International Journal of Automation and Computing 05(1), January 2008, 63-66
DOI: 10.1007/s11633-008-0063-6

Computational Intelligence Determines Effective Rationality

Edward P. K. Tsang
Centre for Computational Finance and Economic Agents, University of Essex, Colchester CO4 3SQ, UK

Abstract: Rationality is a fundamental concept in economics. Most researchers will accept that human beings are not fully rational. Herbert Simon suggested that we are "bounded rational". However, it is very difficult to quantify "bounded rationality", and therefore it is difficult to pinpoint its impact to all those economic theories that depend on the assumption of full rationality. Ariel Rubinstein proposed to model bounded rationality by explicitly specifying the decision makers′ decision-making procedures. This paper takes a computational point of view to Rubinstein′s approach. From a computational point of view, decision procedures can be encoded in algorithms and heuristics. We argue that, everything else being equal, the effective rationality of an agent is determined by its computational power – we refer to this as the computational intelligence determines effective rationality (CIDER) theory. This is not an attempt to propose a unifying definition of bounded rationality. It is merely a proposal of a computational point of view of bounded rationality. This way of interpreting bounded rationality enables us to (computationally) reason about economic systems when the full rationality assumption is relaxed.

Keywords: Rationality, bounded rationality, computational intelligence, economics, computational intelligence determines effective rationality (CIDER) theory.

1 Introduction: What is rationality

Many economic theories are built upon the assumption that decision makers (people or institutes) are perfectly rational. Being rational means being able to maximize one′s utility, given all the available information. Here we assume that utilities can be quantified; whether this is true or not it does not affect the argument in this paper.

Suppose a merchant receives two offers to buy one of his commodities: one offers £10 and the other £20. With everything else being equal, the merchant, being "rational", will sell it for £20.

Full rationality is an assumption behind all major economic theories. For example, the efficient market hypothesis states that security prices always fully reflect the information available[1]. Should this hypothesis hold, no investor can consistently beat the market; this means that no investor can expect to consistently get a return on its investment higher than the return that the market can offer. This hypothesis has many consequences in financial analysis, and therefore is very important in economics. For a market to be "efficient", among other assumptions, a sufficient number of the investors in the market must be rational. In recent years, the efficient market hypothesis has become under serious scrutiny, both theoretically (e.g. see [2, 3]) and empirically (e.g. see [4]). The rationality assumption is seriously questioned by researchers in behavioural finance, which studies the cognitive and emotional biases on investors and their impact on investment decisions[5]. In the rest of the paper, we examine what rationality really means, with focus on the computation aspect of decision making.

2 Rationality involves computation

It is easy to choose between selling an item for £10 or £20, everything else being equal. Suppose a merchant is offered the choice between 1) receiving a payment of £100 today, and 2) receiving a payment of £10 per month over 12 months. Which option should he take? With basic mathematical and finance training, and the knowledge of his cost of capital, a "rational" merchant should have no problem choosing between the two offers.

The above example highlights the fact that knowledge is required in making certain "rational" decisions. In fact, "being rational" requires more than basic knowledge. It also requires computation. In the above example, the calculation is relatively simple; it can be performed on a simple calculator.

Let us turn to another scenario. Supposed a merchant has to visit his customers, who are located far apart. Travelling from one customer to another involves a cost, which may vary depending on the time and distance to travel. Some customers may not be available at all times. Suppose the merchant wants to plan an itinerary that visits all his 100 customers, with the objective to minimize travelling costs and satisfying all the customers′ availability constraints.

A "rational" merchant would attempt to find the optimal itinerary in the above problem. The amount of computation required to find the optimal itinerary in this problem is nontrivial. This is a complicated version of a problem known as the "travelling salesman problem", which has been studied extensively in operations research and computer science[6].

Clever heuristics have been invented to tackle the travelling salesman problem (e.g. see [7, 8]), but they typically involve serious computation. The problem is in nature nondeterministic polynomial-time hard (NP-hard), which means that the time required to find the optimal solution grows exponentially as (in our example) the number of customers increases. Given a fixed amount of planning time, one may not be able to find the optimal itinerary (i.e. an itinerary with the minimal travelling cost). In that case, one would have to settle for the best itinerary found within the given time.

It is also worth noting that computation itself involves a cost. Knowledge acquisition (e.g. to find out the travelling costs between two cities) could also involve costs. A rational

Manuscript received September 3, 2007; revised November 12, 2007
E-mail address: edward@essex.ac.uk

agent should not only minimize travelling cost. It should attempt to minimize the travelling cost plus the cost of computation and knowledge acquisition.

3 Bounded rationality

But then what does "being rational" mean? One is not perfectly rational if one cannot find the optimal itinerary. Simon[9] pointed out that most people are only partially rational. He suggested to describe human as "bounded rational", which means they can only make the best decisions within their knowledge and resources. Although most economists would accept that perfect rationality is not a realistic assumption, it is not clear how most of the economic theories should be revised to reflect bounded rationality. Concretely quantifying what bounded rationality means remains a grand challenge to the research community.

Many have attempted to study the interpretation or implication of bounded rationality. Some investigate the psychological aspect of rationality (e.g. see [10]). To enable one to quantify an agent′s rationality, it would be useful if one could provide a mathematical definition of bounded rationality; for example, "agent A is 86% bounded rational". Then one may be able to revise economics theories to reflect the level of bounded rationality. Unfortunately, no such definition has been widely accepted.

As decision makers have to make decisions about how and when to take what actions, Rubinstein[11] proposed to model bounded rationality by explicitly specifying decision making procedures. This is an attractive approach, as it enables one to study the consequences of bounded rationality. Rubinstein′s proposal puts the study of decision procedures on the research agenda. This is the view that we shall follow in this paper.

4 A computational point of view in bounded rationality

From a computational point of view, decision procedures can be encoded in algorithms and heuristics. One way to study the impact of relaxing the full rationality assumption is to study the impact of adopting different algorithms and heuristics in the agents. One important impact that one can study is the equilibrium of a situation when the algorithms and heuristics are specified for the players. One can also attempt to study the equilibrium of a market given a model of the agents′ algorithms and heuristics. It is worth noting that these studies are nontrivial and not always feasible. The studies themselves may involve computational intelligence too. We shall look at some case studies later in the paper.

So far, we have argued that specifying the algorithms and heuristics enables one to study the equilibrium in the market. It is important to point out that the search for algorithms and heuristics itself is interesting: it is interesting from an individual′s point of view, if one wants to find out what algorithms and heuristics will succeed in a given market situation.

Because of combinatorial explosion, finding the optimal algorithms and heuristics is out of reach in most realistic problems. When optimal solutions are out of one′s reach, one′s knowledge of algorithms, heuristics and our computational power determine effectively how good a solution one can find. For example, in the travelling salesman problem, some algorithms and heuristics find better solutions than others.

If one defines full rationality as "being able to find the optimal decisions in every situation", then it is reasonable to say that the "level of optimality" that one achieves defines one′s effective rationality. Thus, designing better algorithms and heuristics helps to extend the rationality boundary. In other words, computational intelligence determines one′s effective rationality – we refer to it as the CIDER theory.

5 Where do decision procedures come from?

So far, we have not asked the question where decision procedures come from. Procedures can be designed, as it is the case in many disciplines of computational intelligence. They can also be evolved, as it is the case in evolutionary computation. In this section, we shall briefly outline some of the relevant disciplines in computational intelligence.

The general problem solver (GPS) was an early attempt in artificial intelligence to mimic human intelligence[12]. The idea is to separate domain-specific knowledge from the reasoning mechanism. GPS is designed as a general reasoning mechanism. GPS opened the field of artificial intelligence planning, which is still an on-going research area in artificial intelligence[12, 13]. Planning involves knowledge representation (how to represent beliefs, actions and their effects), causal reasoning (reasoning about actions and their consequences), and resources allocation (primarily time resources, i.e. when to perform which action).

Rationality is often studied in the context of decision problems. Finite choice decision problems are tackled in constraint satisfaction, a discipline that brings together research in artificial intelligence, logic programming and operations research[14, 15]. Constraint satisfaction is a general problem which appears in practical problems such as industrial scheduling. Search algorithms have been designed to use constraints to find solutions efficiently. Procedures implementing these algorithms could be used to model rationality.

Human beings find strategies by iterative improvements. Starting from their current situations, they look at possible changes or experiment with them. They change their strategies in response to the anticipated or actual changes. Research that fall into this pattern are called local search[16, 17]. It would be reasonable to model human rationality by local search procedures.

Human beings learn from their experiences. Therefore, it would be reasonable to model rationality with dynamic (as opposed to static) procedures. Evolutionary algorithms (including estimation of distribution algorithms[18]) attempt to evolve solutions instead of designing them[19, 20]. It would be a reasonable idea to model human reinforcement learning with evolutionary computation procedures, or procedures generated by evolutionary computation.

6 Case study: Evolutionary computation in game theory

In this section, we present a case study to demonstrate that bounded rationality can be reasoned when procedures are specified. This case study is the two-player alternating-offers bargaining scenario in the field of game theory[21]. In this scenario, the players make alternative offers on how much (in terms of percentage of the overall resources) they demand on the resources. Both players have incentive to compromise as their utilities drop over time.

The traditional approach to game theory is to derive subgame equilibrium mathematically[22]. Perfect rationality is assumed in such derivations. Jin and Tsang[23] relax this assumption, and attach procedures to the players (as suggested by Rubinstein[11] in Section 3). This enables them to study subgame equilibrium under the given procedures[23].

The procedures used by Jin and Tsang[23] are constraint-directed genetic programming procedures (a branch of evolutionary computation). Constraint-directed genetic programming is novel in its own right, but in this paper, we focus on the fact that the constraints that Jin and Tsang[23] used implement "common sense" that one would expect a human player to use. Examples of these constraints are: 1) one does not ask for more than 100% of the resources; and 2) the faster one′s utility drops, the less aggressive one would be in bargaining[24]. Another aspect of Jin and Tsang′s work is that the players co-evolve[23, 24]. Co-evolution implements reinforcement learning or arm-races, which is quite common in human society.

In many game models where mathematical results have been derived, subgame equilibrium produced by the evolutionary approach produced very similar results[23−25]. This means, at least for these simple game models, the perfect rationality assumption has not made any difference to the equilibrium. It would be interesting to see whether this is still the case in more complex game models.

This case study support′s Rubinstein′s proposal: it demonstrates that once the decision procedure is specified, one can study the subgame equilibrium without assuming perfect rationality. In this case, the decision procedure and the constraints that were adopted implement realistic behaviour in human societies.

7 Case study: Computational intelligence determines agent performance

In this section, we use a few examples to demonstrate that market equilibrium can be studied through procedural attachments. Besides, we demonstrate that procedures can be learned. Furthermore, the algorithm that an agent uses determines how successful it could be.

Selten[26] pointed out that quantitative reasoning is typically infeasible. To better understand bounded rationality, Selten[26] proposed to better understand the "structure" used by business and public administration. He promoted the use of qualitative reasoning. Special attention was paid to causal reasoning, in the form of causal diagrams, which are directed graphs with edge labels.

Alexandrova et al.[27, 28] modelled the card payment market in both the structure and the quantitative implications. An agent-based approach is used to study the card payment market. The agents have their decision process defined. This allows Alexandrova to study market equilibrium. In this case, the decision procedure is evolutionary computation using an extended population-based incremental learning (PBIL)[29, 30].

Alexandrova[28] also studied the market equilibrium when the agents do and do not evolve. Results showed that, unsurprisingly, agents that evolve out-perform those who do not[28]. This supports the argument that computational intelligence determines effective rationality (under our definition in Section 4 that optimality is a measure of rationality).

Alexandrova′s results are also echoed by related research. Gosling and Tsang[25, 31] modelled middlemen strategies in a supply chain. They demonstrated that effective middlemen strategies could be evolved. Martinez and Tsang[32, 33] modelled agents in a stock market. By specifying procedures in different agents (some of which being stationary and some evolutionary), one is in a position to study the market behaviour. This enables one to identify conditions under which the market behaviour resembles real financial markets. Martinez and Tsang also showed that trading strategies can be evolved to beat naive strategies, which supports the CIDER theory.

8 Remarks: Constraint satisfaction

Rubinstein[11] pointed out that, in reality, agents do not necessarily try to find an optimal decision. An agent′s task often involves picking from a finite set of options a decision that satisfies all the constraints. In an enterprise, an agent′s action will restrict the actions that others can take. In a centralized decision problem, the coordinator has to decide for all its agents their choices of decision, satisfying all the constraints. This is known as a constraint satisfaction problem in the literature[15].

The task of "finding any set of decisions that satisfies all the constraints" is typically easier than "finding the optimal set of decisions", especially when many solutions exist. This may have significant impact to the computation procedures (see [14] for a catalogue of constraint satisfaction tools). However, it has no fundamental impact to the analysis above: the problem is still computationally hard. Given a particular problem, some algorithms are more efficient than others in finding solutions – sometimes by several orders of magnitude in terms of solving time. Knowledge of what algorithms to use under what situations could make all the difference, as pointed out in Section 2. Thus, the CIDER theory still applies.

9 Conclusions

Rationality is a fundamental assumption in economics. With the full rationality assumption relaxed, many economic theories must be revised. The applicability of many economic theories to the society must be scrutinized. Unfortunately, we are not in a position to revise and scrutinize these theories, as we do not know how to assess the impact of relaxing this assumption. Modelling bounded rationality by procedures is one way forward. It allows one to study the impact of different procedures. While one′s model of individuals′ procedure may not be perfect, modelling allows one to scientifically study the impact of different procedures. By adjusting the procedures, one may attempt to model the underlying mechanism that drives the economy – a methodology similar to that used in physics. In this paper, we have given concrete examples on how bounded rationality could be modelled by evolutionary procedures, which are basically reinforcement learning procedures. It is reasonable to believe that reinforcement learning is what ordinary people use in daily life decision making. We have also argued that computational intelligence determines an agent′s effective rationality.

To summarize, attaching procedures to decision making is a practical way forward in advancing bounded rationality research. High on the research agenda is how to design or evolve these procedures and how to evaluate their impact. Computational intelligence could play a big part in answering these questions.

Acknowledgements

This research has benefited from my co-workers and colleagues. The work on evolutionary game theory was a joint effort with Nanlin Jin, with advices from Abhinay Muthoo. The work on card payment systems was a joint effort with Biliana Alexandrova-Kabadjova (Mexican Central Bank) and Andreas Krause (Bath University). Research on evolving middlemen strategies was conducted with Tim Gosling, with support from Chris Voudouris at British Telecom. Research on artificial stock markets was conducted with Serafin Martinez-Jaramillo (Mexican Central Bank). Other co-workers and colleagues who have contributed to this line of research include Alma García-Almanza, Sheri Markose, Dietmar Maringer, Simon Lucas, Qingfu Zhang, and Hani Hagras.

References

[1] E. F. Fama. Efficient Capital Markets: A Review of Theory and Empirical Work. *Journal of Finance*, vol. 25, no. 2, pp. 383–417, 1970.

[2] F. Black. Noise. *Journal of Finance*, vol. 41, no. 3, pp. 529–543, 1986.

[3] D. Kahneman, M. W. Riepe. Aspects of Investor Psychology. *Journal of Portfolio Management*, vol. 24, no. 4, pp. 52–65, 1998.

[4] R. Shiller. Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividents? *Americal Economic Review*, vol. 71, no. 3, pp. 421–436, 1981.

[5] A. Shleifer. *Inefficient Markets: An Introduction to Behavioral Finance*, Oxford University Press, UK, 2000.

[6] E. L. Lawler, J. K. Lenstra, A. H. G. Rinnooy Kan, D. B. Shmoys. *The Travelling Salesman Problem: A Guided Tour of Combinatorial Optimization*, John Wiley & Sons, New York, USA, 1985.

[7] S. Lin, B. W. Kernighan. An Effective Heuristic Algorithm for the Traveling-Salesman Problem. *Operations Research*, vol. 21, no. 2, pp. 498–516, 1973.

[8] C. Voudouris, E. P. K. Tsang. Guided Local Search and Its Application to the Traveling Salesman Problem. *European Journal of Operational Research*, vol.113, no. 2, pp. 469–499, 1999.

[9] H. A. Simon. *Models of Man*, Wiley, New York, USA, 1957.

[10] G. Gigerenzer, R. Selten. *Bounded Rationality: The Adaptive Toolbox*, MIT Press, USA, 2001.

[11] A. Rubinstein. *Modeling Bounded Rationality*, MIT Press, USA, 1998.

[12] A. Newell, H. A. Simon. *Human Problem Solving*, Prentice Hall, Englewood Cliffs, NJ, USA, 1972.

[13] S. Russell, P. Norvig. *Artificial Intelligence: A Modern Approach*, Prentice Hall, NJ, USA, 1995.

[14] F. Rossi, P. Van Beek, T. Walsh. *Handbook of Constraint Programming*, Elsevier, 2006.

[15] E. P. K. Tsang. *Foundations of Constraint Satisfaction*, Academic Press, London and San Diego, 1993.

[16] E. K. Burke, G. Kendall. *Search Methodologies: Introductory Tutorials in Optimization and Decision Support Techniques*, Springer, New York, USA, 2005.

[17] H. Hoos, E. P. K. Tsang. Local Search for Constraint Satisfaction. *Handbook of Constraint Programming*, F. Rossi, P. Van Beek, T. Walsh (eds.), Elsevier, pp. 245–277, 2006.

[18] P. Larranaga, J. A. Lozano. *Estimation of Distribution Algorithms: A New Tool for Evolutionary Computation*, Kluwer Academic Publishers, New York, USA, 2001.

[19] D. E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning, Reading*, Addison-Wesley Pubulish Company, MA, USA, 1989.

[20] J. H. Holland. *Adaptation in Natural and Artificial Systems*, University of Michigan Press, Ann Arbor, USA, 1975.

[21] A. Rubinstein. Perfect Equilibrium in a Bargaining Model. *Econometrica*, vol. 50, no. 1, pp. 97–110, 1982.

[22] A. Muthoo. *Bargaining Theory with Applications*, Cambridge University Press, UK, 1999.

[23] N. Jin, E. P. K. Tsang. Co-adaptive Strategies for Sequential Bargaining Problems with Discount Factors and Outside Options. In *Proceedings of Congress on Evolutionary Computation*, Vancouver, BC, Canada, pp. 2149–2156, 2006.

[24] E. P. K. Tsang, N. Jin. An Incentive Method to Handle Constraints in Evolutionary Algorithms with a Case Study. In *Proceedings of European Conference on Genetic Programming, Lecture Notes in Computer Science*, Springer, Budapest, vol. 3905, pp. 133–144, 2006.

[25] T. Gosling, N. Jin, E. P. K. Tsang. Games, Supply Chains and Automatic Strategy Discovery Using Evolutionary Computation. *Handbook of Research on Nature-inspired Computing for Economics and Management*, J. P. Rennard (ed.), Idea Group Publishing, vol 2, Chapter 38, pp. 572–588, 2006.

[26] R. Selten. Boundedly Rational Qualitative Reasoning on Comparative Statics. *Bounded Rationality, the Adaptive Toolbox*, G. Gigerenzer, R. Selten (eds.), MIT Press, USA, pp. 1–8, 2001.

[27] B. A. Kabadjova, A. Krause, E. P. K. Tsang. Competition among Payment Networks Using Generalized Population Based Incremental Learning. In *Proceedings of the 12th International Conference on Computing in Economics and Finance*, Limassol, Cyprus, pp. 22–24, 2006.

[28] B. A. Kabadjova. Artificial Payment Card Market: An Agent Based Approach, Ph. D. dissertation, Centre for Computational Finance and Economic Agents (CCFEA), University of Esesx, UK, 2007.

[29] S. Baluja. Population-based Incremental Learning: A Method for Integrating Genetic Search Based Function Optimisation and Competitive Learning, Technical Report CS-94-163, Carnegie Mellon University, USA, 1994.

[30] M. Kern. Parameter Adaptation in Heuristic Search: A Population-based Approach, Ph. D. dissertation, University of Essex, UK, 2005.

[31] T. Gosling, E. P. K. Tsang. Tackling the Simple Supply Chain Model. In *Proceedings of IEEE Congress on Evolutionary Computation*, Vancouver, Canada, pp. 2179–2186, 2006.

[32] S. Martinez-Jaramillo. Artificial Financial Markets: An Agent Based Approach to Reproduce Stylized Facts and to Study the Red Queen Effect, Ph. D. dissertation, Centre for Computational Finance and Economic Agents (CCFEA), University of Essex, UK, 2007.

[33] S. Martinez-Jaramillo, E. P. K. Tsang. An Heterogeneous, Endogenous and Co-evolutionary GP-based Financial Market. *IEEE Transactions on Evolutionary Computation*, to be published.

Edward P. K. Tsang is a professor in computer science at University of Essex. He is also the deputy director of Centre for Computational Finance and Economic Agents (CCFEA) and a founding member of the Centre for Computational Intelligence at University of Essex.

He has broad interest in business application of artificial intelligence, including, computational finance, computational economics, and constraint satisfaction. Main techniques used included heuristic search, optimization, and evolutionary computation.