: beyond the rational expectations equilibrium. Economic Theory 38, 385–397.
Glycopantis, D., Muir, A., Yannelis, N.C., 2005. Non-implementation of rational expectations as a pefect Bayesian equilibrium. Economic Theory 26, 765–791.
Glycopantis, D., Muir, A., Yannelis, N.C., 2009. On non-revealing rational expectations equilibrium. Economic Theory 38, 351–369.
Koutsougeras, L., Yannelis, N.C., 1993. Incentive compatibility and information superiority of the core of an economy with differential information. Economic Theory 3, 195–216.
Koutsougeras, L., Yannelis, N.C., 1999. Bounded rational learning in differential information economies: core and value. Journal of Mathematical Economics 31, 373–391.
Krasa, S., Yannelis, N.C., 1994. The value allocation of an economy with differential information. Econometrica 62, 881–900.
Kreps, M.D., 1977. A note on the fulfilled expectations equilibria. Journal of Economic Theory 14, 32–43.
Peleg, B., Yaari, M.E., 1970. Markets with countable many commodities. International Economic Review 11, 369–377.
Pesce, M., Yannelis, N.C., 2009. Learning and Stability of the Bayesian–Walrasian Equilibrium. Manchester EDP-0923, see http://www.socialsciences.manchester.ac.uk.
Podczeck, K., Tourky, R., Yannelis, N.C., 2010. Generic non-existence of expectations equilibrium with infinitely many states. Unpublished preliminary draft.
Radner, R., 1979. Rational expectation equilibrium: generic existence and information revealed by prices. Econometrica 47, 655–678.