D'Arms, J, Robert Batterman, and Krzyzstof Górny. 1998. Game theoretic explanations and the evolution of justice. *Philosophy of Science* 65: 76–102.
Kitcher, P. 1999. Games social animals play: commentary on Brian Skyrms' *Evolution of the social contract*. *Philosophy and Phenomenological Research* 59(1): 221–8.
Skyrms, B. 1996. *Evolution of the social contract*. Cambridge University Press.
Skyrms, B. 2004. *The stag hunt and the evolution of social structure*. Cambridge University Press.

doi:10.1017/S0266267106221126

Cognitive Economics. An Interdisciplinary Approach, Paul Bourgine and Jean-Pierre Nadal, eds. Springer, 2004, xiv + 479 pages.

'Cognitive Economics' is a newcomer to economic research. As of now, only a few publications bear its name, and it saw its first European and its first international conference in 2004 and 2005 respectively. This book, carrying the new subdiscipline's name as its title, collects 27 articles from fields as diverse as economics, artificial intelligence, logic, psychology and physics. With many of the articles being surveys, the book serves both as an introduction to the field – nine essays explicitly cover the 'disciplinary bases for cognitive economics' – as well as a 'tool for future research'. The anthology fulfils these two purposes well. Researchers interested in bounded procedural rationality, social influence on individual decision-making and the dynamics of adaptive social systems can learn modelling techniques from outside their fields, and they are offered a wealth of suggestions on how to apply them fruitfully to economic problems.

Cognitive Economics, the editors of this book suggest, is a unified research program that brings a *cognitive turn* to economics. 'It aims to take into account the cognitive processes of individuals in economic theory, both on the level of the agent and on the level of their dynamic interactions and the resulting collective phenomena' (Bourgine and Nadal, v). Now, mainstream economics also takes into account agents' cognition, attributing preferences and beliefs, and modelling deliberation under uncertainty or incomplete information. However, these models rely on two strong assumptions. First, agents are assumed to be *substantively rational*: they deliberate in whatever necessary way to arrive at an optimal choice. Second, the equilibria are presumed to emerge directly from the agents' reasoning. No concrete equilibration process leading to the coordination state is modelled. Rejecting these assumptions, Cognitive Economics focuses on the agents' cognitive constraints, and the deliberative procedures resulting from these limitations. Because agents' information processing capacities are limited, substantive rationality in the sense of universal optimising capacities is excruciatingly costly or simply unreachable. Instead, agents employ *cognitive procedures* that yield at least good-enough results for specific environments. Which specific procedures the agent

CAMBRIDGE JOURNALS
http://journals.cambridge.org Downloaded: 23 May 2011 IP address: 128.214.51.2

employs will strongly influence her behaviour in changing environments; hence the procedures themselves and the way agents acquire them through different kinds of learning will be of interest to economists.

Of course, cognitive science has pursued this type of research for more than four decades. Economists have ready access to these results, and do not need to replicate the research. But cognitive science is often criticised for failing to model human thought as inherently social, and here Cognitive Economics makes its central contribution. Boundedly rational agents in general cannot coordinate their actions with the actions of others in such a way that the optimal equilibrium is instantly reached. Recent research, instead, has investigated possible *trajectories* towards equilibrium in these cases – agents *learn* to adjust their behaviour in repeated interaction in order to achieve a social optimum. But these learning models are based on individual rationality: each agent has the optimum as her goal, and searches (with limited capacities) for information so as to adjust her behaviour to best reach this goal (Kalai and Lehrer 1993; as they point out, their model presents learning not as a goal in itself, but as an implication of utility maximisation). Cognitive Economics, in contrast, develops models where agents strictly rely on certain deliberation and learning rules, without having the overall optimum as a goal in mind, but where aggregate behaviour still converges towards this optimum. Cognition thus is not solely cognition of the individual. In fact, 'individual behaviour. . .is not [Cognitive Economics'] main subject of interest' (Walliser, 196). Instead, Cognitive Economics often models cognition as *distributed*: as information processing distributed over a large number of individuals, who interact in social networks, and influence each other. Consequently, it studies economies as *complex adaptive systems*, and investigates their stability conditions, adaptation dynamics and equilibrium paths.

The book is divided into three parts: three programmatic essays, nine introductory essays and 15 essays on areas of advanced research. The economic introductory essays (essays 2-4) mainly review textbook material, and focus surprisingly little on issues of interest here, be it non-expected-utility, epistemic justifications of game equilibria or models of learning. Makinson's essay on non-monotonic reasoning (essay 6) discusses qualitative logics that allow inferring more conclusions from a set of premises than is classically authorised. His essay provides a wealth of structure that modellers of individual reasoning may find very useful, and it also clarifies the relation of these structures with classical logic. Unfortunately, no other essay in the book makes use of these results (a fate that this essay shares with essay 13 by the same author on conditional statements and directives). In particular, it is a pity that the relation of qualitative logics to logics of belief change remains unexplored, despite their importance for decision and game theory.

CAMBRIDGE JOURNALS
http://journals.cambridge.org Downloaded: 23 May 2011 IP address: 128.214.51.2

Alexandre and Frezza-Buet (essay 7) give an overview of several classes of numerical AI models used for modelling human cognitive abilities. Genetic Algorithms (GA) support determining optimal cognitive procedures for a well-defined search space by simulating an evolutionary process (these may be familiar from evolutionary game theory). Artificial Neural Networks (ANN) are often used to model associative learning. They consist of functioning rules that define the computation performed by the network's units, learning rules that specify how the units' and network's parameters are adapted as a result of learning, and the architecture that defines the way units are connected. Cases of ANN where agents are modelled as units and learning as social influence will be discussed below. Stochastic Behavioural Models, and more specifically Markovian Decision Processes (MDP), allow modelling complex reinforcement learning, where reward is delayed. All these models are part of the numerical paradigm. They stand in contrast to symbolic models, exemplified in standard decision theory, belief revision, indeed the whole propositional attitude tradition, which uses systems of symbol manipulation to model cognitive processes. The authors contend that numerical techniques are 'better adapted to such fields as economics, where expertise and knowledge are too often unconscious and hard to formalise precisely' (Alexandre and Frezza-Buet, 114). Unfortunately, they do not provide more arguments for this interesting but controversial claim.

The research topics section covers individual deliberation, market dynamics and social networks. Starting with individual deliberation, Orléan (essay 12) proposes a concept of collective belief that focuses on the group as an autonomous entity. To believe that group G believes p, according to this proposal, means to believe that the majority of G's members believe that group G believes p. As the author notes, such a concept of collective belief has been fruitfully applied to pure coordination games. Here, each agent faces the problem of identifying the *salient* equilibrium out of many equilibria that every player can identify as such. According to the proposed concept of collective belief, an agent A chooses equilibrium E because A believes that all players believe that the group believes in the salience of E – not because A believes that all players believe in the salience of E on the basis of their beliefs in the others' individual tastes and believes. Choosing on the basis of this kind of collective belief, A has an advantage. It is far more plausible that cultural traits and group identities are common knowledge (think of stereotypes) than that individual tastes and beliefs are. Hence, players choosing on the basis of their beliefs about what the group believes will lead more readily to a coordinated result than players choosing on their beliefs about what all the other players' individual beliefs are. Orléan then proposes to transfer the concept of collective belief to investment decisions in financial markets. In financial bubbles, he claims, a 'strange and enigmatic'

CAMBRIDGE JOURNALS
http://journals.cambridge.org Downloaded: 23 May 2011 IP address: 128.214.51.2

disconnection between individual and collective beliefs occurs. Investors may individually believe that an asset is overvalued, but continue to buy, because they believe that the market will continue to rise. Employing the new concept of collective belief, Orléan suggests, helps to explain the bubble without having to assume the presence of irrational agents.

However, the avoidance of irrationality comes at a price. Financial markets, after all, are not games of pure coordination. In coordination games, collective beliefs are stable because they are self-enforcing. Once agents come upon a collective belief that supports a coordinated result, they have no reason to deviate from such successful collective beliefs. Such a self-reinforcement does not exist in financial markets – bubbles eventually burst, destabilising any collective belief that led to their existence. To invoke collective beliefs in explaining financial market dynamics means invoking common knowledge of p until p isn't common knowledge anymore. Such an explanatory strategy remains *ad hoc* until an explicit model of the dynamics of such a collective belief is provided.

Tallon and Vergnaud (essay 14) develop a non-standard expected utility (EU) model that does not require the sure-thing principle but satisfies the requirement that the decision maker positively values information. They follow Hammond's strategy of justifying the axioms of their EU model by deducing them from axioms of dynamic choice, but relax consequentialism and instead derive their model from *separability* and *selection of optimal strategies* (SOS). Whereas consequentialism requires that an agent's choices are identical in a decision tree and its strategic equivalent form, SOS only requires that the choice in the decision tree is a subset of the optimal strategies in the equivalent strategic form. The 'weak sure thing principle', derived this way, only requires that if an agent prefers betting on $A\grave{E}C$ to betting on $B\grave{E}C$, and both A and B are disjoint with C, then the agent also prefers betting on $A\grave{E}D$ to betting on $B\grave{E}D$ for all other events D that are disjoint with both A and B. The authors further show that with their axioms, the agent always has a positive value of information – a conclusion that distinguishes their model from other non-standard expected utility models. It is difficult, however, to see the relevance of this weakened sure thing principle. The authors point to a family of models that employ possibility measures for a qualitative description of uncertainty, which violate the standard sure thing principle while satisfying its weak form. The real point of contention, though, which set off the whole debate about the principle in the first place – Ellsberg's Paradox – remains a powerful counterexample to the weakened sure thing principle.

Turning to markets, Kirman (essay 18) argues that rational collective behaviour often cannot be directly related to rational behaviour of the collective's individual members. At the example of the Marseille fish market, he shows how the aggregate demand exhibits a standard downward sloping relationship between prices and quantities of fish

CAMBRIDGE JOURNALS
http://journals.cambridge.org Downloaded: 23 May 2011 IP address: 128.214.51.2

transacted, while the demand curves of the studied individual buyers does not exhibit this standard property. Thus, the aggregate data can be rationalised under the standard rationality axioms, while individuals constituting the aggregate do not behave rationally. To explain the macrophenomenon, Kirman concludes, one cannot employ a 'blown up version' of the microbehaviours. Instead, he suggests an analogy between human institutions and a beehive or an ant's nest. There, individual ants' or bees' cognitive abilities are strictly limited. They operate in a restricted neighbourhood, obtaining most of their information from those with whom they interact. Despite the simplicity of their behaviour and their reasoning rules, however, the aggregate outcome of their behaviours is surprisingly sophisticated. Following the analogy, Kirman suggests modelling the economy as a complex system, drawing on techniques from statistical mechanics.

Thankfully, the book provides comprehensible introductions (essays 8 and 9) to the most important class of these, the Ising model. Originally developed for the explanation of ferro-magnetism, this model allows inferring interesting properties of a system that cannot be deduced from the bare properties of the system's components. This invites its application to social phenomena, as long as the system in question can be described by something structurally equivalent to the energy function. Phan *et al*. (essay 20) survey and extend some microeconomic models that use the Ising model to investigate *social influence* on individual decisions. The standard model of an agent's discrete choice is turned into an Artificial Neural Network (ANN) by including a social influence component as an additive element to the private utility component. One such model they discuss takes the agents' utility as the equivalent to the energy function of the original Ising model, and models social influence as the agent's adaptive expectations of her neighbours' choices. All agents i simultaneously maximise V with their binary choice (buy, $\omega_i = 1$, not buy $\omega_i = 0$):

$$V_i = \max \omega_i \left(h_i + J_\vartheta \sum_{k \in \vartheta} \omega_k - p \right)$$

where h_i represents the individual preference of the agent, J_ϑ is the social influence factor from the agents in the neighbourhood ϑ, and p is the price of one unit. That is, their individual choice makes V_i positive if the agent buys and null otherwise. This model shows interesting 'avalanche' effects. Take for example an incremental price decrease. This will make some additional agents buy – namely those whose individual evaluation $h_i - p$ has changed given the price change. But those who choose to buy for that reason will change the situation in their neighbourhood ϑ, changing the social influence on some of the neighbours. Some of those influenced will consecutively also buy, further changing the social influence and possibly

CAMBRIDGE JOURNALS
http://journals.cambridge.org Downloaded: 23 May 2011 IP address: 128.214.51.2

triggering significant change in the whole population. The model's capacity to capture such processes are of great interest; however, the results that the authors present require quite restrictive assumptions on the models, e.g. symmetry of social influence, same size of neighbourhoods, and very specific distributions of preferences.

The models discussed assume that all agents are connected to a local neighbourhood of homogenous size. In his survey of Agent-based Computational Economics (ACE), Phan (essay 22) shows how this restricting assumption can be relaxed by introducing *social networks*. A social network represents the interconnectedness of a population of agents, specifying each agent's 'neighbour' with the help of a graph. A network is called *regular*, if each agent is connected to his closest neighbours. Through increased random replacement of connections, a network loses its regularity in degrees. Neighbourhoods of complex social systems can be modelled as such social networks. Their stability against external shocks or entropic disturbance, and their dynamic behaviour out of equilibrium, depends on the network's degree of regularity. Through simulations, ACE explores these connections between network regularity, the system's stability conditions and its out-of-equilibrium dynamics.

Zimmermann (essay 23) discusses a further expansion of these systems. He envisages a notion of social learning that goes beyond the ability of agents to adjust in the light of their neighbours' influence. It allows the agent to revise the existence and strength of her neighbourhood links as a consequence of the evolving degree of affinity she feels for, or credibility she accords to, her different neighbours. Agents reallocate their 'closeness' to those neighbours who most frequently have agreed with them in the past. Simulation of an evolving network starting from randomly drawn closeness connections after 10.000 steps of learning then yield an interesting result: a very small number of agents have the power to trigger large avalanches at the level of the whole population. 'Expert leaders' have emerged due solely to their structural position, as the result of a social process.

This ambitious book contains many more interesting chapters on viability theory, stochastic game theory, the evolutionary analysis of communication, social influence in social choice, strategic models of coalition and network formations, a dynamic voter behaviour model in a population with bimodal conflicting interests and a discussion of cognitive efficiency of social networks, which unfortunately cannot be covered here. Instead, a few critical remarks are in order.

The programmatic essays make promises that the later chapters do not bear out. For one, they claim that Cognitive Economics studies *both* adaptation *and* reasoning processes implemented by economic agents in their interactions. The claim that 'it conserves both the time of evolution and the time of eduction' (Bourgine 2) suggests a synthesis of the two

CAMBRIDGE JOURNALS
http://journals.cambridge.org Downloaded: 23 May 2011 IP address: 128.214.51.2

approaches. But the book offers no such synthesis; rather, it puts numerical and symbolic models of cognitive processes side by side, and the many chapters treating complex interactive systems focus on numerical models alone. Of course, it may well be that numerical models also are capable of modelling reasoning, but little discussion beyond a brief proclamation of Alexandre and Frezza-Buet can be found. In fact, one author warns that 'cognitive economics, which provides powerful models separately in an eductive and an evolutionist perspective, fails at this time to provide an integrated analytic framework of reference' (Phan 393).

Further, the introductory essays claim that Cognitive Economics can be empirically justified: 'cognitive economics is not armchair economics. The links between cognition, evolution and institutions must be tested by means of field surveys, laboratory experiments, computer simulation and the analysis of models' (eds., vi). Consequently, the book offers two chapters on experimental studies. Politzer (essay 5) surveys the relatively well-known shortcomings in individual reasoning and decision-making, and he cautions about the methodological soundness of many experiments of this sort. The results he surveys are clearly an inspiration for cognitive economists, but do not show how the cognitive economists' own models can be put to the test. Noussair and Ruffieux (essay 19) survey experimental research on markets. Again, the behavioural patterns they report invite the construction of new models; but they say little about how such models of complex interactive systems can be put to the test. In the concluding essay, Lesourne addresses this issue: 'Having to describe complex stochastic processes, the model builders are compelled to introduce numerous assumptions concerning the sequence of events, the way in which information is drawn, the data concept in memory, the size of adaptations, etc.... There is a risk of multiplying ad hoc models based arbitrarily on debatable assumptions' (Lesourne, 468). It is helpful to contrast Cognitive Economics with Behavioural Economics in this regard. Both programmes acknowledge that economics rests on *some* sort of implicit psychology; both relax simplifying assumptions for greater psychological realism and modify other assumptions to acknowledge human limits on computational power, willpower and self-interest. Behavioural economics, however, constructs models that are 'generalizations of standard ones' (Camerer and Loewenstein 2003: 47). Cognitive Economics, as the book portraits it, is quite willing to accept substantial deviation from the standard models. Further, Behavioural Economics justifies introducing psychological assumption as an improvement of economics *on its own terms*: 'The ultimate test of theory is the accuracy of its predictions' (Camerer and Loewenstein 2003: 5). Cognitive Economics concentrates more on *plausible* models of cognition, at the price of less empirical testability and verification – as expressed in one of the programmatic essays: 'the aim is not so much to explain certain realized phenomena as

CAMBRIDGE JOURNALS

http://journals.cambridge.org Downloaded: 23 May 2011 IP address: 128.214.51.2

to show that certain phenomena are possible' (Walliser, 196). This need not be a disadvantage, and indeed may be a necessity in order to develop the new programme – but it should be made clear that the models are not well empirically founded at the current stage.

Lastly, many of the essays would have greatly benefited from a scrupulous proof reading. Repeatedly, sentences are fragmented, graph and section references are mismatched, and bibliographic references are incomplete. This can make for a frustrating read.

All in all, this anthology gives insight into a fascinating research area. It confronts a cognitive science approach with the explanatory aims of the social sciences, and for this purpose presents interesting novel modelling techniques. It will be of interest to a wide range of researchers from cognitive science, economics, the social sciences and AI.

Till Grüne-Yanoff
Royal Institute of Technology, Stockholm

REFERENCES

Kalai, Ehud and Ehud Lehrer. 1993. Rational learning leads to Nash Equilibrium. *Econometrica* 61(5): 1019–45.

Camerer, Colin F. and George Loewenstein. 2003. Behavioral economics: past, present, future. In *Advances in behavioral economics*, ed. Colin F. Camerer, George Loewenstein, and Matthew Rabin. Princeton University Press.

doi:10.1017/S0266267106231122

Just Work, Russell Muirhead. Harvard University Press, 2004, 209 pages.

As its title suggests, *Just Work* is an attempt to articulate an account of the justice of work. But the title is not just descriptive, it is also imperative – just work! – and this leads Muirhead to examine the question, why we work. Muirhead argues that although Americans work for the instrumental reasons of monetary and material sustenance this cannot entirely explain our working life.[1] For besides how much money we make, we also evaluate our work in terms of how it fits us – the way it brings meaning to our lives through developing our talents and capacities. For Muirhead, the justice of work requires not just that we fulfill socially useful roles, for example teaching as opposed to stealing, but also that our work personally fits us in a certain way.

[1] This book is specifically directed to Americans and American working culture, but it draws on sufficiently broad concepts of liberal democracies in general to be accessible and interesting to a wider audience.

CAMBRIDGE JOURNALS
http://journals.cambridge.org Downloaded: 23 May 2011 IP address: 128.214.51.2