



Quaderni di Dipartimento

Intelligent people vs. intelligent design. On T-equilibria as a tool for interpreting both social order and social change

Giorgio Rampa
(Università di Pavia)

110 (01-10)

Dipartimento di economia politica
e metodi quantitativi
Università degli studi di Pavia
Via San Felice, 5
I-27100 Pavia

Gennaio 2010

**Intelligent people vs. intelligent design.
On T-equilibria as a tool for interpreting both social order and social change**

January 2010

1. Motivation

Adam Smith emphasised on many occasions the fact that many people, when looking at marvellous or impressive phenomena, tend mistakenly to attribute their source to the intelligent design of some invisible hand. If the impressive phenomenon relates to the apparent order of social affairs –Smith seems to argue– again many people tend to believe that such state of affairs derives from some intelligent plan¹. In a similar vein, Hayek (e.g. 1937) insisted diffusely that what economists call “equilibrium”, brought about by the intelligent design of rational beings, should be properly interpreted as a “social order”, brought about by the unintentional forces of evolution that shape the intentions of intelligent people.

The aim of this paper is not to pursue the connections between Smith, Hayek and evolutionism, an otherwise interesting topic (see e.g. Marciano, 2009, for interesting arguments related to this theme). Instead, it tries to offer a sort of general and abstract setup which can hopefully capture, at least partially, the sense of Smith’s and Hayek’s social order, and at the same time can account for endogenous social change.

We start from the well-known observation that a basic characteristic of the economy, namely interaction among individuals, when tackled by means of the standard tools of game theory, leads to the requirement of an infinite degree of rationality, which is unfeasible for real intelligent individuals. Indeed, when interacting individuals act purposively, they should consider strategic interaction, and this setting is modelled as a game. Apart from simple textbook cases, *incomplete information* is the normal condition, and this can lead to a high degree of indeterminacy of the possible outcomes, which is disliked by game theorists. Harsanyi (1967-68) proposed an ingenuous trick to bypass this problem, namely the *common prior assumption*: each player’s opinion about the surrounding environment, in particular about other players’ strategies, is constrained to be the conditional distribution extracted from a larger common prior, conditional on which of a set possible *types* that player actually is; and all this, including the common prior, is *common knowledge*. Under this assumption, one can resort to the usual Nash equilibrium concept to restrict the set of the possible outcomes of games of incomplete information played by rational players.

Now, that a common prior exists and is commonly known is already a questionable tenet. Aumann (1987) argues informally that people grown in a common environment will converge to common opinions, even if they start from differentiated ones. However, in order not to fall into tautologies², this property should be proved, and the present paper will precisely argue that it is not necessarily true. But, even accepting for the sake of argument that a common prior exists, a more fundamental issue is raised by it. A prior is a distribution over some support, and a common prior requires a *common support*. Under incomplete information, each player understands that all other players are incompletely informed as well, and are thus endowed with distributions over the surrounding environment, including other players' characteristics. On the other side, a player's characteristics include her/his distribution over the environment: hence we are lead to an infinite regress of distributions over distributions (infinite "belief hierarchies"). From this it follows that the dimension of the common support of a common prior must be *infinite*: this was formally proved by Mertens and Zamir (1985)³, and confirmed by Aumann and Dreze (2008).

Since we are interested in real interacting people, and not in hyper-rational beings, this perspective is untenable. Paralleling the arguments of Keynes' beauty contest⁴, it is impossible that individuals proceed much further into this reasoning, and ore research is needed to capture some general properties of social systems. It is as if, having dared to taste the fruits of the Tree of Knowledge of Evil and Good, we are now condemned to a hard job. Or, put differently, it is as if, having tried to build a Babel Tower, we must accept its collapse and the ensuing fact that people speak many languages: indeed, if our individuals cannot ground their interactions on a common prior, then they necessarily speak different languages.

This does not mean that one must abandon any hypothesis of intelligence on the part of agents, nor that individuals do no 'understand' each other. We simply drop the assumption of infinite rationality (if you like, call 'bounded rationality' the present alternative), considering individuals who tackle their choice problems in an interacting context being endowed with theories of a *finite* nature. On the other hand, we need to enquire the conditions under which individuals, while maintaining theories that are not constrained by a common prior, are lead to believe that they understand each other. This means that what they can observe is no surprise to them, at least during a finite time-span T , hence they keep maintaining those theories and acting accordingly. We term *T-equilibrium* a situation like this.

In order to develop out argument in an abstract way, and to capture some general properties of *T-equilibria*, we need to resort to some simple technical machinery. In Section 2 we describe informally how the setup can be construed. Section 3 presents a slightly more formal treatment, and offers a Main Proposition, stating the following: (a) the set of *T-equilibria* is a continuum whose

dimension increases as individuals become more ‘sophisticated’, i.e. there is not a unique, or ‘natural’, ordered state of the system (*indeterminacy*); and (b) that dimension decreases as the time-length of a T -equilibrium increases, i.e. the system is bound to exit eventually any ordered state it was in (*endogenous social change*). In Section 4 some closing interpretations and implications are put forward.

2. Intuition

Consider a set of individuals who interact in a certain environment: i.e., they purposively make choices at subsequent dates in time, knowing that these choices affects each other’s welfare in some way. *Actions* chosen by individuals are not to be interpreted restrictively as simply market acts, such as supplies of and demands for goods: gestures, signals, words or phrases are actions as well. An action is in general a multidimensional object at a single date: besides being interested in different objects at the same time, an individual might, e.g., offer a certain quantity *and* announce she/he will never produce that good again *and* look away knowing that the announcement is false.

Individual actions unfold in time, that is, individuals formulate *intertemporal* action plans. Consider a given time, and call it “the planning date”: at this date each individual is formulating a sequence of actions for subsequent periods: the length of such sequence is presumably finite, though it might be longer or shorter depending on the individual being more or less farsighted. Given that a single-period action is multidimensional, the whole plan might well be a very high-dimensional one.

As regards the passing of time, we assume that there exists a common clock, a typical human artefact that everyone understands: the indexing of dates (no matter whether they are minutes or weeks) is thus recorded by the clock’s ticking. It is not necessary to assume that all individuals plan to perform a fully dimensional action (or any action at all) at each date: some people have a more frenetic way of life than others have, hence the former act more often than the latter do⁵. Being the clock–time common to all individuals, if an agent plans not to implement (or does not plan to implement) a one- or a higher-dimensional act at any date, simply put a zero in the corresponding coordinate(s) of her/his action vector at that date.

In order to formulate an action plan an individual must hold opinions and expectations about the working of the surrounding environment, that is about actions that will be undertaken by other

individuals⁶. We will call of an individual's opinions her/his "*individual theory*" (or "individual model") at the planning date.

An individual theory plays a *twofold* role: one the one hand it produces expectations about the other individuals' actions during the planning horizon, and on the other hand it prescribes actions to be implemented by the individual during the same planning horizon. Besides depending on expectations, planned actions must be thought of as being "optimal" with respect to some personal welfare criterion of the planning individual, and with respect to expectations. Here we do not want to go into any technical detail of the individual "maximisation" problem: we simply assume, for the sake of argument, that starting from their personal models all individuals are able to formulate expectations and to choose the ensuing preferred actions.

What sort of objects are personal theories? One might want to think of them in terms of standard statistical models, as is common in most traditional economic theorizing. Of course this cannot be excluded for some individuals, for instance those trained in statistics or in post-graduate economics. It is however reasonable to think that individuals usually give their opinions a less explicit and structured form. This does not mean that it is impossible to offer a formal abstract representation of personal theories: a promising route would be in terms of neural networks, since this kind of modelling appears to be more affine to some principles of the basic physiology of reasoning⁷. What is suggested here is that individuals *are* neural networks, and not that they not *use* them as computing devices (as is fashionable today in, e.g., financial forecasting).

In both cases, i.e. whether one considers standard statistical models or neural networks, a personal theory has perforce a finite nature, as argued in the introductory remarks. In other terms, a personal theory can be represented by some model characterised by a finite number of parameters. Assume in addition that the model takes the form of a function mapping from parameter values to planned actions and expectations: that is, to each different parameter configuration there correspond different intertemporal actions and expectations. From the present point of view, then, an individual's theory (or model) is equivalent to the configuration of its parameters.

We say that agents (or agents' theories) are more 'sophisticated' if they are characterised by a higher number of parameters, meaning of course that they consider a higher number of variables (own and others' actions). Notice however that the number of parameters increases not only if an individual theory considers the others' actions, but also if that individual tries to "rationalise" their possible actions as a result of their possible theories, leading to some "belief hierarchy", to be considered *finite* anyway: see Remark 3 in the next Section. In addition, individuals can be deemed to be more sophisticated also if they plan their actions for longer time-spans.

At the planning date (call it date “1”), an individual, endowed with a theory (i.e. a parameter constellation), plans her/his future actions *and* has expectations on the others’ actions⁸. If, after observing the others’ actual actions of the first date, she/he finds that her/his previous expectations are deluded, then she/he will be induced to modify her/his theory. This sets a *learning* story in. However, we do not want, and are presently unable, to study a learning process in analytical terms: hence, apart from some observations put forward in the concluding Section, we leave this problem aside. Suppose, on the contrary, that the individual’s expectations are fulfilled, given other agents’ choices: in this case her/his expectations on subsequent actions will remain unchanged, and the same will hold for planned actions. We reword this situation by saying that the individual’s theory is confirmed.

Suppose now that our individual’s expectations are fulfilled for T subsequent dates: then her/his expectations and planned action remain equal to those formulated at the original planning date, and she/he keeps maintaining her/his theory for all this period. Finally, suppose that *all* individuals happen to be in such a situation starting from the initial date until date T : we call *T-equilibrium* this situation.

A *T-equilibrium* is a state of affairs such that the actions undertaken by the individuals do not induce anyone to modify her/his theory, which in turn informs her/his actions, along the whole T -long time span. This equilibrium notion is consistent with Hayek’s one⁹.

In a sense, one might say that along a *T-equilibrium* individual are allowed to believe that they are endowed with “rational expectations”, or that they live in a Nash equilibrium. However, this does not imply that they ‘know the true model of the economy’: in the present setup there is *no* true model of the economy. In fact, the Main Proposition proved in Section 3 states that the number (better, the dimension of the set) of *T-equilibria* increases with agents’ sophistication, i.e. with the number of parameters that characterise their theory, for given T : the more sophisticated agents are, the wider is the set of possible *T-equilibria*. Hence, in a world populated by sophisticated agents there are continuously many possible equilibria, and none of them is superior to the others in terms of ‘rationality’ of expectations. A trait that differentiates the present equilibrium notion from the conventional rational-expectations (or Nash-equilibrium) one is that it has a *temporary* nature (indeed, the “ T ” symbol might stay for “temporary”). We do not see any problem in this, since we do not believe that real people have at their disposal an infinite amount of time to test whether they are Lucas-type (or Nash-type) people.

On the other side, it follows from our Main Proposition that, for given agents’ sophistication, the set of *T-equilibria* shrinks as T increases. Since the dimension of this set decreases monotonically with T , a particular time T^* will be reached in which it becomes null: no T -

equilibrium can exist any longer after that date. Even if individual theories were so lucky as to support a T^* -equilibrium, the system will exit that equilibrium at date T^*+1 , that is agents' expectations will start being deluded. At that point, people will necessarily need to change their minds, i.e. try to learn something different. Social change, then, is *endogenously implied* by our setup. We will come back to this point in the concluding Section.

3. Analysis

Consider k interacting individuals. An action performed by individual i at a single date is a multidimensional array, i.e. a vector that is taken to belong to some subset of the n_i -dimensional Euclidean space¹⁰ for simplicity. Assuming that the dimension of the single-date action set is equal at all dates, it follows that i 's intertemporal plan is formally a vector whose dimension is $n_i T_i$: T_i is the time horizon of i 's plan, which is finite by assumption.

A personal theory extended for a T_i -long time span can be represented by some model G_{i,T_i} characterised by a finite number of parameters, that is real numbers. Let m_i be the number of parameters of i 's personal theory. Individual i 's personal theory is uniquely defined by its parameter vector: each different parameter configuration implies a different theory. Call M_i the set of possible theories (that is, parameter configurations) of individual i at the planning date, with $M_i \subseteq \mathbf{R}^{m_i}$: a single parameter configuration for agent i will be denoted by $\mu_i \in M_i$.

The model G_{i,T_i} can be construed as mapping from M_i to expectations *and* planned own optimal actions over a T_i -long time span. This requires, of course, that many underlying assumptions be satisfied in order that a possibly very complex optimal control problem is solved by the individual¹¹. We assume that those assumptions hold; we assume in addition that G_{i,T_i} is a *function*, i.e. it is one to one, and that it is (at least piece-wise) smooth.

Now, split G_{i,T_i} into the expectation and the action components: call G_{i,T_i}^E the component of G_{i,T_i} that outputs i 's expectations, and call G_{i,T_i}^A the component of G_i that prescribes i 's own optimal actions (to be sure, also own actions are expressed in expected value terms: see Remark 2 below). Hence, G_{i,T_i}^A maps from a m_i -dimensional space to i 's individual actions, whose dimension is $n_i T_i$; and G_{i,T_i}^E maps from a m_i -dimensional space to i 's expectations about all

individual actions over a time horizon of length T_i . Finally, setting $n = \sum_j n_j$, define $p_i(T_i) = nT_i$:

this is the total number of variables about which individual i might entertain expectations, i.e. the number of joint intertemporal actions on the part of *all* individuals, extended to i 's personal time horizon. We have thus $G_{i,T_i}^E : M_i \rightarrow \mathbf{R}^{nT_i}$, and $G_{i,T_i}^A : M_i \rightarrow \mathbf{R}^{nT_i}$.

Some cautionary Remarks are in order.

Remark 1. We have assumed that each individual maintains expectations about each intertemporal action of any other individual: in fact the space of i 's expectations coincides with the joint space of all individuals' actions. This entails assuming that n_j is known to *all* i 's. However, is it not necessary that an individual considers the whole space of each other individual's actions. Indeed the individual, while forming expectations, might omit some variables, or might even be unaware of their existence. In this case, one could simply put a 'blank' in the corresponding element of her/his expectation vector: then this vector stays of the same formal dimension as before, namely $p_i(T_i) = nT_i$. Alternatively, one might consider explicitly the case $p_i(T_i) < nT_i$: our result, however, would not be modified significantly, as we shall see. In any case, we will disregard the case in which an individual forms expectations about actions *not* included in the actual action spaces of the other individuals.

Remark 2. Observe that i 's expectations are extended to include *own* actions: in fact the sum in the expression $p_i(T_i) = \sum_j n_j T_i$ runs over all individuals, including i her/himself. This means simply

that, since own planned actions depend upon expectations about the others' actions, they are themselves expected and not certain, given own theory at the planning date. Expected own actions are simply the "best responses" to expected actions on the part of other individuals. It follows that G_{i,T_i}^A , the action part of i 's model as defined above, describes *expectations* of own actions, not simply own actions, from date 2 onward (of course, first-date actions are planned in a deterministic way).

Remark 3. We must now reason on the number m_i of parameters defining i 's theory. Roughly speaking, one would argue that for each variable on which expectations are formed there is *at least* one parameter in the model. Hence, i 's model should include at least $m_i = nT_i$ parameters: but such number would be greater than this, if the model incorporated affine or non-linear forms¹²: it could be e.g. $m_i = a_i nT_i$, where $a_i > 1$ is the average number of parameters for each expected variable in i 's model. Hence we have $m_i > nT_i$. In the Introduction we argued however that individuals, while not being infinitely rational, might try to "outsmart" their fellows: that is, they might elaborate on

the other individuals' theories, in order to better rationalise their actions¹³. This means that other people's theories become objects of own theories, giving rise to finite "hierarchies of beliefs". Now, we know that individual models coincide with their parameter configurations: therefore if individual i wishes to account for other people's models, the number of the parameters in i 's model itself must increase accordingly. This effect will be the greater, the more is i 's theory "sophisticated", i.e. the higher is degree of her/his belief hierarchy. If for instance i thinks that the others are "basically" as intelligent as i her/himself, then the deduction is that every other's model should have $a_i n T_i$ parameters. To account for all these parameters about which i is uncertain, i 's model should include at least $(k-1)a_i n T_i$ parameters, where k is the number of interacting individuals: this is the implication of having a formal theory over the others' theories. If i is even more sophisticated, believing that all other individuals have theories over theories, and wishes to account for this, then i 's model should contain at least $(k-1)^2 a_i n T_i$ parameters; and if you go h steps ahead into this logical process, you will get $(k-1)^h a_i n T_i$ parameters. Our discussion shows that the number of parameters of each individual theory, m_i , can be as large as $(k-1)^h a_i n T_i > (k-1)^h n T_i$, where $h \geq 0$ depends on how much sophisticated an individual is¹⁴.

Now, we come to *T-equilibrium*. Define first $T = \min_i T_i$, that is the time-horizon of the less farsighted individual. Assume that each individual forms expectations about the *actual* actions spaces of all individuals, i.e. her/his expectation function contains no 'blanks' (see Remark 1 above). The implication of this is that $p_i(T) = nT$ is *common* to all individuals, hence we can write $p(T)$ to denote this common dimension of the expectation space of any individual. Call $G_{i,T}^E(\mu_i)$ the 'projection' of the image of $G_{i,T_i}^E(\mu_i)$ on its first T time-coordinates. Define further $G_{i,T}^A(\mu_i)$ as the projection of i 's action function on its first T time-coordinates: the image of this function has dimension $n_i T$. Setting $\mu = (\mu_1, \mu_2, \dots, \mu_k)$, define finally the *joint* action function of all individuals as $G_T^A(\mu) = (G_{1,T}^A(\mu_1), G_{2,T}^A(\mu_2), \dots, G_{k,T}^A(\mu_k))$. One has thus $G_T^A : M \rightarrow A_T$, where $M = \prod_i M_i$ and $A_T = \prod_i A_{i,T}$. The dimension of M is $m = \sum_i m_i$ and that of A_T is $p(T) = nT$.

It is reasonable to think that if expectations are deluded, then an individual will change her/his mind, i.e. expectations and plan. This would lead to some form of *learning*, on which we cannot elaborate here in an analytical way. If, on the contrary, expectations are fulfilled up to a certain date t , then we can say that the individual does not change her/his expectations and plan for subsequent dates. Indeed, calling $E_i(x)$ the prior expected value of x and $E_i(x|y)$ the posterior expected value

after observing a realization y , one has $E_i(x|E_i(x)) = E_i(x)$ under any reasonable learning process. Since we are assuming that individuals are so bright as to plan initially their future actions conditional on any event they might happen to face, it follows that expected values of subsequent events are computed recursively, based on expected values of previous events: if the latter remain unchanged after observing actual actions up to a certain date t , the former, concerning the dates from $t+1$ onward, remain unchanged as well. In particular, the *own* action actually undertaken at date $t+1$ stays equal to what the individual expected beforehand. Thus: if actual actions undertaken by all other individuals at all dates from 1 to T are equal to what individual i expected initially, namely they are equal to $G_{i,T}^E(\mu_i)$, then actual actions undertaken by i are equal to what she/he expected initially, namely $G_{i,T}^A(\mu_i)$. This situation might be called an *individual* equilibrium of length T .

We are led to define a *general T -equilibrium*, or simply a *T -equilibrium*, as a state of affairs in which all individuals find themselves in a individual equilibrium of length T . From the argument of the last paragraph it follows that along a T -equilibrium actual actions on the part of all individuals are equal to $G_T^A(\mu)$, as defined above. In order that a T -equilibrium holds, it necessary that each individual's expectations on all agents' choices are equal to their actual realisations. Hence one can put forward the following

Definition 1a. *The system is an a T -equilibrium if $G_{i,T}^E(\mu_i) = G_T^A(\mu)$, $\forall i$.*

However, we saw in Remark 2 above that i 's expectations are extended to include own actions. Now, under an optimal plan, expectations on own actions *coincide* with own planned actions (taken, of course, in expected value terms). In other terms, we can say that the projection of $G_{i,T}^E(\mu_i)$ on i 's own actions is equal to $G_{i,T}^A(\mu_i)$. A short reflection will then show that we can give the following alternative definition of a T -equilibrium:

Definition 1b. *The system is an a T -equilibrium if $G_{i,T}^E(\mu_i) = G_{j,T}^E(\mu_j)$, $\forall i, j$.*

Whether the system is in T -equilibrium or not depends clearly on the configuration of all individual theories, i.e. on the configuration of all individual parameters described by the vector μ . In addition, in our setting all relevant expectations and actions depend on those parameters. For this

reason we attach the T -equilibrium property to μ vectors, and say that a vector $\mu \in M$ is a T -equilibrium if the condition of Definition 1b is satisfied. Now we enquire whether T -equilibria can exist, and ‘how many’ they are in the parameter space M . As regards this, we can offer following

Main Proposition. *The set of T -equilibria has generically topological dimension $m - (k - 1)nT$ in the parameter space M .*

Proof: see the Appendix.

Only if the dimension quoted in the Proposition is non-negative¹⁵ T -equilibria exist generically¹⁶. In particular, if this dimension is equal to zero, equilibria are generically *finite* in number, while if it is positive the set of equilibria forms generically a *continuum*.

One may ask how likely the continuum case is. Recall the following: k is the number of interacting individuals; $m = \sum_i m_i$ is the number of parameters of all individual theories taken together; $n = \sum_j n_j$ is total number of (scalar) actions undertaken by all individuals at each date; and T is a common time-horizon to which individual plans and expectations are extended, corresponding to the time-horizon of the less far-sighted individual. Thus, nT is the total number of intertemporal actions on which individuals form expectations by means of their theories, restricted to the time-horizon of the most myopic of them. On the other side, recall what we observed in Remark 3 above, and in the ensuing definitions: from those arguments it follows that for each individual one has $m_i > nT$; better, m_i can well be much higher than nT , considering that she/he might want to speculate on other people’s theories (up to a finite degree), and given that a joint distribution on a certain number of variables requires an even higher number of parameters (see footnote 12 above). Hence, $m = \sum_{i=1}^k m_i$ can well be much higher than $knT > (k - 1)nT$, leading to $m - (k - 1)nT > 0$. As a consequence, the case of a continuum of equilibria is fairly probable. If our individuals become more sophisticated, i.e. if m increases, the dimension of the T -equilibrium set grows, for given T . Finally, the case of an individual ‘not caring’ of some action undertaken by some other individual (see Remark 1 above) imposes fewer constraints on the equilibrium condition, and hence raises further the dimension of the equilibrium set.

Notice: given our assumption of smoothness of the functions $G_{i,T}^A(\mu_i)$ mapping from parameters to actions, one expects that to different equilibrium parameters there correspond different equilibrium action plans. Hence, one cannot hope to ‘refine’ the set of equilibria in terms

of observational equivalence. In other words, we do not expect that the set of equilibrium actions is ‘thin’, independently of how large the set of equilibrium parameters is. In addition, at the present level of generality it is not feasible to refine the equilibrium set using Pareto-ranking arguments.

Another important implication of the Main Proposition is the following one: given the single-date joint action space (i.e., given n) and given the ‘sophistication’ of individual theories (i.e., given m), an increase in the time-horizon T reduces the dimension of the T -equilibrium set, since it lowers the number $m - (k - 1)nT$ until it becomes null or negative. For high T ’s, a set of theories of given sophistication cannot support generically any T -equilibrium.

4. Interpretation and implications

In this Section we offer some brief comments, grouped under ten headings, whose aim is to work out some interpretation of the above analysis, and to show how it might direct future research.

Intelligence and social order. Our intelligent agents try somehow to outsmart each other elaborating on each other’s opinions, in order to better forecast their fellows’ future actions and hopefully obtain higher returns. They are lead to do this, precisely because they are conscious of interaction. However, it is impossible that everyone outsmarts everyone else, since this would require an infinite amount of silicon chips, or of grey matter. Along a T -equilibrium –or, in Hayek’s terms, a *social order*– what makes agents satisfied with their theories is that they do not happen to observe anything that comes as a confutation. “Full rationality”, as understood by some, would require infinitely more than this; “intelligence” requires instead a finite amount of thinking. We would take it as an axiom –it comes also from our daily experience– that intelligent people have a preference for ‘psychological tranquillity’: if a theory ‘works’, meaning that nobody sends me signals contradicting it, one prefers to stick to it, and not spend resources inventing a new one. After all, we know of very long-lasting social orders in which many intelligent people believed that the Sun goes around the Earth. Closer to our matters, we know of social orders in which many intelligent people believe that economic policies are (are not) effective, and they *see* that they are (are not) so.

Indeterminacy and relativity. The existence of a continuum of possible social orders is called *indeterminacy* by economists. In our setting, its meaning is that there are (infinitely) many different possible ways in which interacting people might conceive correctly their interaction, and act accordingly. In other terms, there is *not* a single model (the ‘true model’, as Rational-Expectation

scholars would call it) of the economy. To each different constellation of individual models, if they form a social order, there corresponds a different evolution of the economy that confirms those models (see also the quotation from Hayek in footnote 9 above). This relativity might disturb someone pretending to ‘refine’ the set of equilibria to reach uniqueness. However, this is a problem of our economist, not of the agents populating a certain social order: the latter are living in a *single* social order, and are not so lucky as to experience something which might suggest to them the existence of other orders. In addition, given the preference for ‘psychological tranquillity’, they are hardly motivated to ‘tremble’ around the equilibrium they are living in. It seems odd that a theorist pretends that the agents living in her/his model be endowed with the same theory that she/he is endowed with. Put in a different, and perhaps interesting, way: the economist her/himself should be viewed as an agent, maintaining a certain individual theory and interacting with other people in the economy.

Indeterminacy and nature. Basically, our indeterminacy derives from the high number of *degrees of freedom* that characterises agents’ theories, compared with the set of events that they might observe, namely their own actions. There is an argument put forward by Rational-Expectations theorists to the effect that agents should use *parsimonious* models: hence, ‘over-parametrised’ models are not welcome by those theorists. On the one side, the argument is motivated by the econometrics practice, as if it were pretended that all agents be themselves econometricians, which is not necessarily the case; on the other side, it purports to urge that agents use the same model the theorist uses, which again is questionable as we argued before. In addition, if it were the case that we know the true model of the world, and handled it in a fully rational way, it would be impossible to explain why theories (and scientific theories) change in time and evolve: if that were the case, we should be in a position to know all the theorems of algebra instantly without any ‘sweat of our brow’, and without any need to learn new things. But, more importantly, that argument seems to be contradicted by the mere observation of our brain and of its functioning: its capabilities, even if never exploited hundred percent, are fantastic in terms of degrees of freedom (see footnotes 7 and 14 above). A related example comes from biology: if organisms were projected in the same way as engineers project their tools (i.e. without any degree of freedom), and were not the result of DNA replication –a good instance of high redundancy– we would know no evolution.

Social order and language. Along the unfolding of any one social order the actions performed by individuals, and circulating among them, can be interpreted as signals, hence as a (more or less sophisticated) language: an ‘ordered’ dialogue is one where the words emitted are understandable, i.e. are no surprise, for the receiving part. So, different social orders imply different languages: if a population is split into different parts by the collapse of a giant Babel Tower, and these parts are

confined in different lands, they would end up speaking different languages even if they faced very similar experiences. True: if contacts among them restart, some form of translation can be devised, which is however imperfect; but most probably a new language will be born (more on this point below, under the headings ‘learning’, ‘nurture’ and ‘change’). Given the argument put forward above on the high number of degrees of freedom, one expects that a language is ‘less complex’ than any underlying individual theory: it is the tip of the iceberg, and the ‘intimate’ meaning of words might well differ significantly among different individuals. However, along a social order nothing happens that unveils these differences.

Indeterminacy and learning. Since a society of interacting intelligent individuals can end up in one of infinitely many possible social orders, *which one* of them will be picked up depends necessarily on learning. If their expectations are deluded, individuals will modify their theories, and this sets up a dynamical system: the stability properties of this system are obviously very hard to be worked out analytically, but we can easily state a couple of general properties thereof. First, given multiplicity, the final state of the learning process depends on initial conditions, and if some noise is present on actual history as well: as a consequence, what is learnt in an interactive setting is not something objectively ‘true’, differently from what happens when we say we have learnt the length of the Earth day. Second, as non-linearity is surely a significant feature of the dynamical system, small differences in initial conditions can lead to large differences in final states (the space covered by the ruins of the Tower is very small with respect to the Earth surface, but languages are different from each other).

Indeterminacy and nurture. Not only learning, but also teaching is a fundamental activity within social systems. Individuals do not enter their interactions starting from nothing except a given set of free parameters: they come from many past experiences during which older members tried to educate them in different fields. This raises the interesting question addressed to traditional economic theorists: at which age do rational individuals start being rational? Apart from our preference for ‘intelligent’ instead of ‘rational’, a reasonable answer could be ‘short after conception’. Not only: a longer-term genetic learning takes place in the population, shaping the biological bases of individual theories. Hence, many constraints contribute to actually ‘refine’ the set of possible social orders. This notwithstanding, any new problem faced in an interactive setting raises the questions we have been discussing in this paper. Even supervised training is no guarantee that the pupil learnt ‘exactly’ what the teacher taught: if this were always the case, we could be back to the ‘no sweat’ ex-ante knowledge of all theorems, and we could not explain cultural and scientific evolution and differences.

Endogenous social change. We saw in Section 3 that, given the dimensions of the joint theory and action spaces, the set of T -equilibria shrinks as the time T increases, until when *no* social order can exist generically. Thus, a social system where all individual theories support a social order is *bound* to exit that order sooner or later. Some of the individuals, and more of them as time elapses, will start observing unexpected facts, i.e. facts that delude the expectations formed on the basis of those theories. This will require a change of mind and a new learning process, whose outcome is fairly indeterminate *ex ante*, for the very same reasons discussed until now. This is why we contend that the notion of T -equilibrium can explain, not only social order, but also social change as an *endogenous* phenomenon. On the one hand, sophistication implies indeterminacy, resolved only by learning and history; on the other hand, finiteness implies that no mind can forecast everything in the indefinite future, and that every theory is in essence provisional. To be sure, there can exist very simple-minded social orders in which things repeat equal to themselves day by day, and this can last for centuries. But nothing in our arguments prevents that planned actions follow very complicated paths¹⁷, and the more so if individuals are more sophisticated. Hence, there will be someone who will stop understanding them sooner or later.

Learning, precision and catastrophes. It is a property of almost all known learning processes that what Bayesian statisticians call ‘subjective precision’, that is the confidence individuals attach to their own theories, increases as they accumulate more observations. It corresponds also our normal experience. This entails that individuals find it more and more difficult to change their minds as time elapses. If, then, the system exits eventually a given social order, people will start observing ‘strange’ things: they will be tempted to consider those things as ‘outliers’, and give them no importance, coherently with their high precisions. But if outliers keep accumulating, this situation will start being intolerable. Traditional statistical practice offers no tool to handle this problem, but one can argue that on these occasions sudden jumps, or ‘catastrophes’, take place in individual minds and in the whole system. If one were to model theories and learning by means of neural networks, this would precisely be the case also in technical terms, given the typical non-linearity incorporated in these objects (see footnote 7 above). The width of these jumps is higher, the deeper our theories are rooted in our minds, i.e. the higher their precisions are, and the more the observations are at variance with theories.

Innovation. The exit from a previous social order is a source of uncertainty both for individuals and for the whole system. Given our ‘axiom of psychological tranquillity’, one expects that people are gratified if they can resolve that uncertainty in a short time¹⁸. Suppose that one individual succeeds in announcing a new explanation of how things are going, *and* in convincing a critical mass of other people of its soundness: there is an *innovator*. Three seem to be the conditions for this

success: *brightness*, or sophistication on our sense, in that one needs to anticipate the reactions of the other individuals to the new theory; *luck*, in that one cannot know the others' minds in detail and anticipate their reactions; and, above all, *power*, in that the new theory must arrive loud and clear to the recipients in order to overcome their current uncertainty.

Strong ties and open societies. As we said under the heading 'endogenous social change', there exist 'simple-minded' social orders in which things repeat themselves almost equal day after day. This case is typical, mainly, of small communities that are relatively closed with respect to the outside world. In these communities the 'ties' tend to be strong (Granovetter, 1973 and 1983; cp. also Burt, 1995), not only because everyone is linked to everyone else in the group and to none of the outside world, but mainly because the repetition of commonly understood experiences tends to reinforce everybody's confidence in her/his personal theory, seemingly shared by everyone (culture?). These communities are particularly fragile to powerful enough external shocks, like the opening of new links with other communities entertaining different views and speaking unintelligible languages. Observe, however, that the same can be true of larger communities governed by simple and strong ideas embedded in people's minds by means of an artful propaganda. In both cases, the contamination with hitherto unknown 'viruses' can lead to very keen breakdowns in the previous prevailing social order (including wars), which are painful for society. On the contrary, 'weak ties' imply that people are prepared to weigh new ideas and theories. For this reason, some recommend that societies are always 'open' to external stimuli, in order to minimise their potentially disruptive effects on an exceedingly hard-shell social system. One should also compare this view with the fundamental tenets of classical Liberals: free entry, social mobility and equal opportunities.

¹ Most noticeable is Smith's treatment of this point in Section III of *The History of Astronomy* (Smith, 1795, already written by 1773); the same point, as is well known, is resumed scantily in *The Theory of Moral Sentiments* (1759), Part IV chap. 1, and in *The Wealth of Nations* (1776), Book IV chap. 2. See also Macfie (1971) and Rassekh (2009); the latter argues in a vein similar to that put forward in the present paper.

² Aumann (1976) showed that interacting people, sending some signal to each other, cannot eventually agree to disagree: they will arrive to common opinions in finite time. However, in Aumann's argument the complete description of the possible states is known to all them initially, together with a map from what they can observe and the possible refinements of their posteriors. In other terms, the common prior is already there at the outset.

³ More precisely, these authors proved that the dimension of the support must be higher than infinite, since it must be equal to the transfinite number ω , i.e. the lowest number which is greater than any natural number.

⁴ One should notice that a similar argument was put forward by Hayek as well (1952, pp.185-ff.), when he observed that a single mind cannot understand any other individual mind in detail, since this would lead to an infinite regress.

⁵ This way of putting things might be interpreted as agreeing with an oft-maintained subjectivity of the perception of the flow of time on the part of different individuals.

⁶ One could conceive opinions and expectations as being defined about other variables as well. These "system" variables might be entirely exogenous, which are fairly uninteresting from our point of view; or they might depend in turn on individual actions: think e.g. of "GDP", or of prices formed as results of market demands and supplies (some prices should however be included among individual actions, if a subset of individuals are price-makers). This extension would not alter our setting significantly: one should (a) define the "structural" function mapping from individual actions to system variables, and (b) extend individual expectation formation to an enlarged space including the same system variables. However for simplicity we prefer to concentrate on action-variables only.

⁷ In this paper it is not possible to go into the details of neural network modelling, and we refer e.g. to Serra-Zanarini (1990), Hertz *et al.* (1991), Haykin (1994), Golden (1996), for introductions. We only recall the following properties of neural networks: (a) a network is composed of a finite number of units ("neurons") receiving signals from (and sending signals to) the other units; (b) the signals being sent between units depend on the value taken on by the sending units, and on the strength of the synaptic links between them; (c) a subset of the units are connected with sensor units

recording external stimuli, or with actuator units performing external “actions”; (d) to any external stimulus, which is a multidimensional array hitting many sensory units simultaneously, the corresponds a “final state” of the whole networks, a value for each unit, which is possibly translated into some action (again a multidimensional array); (e) stimuli giving rise to the same final state (or action) have the same “meaning”; (f) a final state induced by a stimulus is not necessarily constant: it can be an oscillation, if the network is a so called “recurrent network” (as opposed to a “feed-forward” one), that is if there are feedbacks in the net; (g) the working of a network, i.e. how stimuli are translated into final states and actions, is governed by a finite set of parameters, representing the strengths of the synaptic links between any pair of units; (h) although units can take on continuous values, under certain parameter values final states can be seen, for all practical purposes, as *discrete*, since units behave mainly as binary variables; (i) the configuration of parameter values at a certain date depends on the past history, that is on *learning* on the part of the network: learning takes the form of changes in parameter values; (j) given that neural networks incorporate significant non-linearities, a small change in parameter values can induce important changes in the working of the net, that is in the way different stimuli are given different meanings; (k) finally, and somehow interestingly for Austrian economics, we are convinced that neural networks are a good representation of some of Hayek’s (1952) ideas on “sensory order”.

⁸ To be precise, she/he plans own action of date 1, and has *expectations* on own actions of subsequent dates, since these depend on her/his expectations on others’ actions.

⁹ Consider the following famous quotation from Hayek (1937), p. 42: “Correct foresight is then not, as it has sometimes been understood, a precondition which must exist in order that equilibrium may be arrived at. It is rather the defining character of a state of equilibrium. Nor need foresight for this purpose be perfect in the sense that it need extend into the indefinite future or that everybody must foresee everything correctly. We should rather say that equilibrium will last so long as the anticipations prove correct and that they need to be correct only on those points which are relevant for the decisions of individuals”. See also Hahn’s (1973) notion of conjectural equilibrium.

¹⁰ This means that action are seen as continuous, not discrete objects. This might appear a strong assumption; for instance, words seem to be discrete in nature. See however footnote 10 above, related to neural networks, property (h), for a justification of this assumption. On careful reflection it will be understood that an individual rarely pronounces a given word in exactly the same way, depending on the detailed configuration of her/his neural network at each instant. Hence, ‘actions’ may well be deemed to be continuous variables.

¹¹ These assumptions include continuity of the personal welfare function, compactness of its domain, proper convexities, discount rate lower than one. We leave, however, the detailed study of the analytical solution to this problem to better equipped minds, if they are so interested.

¹² Consider in addition that, apart from the ‘structural’ form of the model, in purely statistical terms a distribution on n variables requires $m > n$ parameters. For instance, if one posits a uniform joint distribution on n independent variables, one needs $2n$ parameters, two endpoints for each dimension of the support; if one posits a joint normal distribution on n variables, one needs $n + n(n+1)/2$ parameters, the former addend for the means, and the latter for the symmetric covariance matrix.

¹³ Indeed, when we try to understand other people’s behaviour we are usually not interested in their actions point-wise, but we aim at finding *relations* between these actions, in order to be able to better anticipate their evolution. This is a clear symptom that we are trying to understand somehow the theories which underpin their actions.

¹⁴ Think of individual theories as stemming from natural neural networks, like human brains: in order to appreciate how large the number of parameters can be, refer to note 7 above and consider that the number of units in such networks is of the order of 10^{11} , while the number of their connections, i.e. of their parameters, is estimated to be of the order of (at least) 10^{14} .

¹⁵ Following the arguments of the Appendix, one understand the meaning of a ‘negative’ dimension from the following example. Consider the 3D space, and consider two 1D lines embedded in it. From the Theorem used in the Appendix, the so called ‘co-dimension’ of their intersection is 4, since it is the sum of the two co-dimensions of the lines, each of which is the complement of its dimension in the reference 3D space, and hence is 2. It follows that the dimension of the intersection is $3 - (2+2) = -1$, meaning that two lines do *not* intersect generically in the 3D space, an obvious fact.

¹⁶ A *generic* property is one that, *if* satisfied, keeps holding true in a whole full-measure subset of the reference space (here, the parameter space). Hence, what the Proposition proves is *not* that T -equilibria exist always, but that it is generically the case that they do exist, i.e. one can find parameter values such that they exist, if the reported dimension is non-negative.

¹⁷ See Boldrin-Montrucchio (1986) for an example of chaotic paths followed by optimal plans, even in purely individual, non interactive, problems.

¹⁸ There is a whole spectrum of the strength of this urgency, depending on the role one happens to play: the two extremes are possibly the sport or political supporter and the patient scientist.

References

- Aumann, R.J. (1976), Agreeing to Disagree, *The Annals of Statistics*, vol. 6, 1236-1239
- Aumann, R.J. (1987), Correlated equilibrium as an expression of Bayesian rationality, *Econometrica*, vol. 55, 1-18
- Aumann, R.J. and Dreze, J.H. (2008), Rational Expectations in Games, *American Economic Review*, vol. 98, 72-86
- Boldrin, M. and Montrucchio, L. (1986), On the indeterminacy of capital accumulation paths, *Journal of Economic Theory*, vol. 40, 26-39
- Burt, S.R. (1995), *Structural Holes: The Social Structure of Competition*, Cambridge MA, Harvard University Press
- Golden R.M. (1996), *Mathematical Methods for Neural Network Analysis and Design*, Cambridge MA, MIT Press
- Granovetter, M.S. (1973), The strength of weak ties, *American Journal of Sociology*, vol. 78, no. 6
- Granovetter, M.S. (1983), The strength of weak ties: A Network Theory Revisited, *Sociological Theory*, vol. 1, 201-233
- Guillemin, V. and Pollack, A. (1974), *Differential Topology*, Englewood Cliffs, Prentice-Hall
- Hahn, F.H. (1973), *On the Notion of Equilibrium in Economics. An Inaugural Lecture*, Cambridge, Cambridge University Press
- Harsanyi, J.C. (1967-68), Games with incomplete information played by “Bayesian” players, parts I-III, *Management Science*, vol. 14, nos. 3, 5, 7
- Hayek, F.A. von (1937), Economics and Knowledge, *Economica*, vol. 4, 33-54
- Hayek, F.A. von (1952), *The sensory order*, Chicago, University of Chicago Press
- Haykin S. (1994), *Neural Networks. A Comprehensive Foundation*, New York, Macmillan College Publishing Company
- Hertz J, Krogh A. e Palmer E.G. (1991), *Introduction to the Theory of Neural Computation*, Redwood City, Addison-Wesley
- Macfie, A. (1971) The Invisible Hand of Jupiter, *Journal of History of Ideas*, vol. 32, 595-599
- Marciano, A. (2009), Why Hayek is a Darwinian (after all)? Hayek and Darwin on social evolution, *Journal of Economic Behavior and Organization*, vol. 71, 52-71

- Mertens, J.F. and Zamir, S. (1985), Formulation of Bayesian analysis for games with incomplete information, *International Journal of Game Theory*, vo. 14, 1-29
- Rassekh, F. (2009), In the shadow of the invisible hand, *History of Economic Ideas*, vol. 17. no. 3
- Serra, R. e Zanarini, G. (1990), *Complex Systems and Cognitive Processes*, Berlin, Springer
- Smith, A. (1759 [2002]), *The Theory of Moral sentiments*, K. Haakonnsen ed., Cambridge, Cambridge University Press
- Smith, A. (1776 [1937]), *An Inquiry into the Nature and Causes of the Wealth of Nations*, E. Cannan ed., New York: The Modern Library
- Smith, A. (1795 [1980]), *Essays on Philosophical Subjects*, W. P.D. Wightman, J. C. Bryce, Dugald Stewart, I. S. Ross, D. D. Raphael, A. S. Skinner eds.; Oxford, Oxford University Press

Appendix

Proof of the Main Proposition

Firstly, the individual expectation function extended over the T time-horizon is $G_{i,T}^E : M_i \rightarrow A_T$. The *joint* expectation function $G_T^E(\mu) = (G_{1,T}^E(\mu_1), G_{2,T}^E(\mu_2), \dots, G_{k,T}^E(\mu_k))$ can be seen as mapping from the space $M = \prod_i M_i$, whose dimension is m , to the space $(A_T)^k$, the k -th Cartesian power of A_T , whose dimension is knT . Call X_1 the *graph* of this function: given our assumptions, X_1 is (at least piece-wise) a *smooth manifold*. The topological dimension of X_1 is m , and it is embedded in the space $X = M \times (A_T)^k$, whose dimension is $m + knT$.

On the other side, the T -equilibrium condition stated in Definition 1b is $G_{i,T}^E(\mu_i) = G_{j,T}^E(\mu_j)$, $\forall i, j$. It requires that the images of all individual expectation functions intersect in the “diagonal” Δ of the Cartesian power $(A_T)^k$, whose dimension is nT . Define the product space $X_2 = M \times \Delta$: X_2 is clearly a smooth manifold embedded in the space X , and its dimension is $m + nT$. The T -equilibrium condition is thus satisfied at the intersection $X_1 \cap X_2$, which is a subset of X_1 .

We want to evaluate the dimension of the subset of M such that X_1 intersects X_2 . Said differently, we want to evaluate the dimension of the projection of $X_1 \cap X_2$ on M . Since the projection of X_1 on M is of ‘full’ dimension in M , it follows that the dimension of the projection of $X_1 \cap X_2$ on M is the same as the topological dimension of $X_1 \cap X_2$.

Define the *co-dimension* of a set Y in the embedding space Z as $\text{codim}(Y) = \dim(Z) - \dim(Y)$. Consider now the following Theorem (see e.g. Guillemin-Pollack 1974, p. 30):

The intersection of two manifolds X_1 and X_2 in the space X is a manifold, and the following holds:
 $\text{codim}(X_1 \cap X_2) = \text{codim}(X_1) + \text{codim}(X_2)$.

We have $\text{codim}(X_1) = m + knT - m = knT$ and $\text{codim}(X_2) = m + knT - m - nT = (k-1)nT$. It follows $\text{codim}(X_1 \cap X_2) = knT + (k-1)nT$, and thus $\dim(X_1 \cap X_2) = m - (k-1)nT$. **Q.E.D.**