
Computable Cosmos

ERIC M. SCHEFFEL

Department of Economics and Quantitative Methods, Nottingham University Business School China, The University of Nottingham, 199 Taikang East Road, Ningbo 315100, Zhejiang, China

Email: Eric.SCHEFFEL@nottingham.edu.cn

Web: <http://www.nottingham.edu.cn/en/business/people/eric-scheffel.aspx>

Bio-sketch: Eric Scheffel's work covers both the theoretical and empirical domains. He has experience in the fields of applied time series econometrics and international finance. In particular, he has undertaken work to analyse non-linear time series models of the U.K. term structure of interest rates. More recently, his research has concentrated on employing general equilibrium models of the economy, incorporating financial intermediation, seeking to explain the behaviour of monetary factors (such as money demand and various velocity measures), fundamentals and asset prices. While working at the Office for National Statistics he extensively researched confidential micro data using modern panel data econometrics methods. He also shares a great interest in the way the internet is transforming traditional modes of production and exchange, as well as the agent-based computational economics literature.

Abstract: This essay contrasts the Hayekian notion of spontaneous order with mathematically deducible and computationally attainable equilibria within economic systems, which are often the focus of general equilibrium analysis. My analysis focuses on the perspectives these two approaches to social systems take on the evolution of system configurations, which can be interpreted as equilibria, limiting ergodic distributions, attracting states, homeostatic orders, or simply states of the system. While Hayek's notion of a spontaneous order is very broad, suggesting the impossibility of precise systematization of causal forces, equilibria are typically determinate as a result of a narrowly defined Walrasian definition. Combining elements of Hayekian spontaneous order with recent insights from computer science, I introduce the concept of a real-time concurrent computable equilibrium system, or computable cosmos. In contrast to agent-based computational models, computational cosmos are endowed with the property of fundamental system indeterminacy, which originates from the concurrency of agent behaviour.

Keywords: Agent cloud; computational economics; general equilibrium; stability analysis; statistical physics.

INTRODUCTION

A casual study of some of Hayek's works, such as *Law, Legislation and Liberty* (Hayek, 1979/2012), or his article on the role of knowledge in economic systems (Hayek, 1945), makes clear the breadth and scope of his analysis, with the notion of spontaneous order or 'cosmos' at its core. Yet Hayek's treatment of the subject also betrays its complexity and intractability, notwithstanding his own assertion that although

we cannot see, or otherwise perceive the order of meaningful actions, [...] we are [...] able mentally to

reconstruct it by tracing the relations that exist between the elements.

In yet another passage he remarks that although such emergent orders lack a man-made purpose

our awareness of its existence may be extremely important for our successful pursuit of a great variety of different purposes.

While Hayek's notion of social order may seem to resemble the ubiquitous concept of equilibrium, the two are clearly of a different nature, both in the way they are defined as well as treated theoretically. In contrast to Hayek's

definition of spontaneous order, which is associated with the question of how aggregate order can arise out of any apparent chaotic or random behaviour exhibited by its constituent elements, the notion of equilibrium is connected to Leon Walras' (1874/1954) mathematical treatment. Walrasian economics is based on a general equilibrium treatment of supply and demand, as they are thought to interact in interconnected markets.

Each perspective involves conceptual trade-offs. While Hayek analyses the interconnected roles played by individuals, formal and informal institutions, and markets, he also recognizes the complexity and frequent intractability of emergent orders, and therefore proceeds by systematically ring-fencing and tightening his own theoretical structure. He does this by means of an intricate web, consisting of a comparative dialectic chronicled in much of his work on the subject¹.

Contemporary economists' view of markets, by contrast, is mathematical and thus exudes the definitiveness one associates with a series of interlocking mathematical proofs and lemma. What is however sacrificed is the multifaceted nature of a more complex market co-ordination process, a loss which results from the abstraction and simplification implicit in only considering a highly stylized perspective of markets.

This essay compares Hayek's notion of a spontaneous or homeostatic order, in a computable sense, and modern economists' take on this issue. While some may be sceptical towards the juxtaposition of these two theories, due to their different theoretical conceptions, the question of *system stability* is shared by both approaches and can be explored from various angles. Once I have made clear how much narrower and more deterministically treated system equilibria really are as compared with their Hayekian counterpart, I will deepen my analysis by introducing recent advances in computer hardware and software. I then define a real-time concurrent computable equilibrium system, or *computable cosmos*, and compare this with agent-based computational models. Both approaches make use of modern multi-core digital processors.

HAYEKIAN AND PURELY ECONOMIC SELF-ORGANIZING SYSTEMS

A large number of labels describe the configurations toward which systems—if and whenever they do so—self-organize *endogenously* over time. Qualifiers such as 'resting points', 'fixed points', 'equilibria', 'equilibrium growth paths', 'spontaneous orders', 'homeostatic orders', 'limiting ergodic distributions', and very likely a few others come to mind, all beset with their own idiosyncrasies. Within the specific context of the Hayekian notion of an emergent order, the social order or equilibrium he has in mind manifests itself as an abstract regularity or pattern. This regularity is detectable at a regional or global level in a topological sense, with individual actors' precise adaptive actions exhibiting a comparatively low degree of predictability in any given instance. At the same time, however, a certain degree of constancy in conduct and reaction to stimuli is expected to occur at the individual micro level, so as to make it possible for an abstract macro order to emerge.

A notion of equilibrium from the mathematical-statistical domain which closely mirrors the Hayekian definition is that of a *limiting ergodic distribution*, in which large numbers of observations occurring (or drawn) over time eventually settle down to a stable *distributional pattern* with ascertainable statistical properties, in spite of the fact that each new observation that is added to the existing pool is again random or stochastic. Hayek himself almost certainly contemplated a system which transcended primitive notions of statistical stability or stationarity, such as a system endowed with seemingly complex or even chaotic behaviour.² Such a system would however only appear to be random or chaotic because of people's inability to grasp the complex forces of cause and effect. In Hayek's view, however, an abstractly defined order is still in existence, and its properties are at least partially traceable, if only in some probabilistic sense.

The perspective of mainstream economic theory on self-organizing order or stability within mathematical and applied general equilibrium analysis is more tractable than Hayek's notion of a spontaneous order, at least in terms of the degree of complexity involved in mathematical modelling. Yet any rigorous analysis of such systems, which seeks precise answers to the questions of existence, uniqueness, and stability of equilibria, can be mathematically demanding, involving fixed-point theorems and investigations of dynamic stability using Lyapunov functions, to give but two examples. Still, tractable solutions and theorems are obtained in most cases, notwithstanding the mathematical sophistication required.

At the same time much of this tractability is only attained by abstracting from a more general problem. The aim is to achieve a dimensionality which makes possible the application of analytical methods and the derivation of so-called 'closed-form' solutions. One example of such a simplification involves reducing the real-world problem of

analysing numerous interconnected markets to a model with only three such markets (Anderson *et al.*, 2004; Hirota, 1981; 1985; Scarf, 1959). This three-market case invites speculation over the extent to which general results derived from such a simplified model carry over to a ‘massively scaled’ version of the same model. The only way to deal with this question is by employing (computer) simulations. Such simulations are particularly important to my discussion and I will discuss them in later sections.

So far I have argued that Hayek’s notion of a self-organizing system is broad, complex and largely mathematically intractable. This is certainly true if we understand tractable as good at replicating most of the elementary factors in an emergent order. Examples of sources of intractability in emergent orders include agents’ locally bounded and idiosyncratic information sets, their boundedly rational processing of such information, and the imperfect functioning of ‘global co-ordination devices’.

Like the laws of motion of gas nebulae or the mutual attraction or repulsion of atoms making up larger compounds, social actors perceive, adapt, and act *concurrently* in *real time*. By contrast, the iterative approach in traditional Walrasian general equilibrium systems is strictly *sequential* and *deterministic*, and so does not take into account any real-time concurrency or fundamental indeterminacy. Before contrasting this with the novel concept of a *computable cosmos*, I will first attempt to explain why the mathematically modelled Walrasian *tâtonnement* mechanism constitutes a rigidly sequential and thus not a real-time concurrent mechanism. I will then explain the implications of this such as how each particular solution reflects considerations of mathematical tractability at the aggregate level.

SIMPLE ITERATED WALRASIAN TÂTONNEMENT DYNAMICS

Walras’ (1874/1954) dynamic market equilibrium mechanism is his *tâtonnement* process. Walras limited his theory to a verbal-intuitive statement that explained the underlying logic of this process. It is however now more commonly formalized as a mathematical system of equations, which consists of *aggregate excess demand functions* for all interdependent markets (e.g. Arrow, 1952; Debreu, 1954; Hildenbrand, 1994; Mantel, 1974; Radner, 1968; Smale, 1976; Sonnenschein, 1973). An excess demand function for a given market sums all consumers’ individual demands net of all their individual incomes (or supplies) of the particular good traded in that market. As is well known, the attainment

of equilibrium in any one market requires that its excess demand function is identically equal to zero for some vector of prices. This vector determines consumers’ demands and incomes.

By analogy, a specific equilibrium vector of prices that results in a zero-valued vector of excess demand in *all markets simultaneously* constitutes a general equilibrium solution to the entire mathematically modelled market system. Economists working with such systems of excess demand functions can approach their analysis by employing two different but interrelated approaches to the problem. The first one involves solving the system of (often non-linear) equations simultaneously in a static sense. The second concerns itself with the question of whether and how a dynamic updating scheme may be capable of transitioning the system towards a general equilibrium price vector, when it starts from a given initial *disequilibrium* price vector and a predetermined distribution of endowments. The question of *existence* addresses the conditions under which at least one equilibrium price vector is guaranteed to exist, while that of *uniqueness* deals with conditions with only one such general equilibrium ‘resting point’.

In this essay I prefer to focus my attention on the *stability* criterion, which implies a solution to the first two problems, and which asks whether and in which way unique or multiple ‘resting points’ can be attained dynamically along a transition path from an initial disequilibrium price vector. A somewhat less demanding but equally important criterion concerns under what conditions the system remains *non-explosive* or *controllable*, interpreted in a strictly dynamic sense³. This requires a more detailed description of the traditional Walrasian *tâtonnement* price-adjustment mechanism, which is a somewhat artificial sequential dynamic process, reflecting the assumed behaviour of a Walrasian auctioneer.

The Walrasian dynamic *tâtonnement* process starts with an assumed disequilibrium vector of market prices⁴ as well as an assumed initial distribution of consumer endowments of all traded goods. In this simple description of an economy, the only activity is exchange, and any notions of production, investment or general economic growth are entirely absent. In addition, many early models restrict the analysis to three goods and three consumers, so as to ensure analytical tractability.

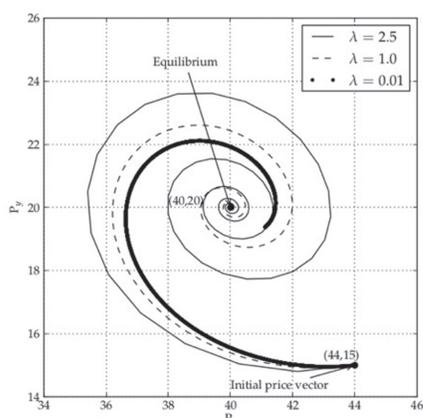


Figure 1: Simulated trajectories of a Scarf-type economy; three dynamic simulations with different linear speed-of-adjustment coefficients, all asymptotically approaching the equilibrium price vector

With given prices and endowments, the second step assumes that each consumer computes the value of her endowments (i.e. her income) and the quantity and value of her demands. These endowments and demands are conditional on a pre-specified stable individual demand function. Each consumer reports her computed supply and demand for goods to a global authority in the guise of the *Walrasian auctioneer*, who can then centrally process all individual consumers' supplies and demands in order to compute the value of the excess demand in each market. If this turns out to be different from zero for any, some, or all markets, a third step is needed.

The third step is key.⁵ So long as the prevailing vector of market prices causes excess demands to be different from zero everywhere and all markets are therefore simultaneously in *disequilibrium*, the auctioneer has to decide how to *update* the current vector of prices. Preferably, the updating should result in a new system of excess demands that is closer to the zero-valued vector. There is a good reason for why this step has become such a closely scrutinized one in the literature. It is because the manner in which the prevailing price vector is updated is the single most important factor in determining the *stability* of the modelled market, aside from the initial conditions and the distribution of consumers' endowments.

In the fourth and final step, the updated price vector is passed back (or publicly 'cried out' as in a realistic auction environment) to all the consumers, enabling them to update the value of their incomes and demands for all goods in the economy. This information will then in turn pass back to the

Walrasian auctioneer for the next round of re-computing excess demands, in which all steps are repeated in a looping iteration. The iterative process stops when the most recently updated price vector is sufficiently close to the static equilibrium. It is however important to note—within the specific context of this iterative process—that the actual exchange of goods is at all times assumed to be *completely suspended* until the general equilibrium price vector is found. It is only after this discovery that trade among agents will occur. Figure 1 illustrates three simulated relative price trajectories obtained from a typical Scarf-type general equilibrium model. Each simulation uses a different linear speed-of-adjustment coefficient. In this case, one common element shared by all three simulations is dynamic stability, since all of them approach the system's static 'resting point' asymptotically.

This completes the description of the iterated Walrasian *tâtonnement* dynamics and also serves as an example that illustrates typical dynamic economic modelling. At the same time it also demonstrates the somewhat artificial set-up of the model, as well as the sequential way it treats the question of the dynamic attainability of a general equilibrium. It accomplishes this through a system of (non-linear) differential equations.

A critical stance is particularly valid if one thinks in terms of the multifaceted, complex and seemingly chaotic character of a Hayekian cosmos. Which features of the sequential dynamic equilibrium process should be viewed with more scepticism from the vantage point of an endogenously emerging spontaneous order or cosmos?

The first and perhaps most serious problem is the assumption of the existence of a Walrasian auctioneer. This assumption has been criticized by the so-called 'Post-Walrasian' school of economics (Colander, 1996, 2006; Holt *et al.*, 2011; Kirman, 2010a, 2010b). The most critical Hayekian insight is the recognition of the central role played by decentralized and locally perceived knowledge, which continuously percolates through the aggregate system. Such a system thus comprises cognitively constrained actors, which however still gives rise to an order with a surprising degree of regularity, vitality, and robustness (Hayek, 1945).

The second problem, which is indirectly related to the first, is the assumption that an orderly and exhaustive collection of *all* relevant market information is conceivable between the various steps in the sequential process. This assumption enables consumers to report their incomes and demands to the auctioneer, who is able to compute excess demands once *all* information from *all* consumers has arrived.⁶ It also enables the auctioneer to 'cry out' the updated

and *uniformly and uniquely perceived* price quotations to the consumers, all of whom are assumed to perceive and employ this vector of prices symmetrically.

In fact, besides the assumed exhaustiveness of information, the whole structure of information exchange is highly artificial in the way the perfectly and globally knowable price vector is assumed to be publicly ascertainable between updating steps. Moreover, individuals supposedly transmit their information through impenetrable paths or silos to the global auctioneer, so that each agent is ignorant of his neighbour's consumption choices.⁷ The implication is that these models disregard temporary misperceptions and bandwagon effects; any notions of localized group-level information pooling, pre-processing or inter-agent knowledge exchange are entirely beyond the reach of this type of model.

One final weakness is the assumed lack of diversity of the information criteria circulating in the economy. All that appears relevant can thus be reduced to the set of prices and endowments.⁸ An example of how additional information criteria can matter in the determination of equilibria is the relative wealth endowments of individuals. This may affect the quantity, quality, and speed of dissemination of information, which in turn can lead to non-negligible consequences for the aggregate system as a whole and in particular for its overall stability, a result which Andersson (2008) explains in conceptual terms.

But does a research programme that focuses on simulating artificial economies with the help of digital processors constitute a more viable option for theorizing about market co-ordination problems? And what are the general implications, if any, which follow from allowing actors to behave and co-ordinate *concurrently* and in *real time*, instead of imposing the methodological straitjacket of iterative processing?

COMPUTABLE COSMOI AND THE SOCIAL-TO-DIGITAL MAP

The application of digital processors as part of the scientific method has been on the rise within many disciplines. In conjunction with this general development, increasing attention is being paid to *algorithms* or *numerical methods*, which are suitable for analysing a broad class of problems. This trend has also affected the social sciences, including economics (Judd, 1998). Another and perhaps even more influential use of the computer-assisted approach has been seen in mathematics, where the new field of 'experimental mathematics' has eschewed the traditional route of explicit derivation in

favour of alternative 'brute-force' methods using digital processors (Bailey *et al.*, 1997).

To this day modern computer hardware tends to be introduced to computer science students through the abstractly defined concept of a Turing machine (Turing, 1936). Turing machines form the theoretical precursors of contemporary computer hardware. Sometimes, digital processors executing algorithms are interpreted as *discrete state automata* (DSA)⁹, a conceptual sub-class of the more general Turing machine that is more suitable for understanding the operation of real-world computers, because of the theoretical capacity of Turing machines to store infinite numbers of states (and thus to possess *infinite memory*), and the self-evident limitation of DSA in this regard. For the purposes of our discussion, we may define computer programs or algorithms as software-encoded sequences of declarations, expressions, control statements and functions, all of which serve to manipulate information states or data (otherwise defined in the relevant literature as the generation of *side effects*). These are stored in alterable form in the computer's memory over the lifetime of a running program.¹⁰

One theory espouses the somewhat more radical view that the *entire universe* may be one grand computation carried out using a discrete state automaton (Zuse, 1970). This conjecture gave birth to the sub-discipline of digital physics, which remains unrefuted (although some would say that it cannot be refuted). The universe-as-computation conjecture, with all its vexing implications for the philosophical determinism-versus-free-will debate, poses a hypothesis that many may find difficult to digest. The less ambitious undertaking of modelling the market co-ordination process using discrete state machines (i.e. algorithms executed on digital processors) should thus be relatively uncontroversial.¹¹

I will now outline such an attempt of formulating a computer-simulated economy or *computable cosmos*. I will be mindful of clearly signposting the differences that exist between my conjectured approach, on the one hand, and that embodied by the application of so-called agent-based computational models (ABMs), on the other. The use of ABMs in any field related to economics is still relatively scant, although it has become more popular recently (Ashraf *et al.*, 2012; Colander *et al.*, 2008). In my proposed implementation of such a computable cosmos, I hope that its constituent characteristics will make self-evident how it is still remotely based on the traditional Walrasian *tâtonnement* mechanism, especially when viewed in its *algorithmic* form. At the same time, the approach is more ambitious in its ultimate goal of approximating, at least in principle, the more

broadly defined concept of a Hayekian emergent-order system. It should be able to achieve this by allowing for rich and varied knowledge exchanges at the micro level. The most distinguishing features of a computable cosmos—which demarcates it from the ABM literature—are agents' action *concurrency* and *real-time* computation, as well as the fundamental notion of systemic *indeterminacy*. The software paradigm I will be employing throughout to illustrate the design of such a system is that of the actor model (Hewitt *et al.*, 1973)¹², which I will argue lends itself ideally to the software-encoded implementation of a computable cosmos.

The Walrasian *tâtonnement* process is a sequentially unfolding one, in which exhaustive information collection was always shown to precede the functional evaluation of the system of excess demands and the vector of spot market prices. To make the conventional mathematical treatment of the market co-ordination problem feasible, this assumption of exhaustive information aggregation turned out to be a necessary one, since both the updating of the price and the consumers' demand functions depended on the supply of a complete, well-defined and knowable argument. This was then communicated to them for computation purposes.

It is therefore clear that the method's insistence on employing a system of dynamic differential equations, governing the aggregate behaviour of the entire system, necessarily implies information processing of the Walrasian type. By contrast, Hayek's work shows that the circulation, mutation, as well as perception of information or knowledge in society should not only be viewed as more complex, but also as richer regarding the amount and type of content which is transmitted between emitting and receiving actors.

Even more important, individuals all act *concurrently*, or at least possess the capacity of doing so in principle, while it is typically only because of institutional or innate factors that some individuals are capable of acting with lower latency, higher frequency, and more impact in practice. At the same time this particular diversity in behavioural properties never implies sequentially executed actions, and *intended* or *planned* actions are certainly always conceived of concurrently, actor-specific frictions in execution notwithstanding. So in contrast to the conventional Walrasian general equilibrium model, in which all agents are effectively homogenous in behaviour if not in endowments, a more realistic modelling approach should account for agent-specific differences. This should be apparent in the manner they process and perceive knowledge, thus introducing heterogeneity in that dimension.

Any attempt at generating a conceptual (and eventually computable) mapping that mimics a capitalist market system, which is translated into a finite state automaton, constitutes a fruitful and feasible exercise. This is so because many of the important features that constitute a snapshot of a capitalist society can be encoded using a finite set of digital information states¹³, as long as the scale and complexity of the system is kept to a manageable dimensionality¹⁴. Indeed, a simple mapping of this type applies to the simulation of a Scarf-type economy with three goods and three consumers (see Figure 1). This is in spite of differences between the two approaches regarding information-processing properties, basic iterated structure, and the extent of simplification.

But what is then the nature of the new actor model of computation, and how would it allow researchers to seek answers to questions associated with a market co-ordination problem that is considerably more complex than the Walrasian one? And how does it differ from agent-based computational models so that it is more deserving of the 'computable cosmos' description?

Herbert Gintis (2007) offers an instructive and recent example of an agent-based computational model, which is remotely based on the Walrasian price-adjustment system, but otherwise possesses more detail and complexity. Gintis' model transforms the Walrasian model into a system that includes analytically intractable non-linearities, such as agent-specific private reservation prices and replicator dynamics (Taylor and Jonker, 1978). Agents may thus copy or imitate more successful agents in a trial-and-error fashion. Moreover, because prices in the economy are modelled as *private reservation prices*, which apply to small-group trade and bargaining, the traditional Walrasian auctioneer is no longer present. The information content that is relevant to the market exchange process is thus local in nature.

Although the computation of the dynamically evolving price and exchange processes possesses certain random elements, such as probabilistic matching of agents in trade and replication (or imitation), the model is still a sequential one; the simulation steps follow the logic of a circuit flow possessing few if any logical branches. And in spite of the fact that the random matching of agents does introduce an element of indeterminacy into repeated simulations, this observed randomness is more closely related to simple 'sampling indeterminacy'. It is not related to the more fundamental and systemic type of indeterminacy one associates with spontaneously evolving orders. The salient regularities that emerge in repeated simulations are essentially constant in a convergent sense. They only differ from one another (in repeated

simulations) through minor variations in the ‘sampling noise’ that random matching generates.

Earlier in our discussion we noted that computers are machines that allow programmed logic to mutate states (or generate side effects). These states, in turn, are accessible in the form of data. The data is stored in the hardware’s random access memory during the lifetime of a running program, which in the shape of a compiled source code encapsulates and executes that logic either on the central processing unit (CPU) or on other special-purpose dedicated processing components (e.g., graphics processing units (GPU)).

The logic of the program as source code may define, *inter alia*, the behaviour and more static properties of agents, institutions, intermediaries and other system-relevant actors. Many of these will generate side effects and thus mutate states. Conversely, the stored information may encompass continuously updated variables such as private and public price signals, the evolving distribution of wealth, technology shock processes, agents’ changing perceptions of market conditions, monetary factors, and other agent-specific or public-domain variables. What is important is whether they are deemed relevant for a realistic simulation of a market co-ordination process.

A genuinely real-time *concurrent* variant of this modelling problem would allow agents to act *in parallel* during the lifetime of the program, instead of computing their behaviours and attendant side effects in *sequential* fashion. Not only would this imply a computational modelling paradigm that exhibits much more realism in agent interaction, but it would also introduce more fundamental indeterminacy than is possible through simple sampling variability. The concurrent-actor approach may thus generate aggregate pattern regularities that one would associate with a Hayekian spontaneous order.

The actor model was originally designed as a framework for concurrent software logic, and is currently implemented as part of the *Erlang* programming language (Armstrong, 2007). It constitutes an almost perfect tool for the implementation of a real-time concurrent computable equilibrium system, or *computable cosmos*, as described above. The model was developed in the mid-1980s at the Ericsson Computer Sciences Laboratory. The goal was to create a programming language for telephony network switches. Such programs have to be highly fault-tolerant and concurrent in operations, using a large number of extremely lightweight threads.¹⁵ The design of the language was therefore highly *domain-specific* and moulded to the specific characteristics, needs, and ob-

jectives of the particular hardware platform on which developed and compiled source code was to be executed.

In *Erlang* and the actor model it encompasses, a very large pool of actors—each endowed with some software-encoded behavioural logic—can be launched concurrently and in massive numbers. This trait allows actors to communicate with each other via the sending and receiving of messages, to which they then can be programmed to react by means of content processing and response formulation. It would not take too great a leap of imagination to recognize the intimate connection between the specific features of this domain-specific programming paradigm and a relevant computable-cosmos model. But what would be the benefits and costs of doing so, especially when compared with the orthodox mathematical approach embodied in Walrasian *tâtonnement* models?

One benefit from employing a real-time concurrent modelling would be an almost exact mapping of the simulation of a system that is governed by random or chaotic knowledge cascades. Such cascades would percolate through a cloud of knowledge-exchanging agents of the type that is so characteristic of Hayekian emergent orders. Furthermore, it is a well-known property of concurrent programming systems that they introduce fundamental indeterminacy, given that the concurrent dissemination of messages among agents is handled through software-controlled arbiters. Such arbiters ensure that the exact order in which messages are received and processed by actor addressees is unknowable *a priori*. The order is in one sense *chaotic* in practice (Hewitt, 2010).

It is chiefly this fundamental indeterminacy, which is also associated with a market economy, which explains why I chose to call the approach a computable cosmos. The randomness and potential richness in information or knowledge transfer facilitated by such a system intimately mirrors the Hayekian ideal of a spontaneous or homeostatic order. Moreover, all side effects and thus state mutations of the system have to be implemented by adhering to the practice of ‘message passing’, a design-specific straitjacket imposed by *Erlang* itself. This constraint forces programmers to use a set of laws or heuristics, which mirrors a market co-ordination system. It is thus natural to approach it as an emergent-order problem.

A fitting analogy to the workings of a simulated concurrent actor market model would literally be that of an autonomously functioning brain composed of a large number of synapses. Each synapse emits, accepts, processes and again

re-emits electrical impulses in rapid succession in a seemingly disorganized fashion at the aggregate level.

Well-defined software-encoded logical rules or heuristics may be incorporated into the source code, describing actor-specific and possibly also *satisficing* behaviour (Simon, 1947). It bears repeating that *Erlang* and its implementation of the actor model were developed specifically for the purpose of developing software systems that could control telephony communication systems or switches. Such systems have the attributes of complexity, fault-tolerance, concurrent operation, and a high throughput of message exchanges. These are all ideal prerequisites for the implementation of a computer-simulated market exchange model exhibiting rich and complex information transfer.

By contrast, one immediate drawback confronting any model builder, at least as compared with the treatment of conventional systems of differential equations, is the handling of time itself. Within the context of Walrasian *tâtonnement*, we discovered that each new time period was clearly demarcated. It was thus implicitly defined by the specific step in which a new market price vector was computed and ‘cried out’, based on the system’s prevailing excess demands. The iterative treatment of the market co-ordination process led to the natural identification of adjacent time periods, and time-series simulations of the model were obtained by iterating forward a first-order non-linear differential equation.

In real-time concurrent simulations of a market economy, no such easily defined demarcations of adjacent time periods obtain. This particular problem—associated with the collation, handling, and feeding back into the system of time-denoted state information—is a well-known one in digital processor-based simulations of market economies, and also tends to crop up with frequent regularity within agent-based computational models (LeBaron, 2001).

Within the specific context of an *Erlang*-programmed computable cosmos, one possible way of obtaining chronologically ordered values of current system states would be via the inclusion of a logically programmed ‘statistical authority agent’. Such an agent could at regular intervals be instructed to survey a random sample of evolving households and firms (within the simulation) exactly in the same way real-world statistical authorities would be instructed by governments. In principle, this would even allow the model builder to introduce some of the vexing real-world complications with which statistical authorities routinely grapple, such as time lags in collection and eventual publication, the introduction

of potential biases, and revisions as part of the data collection process.

The theoretical possibilities of computable *cosmoi* do not end there. As a result of each and every actor’s capacity to have her own digital ‘vessel’ with information on both her own characteristics as well as her *individually perceived* public-domain information states, an extremely rich and varied information landscape is conceivable at the model design stage. This would make it possible to explore numerous ways in which the system endogenously self-organizes over the lifetime of the simulation, including in-depth investigations into the dynamic stability property of the system.

One final but important remark is that one particular mode of inquiry, which necessarily undergoes a significant transformation in the role it plays in any such analysis, is that embodied in the *mathematical* characterization of the properties of any such modelled system of the market co-ordination process. We have seen that in traditional mathematical and applied analyses of general equilibrium systems, mathematical perspectives and tools which originate from the study of dynamic systems, functional analysis and topology have a critical bearing on the entire body of knowledge associated with the relevant subject matter. By contrast, within the context of computable *cosmoi*, but also discrete state automata more generally, the role of mathematics remains indispensable in many ways, yet digital processor-driven applications tend to migrate away from seeking answers to the properties of the system in its entirety, toward the characterization and specification of the behaviour and evolving information states encapsulated by individual, concurrently acting agents.

CONCLUSION

The use of computers within economics is an attempt to introduce a greater degree of both realism and complexity into what would otherwise end up as a whole class of mathematically intractable models. In this essay I have sought to compare and contrast conventional applied general equilibrium models with the novel concept of a *computable cosmos*, showing one feasible way forward for designing a computer-simulated *catallaxy*. The proposed approach would display richer and more varied knowledge exchange processes than those found in a Walrasian world.

I have alluded to the possibility of employing a mature programming language, *Erlang*, and its own built-in implementation of the concurrently capable actor model, in an effort to design a simulated complex market economy. The

proposed model's simulated evolution would not only unfold in real time, but with a fundamental, systemic degree of indeterminacy, thus eliminating the precise and iterated structure of conventional applied general equilibrium models. I have chosen to call this proposed software-encoded implementation a *computable cosmos*, both because of the fundamental indeterminacy that real-time concurrent computer programs exhibit and because the possible mutation of a very large number of digitally stored information states mirrors an information-rich exchange environment.

The biggest challenge faced by adopters of computable *cosmoi* is that of software-encoding behavioural rules and heuristics that govern the conduct of such concurrently acting actors. The biggest question remains why and how spontaneous orders emerge.¹⁶

NOTES

- 1 One part of Hayek's work where this approach is revealed is in the second chapter (*Cosmos & Taxis*) of *Law, Legislation and Liberty* (Hayek, 1979/2012).
- 2 The comparison with chaos is suitable in this context. Chaotic and unpredictable dynamic behaviour arises out of the solution to mathematical equations that are nevertheless determinate and functionally specifiable. A good introductory reference is Hirsch *et al.*, 2012.
- 3 For example, dynamic systems can remain away from equilibrium indefinitely, yet at the same time be non-explosive. This can be achieved by exhibiting stable orbital solution paths which keep oscillating around the equilibrium point (Scarf, 1959).
- 4 Strictly speaking, the process really employs a system of relative prices as one of the goods in the economy is commonly defined as the *numéraire* good which measures prices. In the popular DSGE macroeconomic modelling paradigm, the *numéraire* good is the consumption good.
- 5 In a seminal paper, Stephen Smale (1976) shows that global stability is only guaranteed if the Walrasian auctioneer has recourse to the derivatives (otherwise known as the Jacobian) pertaining to all excess demand functions. Given that real-world economies consist of essentially infinitely many such inter-related markets, the upshot of the result is that a stable updating scheme employs implausibly large amounts of information. This point is particularly relevant when viewing it from a perspective which recognizes Hayek's views on the informationally efficient mechanisms underpinning capitalist systems.
- 6 This is a well-known limitation applied researchers and policy-makers in macroeconomics have to grapple with: crucial economic indicators always arrive with a substantial time lag and are typically liable to be revised at an even later date.
- 7 This assumption is particularly problematic in view of 'keeping-up-with-the-Joneses' effects, conspicuous consumption, and the habit-persistence literature in general (Constantinides, 1990; Duesenberry, 1949; Veblen, 1899/2007).
- 8 One counter-example is provided by a model which dispenses with the Walrasian auctioneer and also allows for the exchange of quantity-information, leading to states of rationing. This model is described in Benassy, 1990 and, more accessibly, in Bridel, 2011.
- 9 Many economists, especially macroeconomists but also algorithmically-inclined game theorists, are likely to have employed a mathematical construct that can be represented in terms of a more general discrete state automaton, which is given by a Markov chain process.
- 10 Genuine data or state persistence in the form of physically recorded states that persist outside the scope of any running program is typically rendered operational through the use of physical hard drives or solid-state discs, or previously through the use of magnetic tapes.
- 11 Figure 1 already provides the visual output from a computer simulation (and is thus an implementation of a DSA approach) for a fairly standard Walrasian general equilibrium exchange economy, obtained by iterating forward through time the non-linear difference equation in the vector of market spot prices.
- 12 All computer programs are alike in their purpose of facilitating state or data changes, or mutations, prescribed by some overarching logic that is encapsulated in the source code of the program. How that logic is ultimately presented, arranged and structured, and thus engineered, can be fundamentally different across programs. It largely depends on the programming paradigm employed by the architect, of which the *functional* and the *object-oriented* are the two most prevalent in use. The former paradigm has dominated the industry for the last 30 years or so, while the latter one has experienced a renaissance of sorts in more recent times, due to the specific needs imposed by parallel computing.
- 13 We may consider a necessarily non-exhaustive set of factors such as wealth and other resources, socio-demographics, stable preference attributes, technology, and to some extent even less tangible characteristics such as cultural and psychological biases as being of relevance in this context.
- 14 The so-called *curse of multi-dimensionality* represents a well-known problem in both the economic and the physical sciences. It is particularly relevant for finite-state or discrete problems that are amenable to computer simulation and analysis. The problem manifests itself specifically in the explosion of the state space with the addition of each additional dimension to any problem under scrutiny.
- 15 Threads in computer hardware design are lightweight processes which can run in parallel. They typically also share some minimum amount of memory through which they can communicate most efficiently. This particular feature is however not supported in *Erlang* and the actor model's implementation of concurrent agent behaviour; agents can only share information states by sending and receiving messages which are copied. They are *not* provided as memory pointers.
- 16 I would like to thank David Emanuel Andersson, Nicolai Petrovsky and an anonymous referee for useful comments. All remaining errors are my own.

REFERENCES

- Anderson, C. M., Plott, C. R., Shimomura, K.-I., and Granat, S. (2004). Global instability in experimental general equilibrium: the Scarf example. *Journal of Economic Theory*, 115: 209-49.
- Andersson, D. E. (2008). The double-edged nature of the Hayekian knowledge problem: systemic tendencies in markets and science. *Studies in Emergent Order*, 1: 51-72.
- Armstrong, J. (2007). A history of Erlang. In: Proceedings of the Third ACM SIGPLAN Conference on History of Programming Languages, New York, 616-26.
- Arrow, K. J. (1952). An extension of the basic theorems of classical welfare economics. Cowles Commission for Research in Economics, The University of Chicago, Chicago.
- Ashraf, Q., Gershman, B., and Howitt, P. (2012). How inflation affects macroeconomic performance: an agent-based computational investigation. National Bureau of Economic Research, Inc.
- Bailey, D. H., Borwein, P. B., and Plouffe, S. (1997). On the rapid computation of various polylogarithmic constants. *Mathematics of Computation*, 218: 903-13.
- Benassy, J.-P. (1990). *Non-Walrasian Equilibria, Money, and Macroeconomics*. Amsterdam: Elsevier.
- Bridel, P. (2011). *General Equilibrium Analysis: A Century after Walras*. London: Routledge.
- Colander, D. (1996). *Beyond Microfoundations: Post Walrasian Economics*. Cambridge: Cambridge University Press.
- Colander, D. (2006). *Post Walrasian Macroeconomics*. Cambridge: Cambridge University Press.
- Colander, D., Howitt, P., Kirman, A., Leijonhufvud, A., and Mehrling, P. (2008). Beyond DSGE models: Toward an empirically based macroeconomics. *American Economic Review*, 98: 236-40.
- Constantinides, G. M. (1990). Habit formation: a resolution of the equity premium puzzle. *Journal of Political Economy*, 98(3): 519-43.
- Debreu, G. (1954). Existence of an equilibrium for a competitive economy. *Econometrica*, 22: 265-90.
- Duesenberry, J. S. (1949). *Income: Saving, and the Theory of Consumer Behavior*. Cambridge, MA: Harvard University Press.
- Gintis, H. (2007). The dynamics of general equilibrium. *Economic Journal*, 117: 1280-1309.
- Hayek, F. A. (1945). The use of knowledge in society. *American Economic Review*, 35: 519-30.
- Hayek, F. A. (1979/2012). *Law, Legislation and Liberty: A New Statement of the Liberal Principles of Justice and Political Economy*. London: Routledge.
- Hewitt, C. (2010). Actor model of computation: Scalable robust information systems. *arXiv*, 1008:1459.
- Hewitt, C., Bishop, P., and Steiger, R. (1973). A universal modular ACTOR formalism for artificial intelligence. In: *Proceedings of the 3rd International Joint Conference on Artificial Intelligence*. San Francisco: Morgan Kaufmann, 235-45.
- Hildenbrand, W. (1994). *Market Demand: Theory and Empirical Evidence*. Princeton: Princeton University Press.
- Hirota, M. (1981). On the stability of competitive equilibrium and the patterns of initial holdings: an example. *International Economic Review*, 22: 461-67.
- Hirota, M. (1985). Global stability in a class of markets with three commodities and three consumers. *Journal of Economic Theory*, 36: 186-92.
- Hirsch, M. W., Smale, S., and Devaney, R. L. (2012). *Differential Equations, Dynamical Systems, and an Introduction to Chaos* (3rd Edition). Waltham, MA: Academic Press.
- Holt, R., Rosser, J.B., and Colander, D. (2011). The complexity era in economics. *Review of Political Economy*, 23: 357-69.
- Judd, K. L. (1998). *Numerical Methods in Economics*. Cambridge, MA: MIT Press.
- Kirman, A. P. (2010a). *Complex Economics: Individual and Collective Rationality*. London: Taylor & Francis.
- Kirman, A. P. (2010b). *Walras' Unfortunate Legacy*. Shanghai: HAL.
- LeBaron, B. (2001). A builder's guide to agent-based financial markets. *Quantitative Finance*, 1: 254-61.
- Mantel, R. R. (1974). On the characterization of aggregate excess demand. *Journal of Economic Theory*, 7: 348-53.
- Radner, R. (1968). Competitive equilibrium under uncertainty. *Econometrica*, 36(1): 31-58.
- Scarf, H. E. (1959). Some examples of global instability of the competitive equilibrium. Cowles Foundation for Research in Economics, Yale University, New Haven.
- Simon, H. A. (1947). *Administrative Behavior: A Study of Decision-Making Processes in Administrative Organization* (1st ed.). New York: Macmillan.
- Smale, S. (1976). Dynamics in general equilibrium theory. *American Economic Review*, 66: 288-94.
- Sonnenschein, H. (1973). Do Walras' identity and continuity characterize the class of community excess demand functions? *Journal of Economic Theory*, 6: 345-54.
- Taylor, P. and Jonker, L. (1978). Evolutionarily stable strategies and game dynamics. *Mathematical Biosciences*, 40: 145-56.
- Turing, A. (1936). On computable numbers with an application to the 'Entscheidungsproblem'. Proceedings of the London Mathematical Society, London.
- Veblen, T. (1899/2007). *The Theory of the Leisure Class*. Fairford: Echo Library.
- Walras, L. (1874/1954). *Elements of Pure Economics: Or the Theory of Social Wealth*. London: Taylor & Francis.
- Zuse, K. (1970). Calculating space. Project MAC, Massachusetts Institute of Technology, Cambridge, MA.