Rationality, complexity and self-organisation

Fabio Boschetti

CSIRO, Marine and Atmospheric Research
Fabio.Boschetti@csiro.au
&
School of Earth and Geographical Sciences
at the University of Western Australia

Abstract

I discuss the definition of a rational agent in a set of game theoretical scenarios commonly used to study competition and collaboration in social and economic interactions. In particular I analyse the relation between rationality and the ability of a community of agents to self-organise into viable configurations. I suggest that a useful definition of rationality depends on the specific structure of a problem and consequently a common definition which applies to all scenarios is not available. Unless rationality is defined a priori or obtained by induction via an extensive analysis of a given problem, it seems sensible to accept an adaptive view according to which the concept of rationality is imported into a problem from the experience accumulated in similar settings and modified if evidence requires it.

1 Introduction

The question of what the rational behaviour of an agent should be, given a certain problem, is not simple and has been long debated in several disciplines from philosophy to history, sociology, economics and evolutionary theory. For example, in a fundamental work Max Weber highlights the contextual nature of the problem and proposes the existence of different types of rationality (Kalberg, 1980).

Independently of the problem faced, the connotation of the word ‘rational’ suggests that the process followed to decide the course of action should be well reasoned, logical and explainable; for example, it should be free of contradictions and should satisfy a number of transitivity rules (von Neumann and Morgenste, 1944). Empirical evidence suggests that humans rarely fulfil these requirements: even when facing relatively simple tasks, they often employ heuristics rather than complete logical analysis and such heuristics at times fail to satisfy even basic laws of probability (Tversky and Kahneman, 1974; 1983). Whether these heuristics are simply a manifestation of human limited processing ability (‘bounded rationality’) or are optimal from an efficiency perspective (allowing to rapidly reach an answer which is most of the times correct and thus evolutionary effective) is a topic still unresolved (Simon, 1956; Stanovich, 1999).

The word ‘rational’ has a further connotation in implying the pursuit of a goal. Because humans are social animals such pursuing often happens within a social context which needs accounting for; as a result, the consequences of an individual’s action on his/her neighbours (and thus on society at large) is of crucial importance in the analysis of rationality; how to design rules so that individual actions lead to effective social outcomes has also been widely studied in philosophy since classical times (for a nice and concise review we refer to (Bowles, 2006)).

Today much of the discourse that permeates the analysis of rationality in complex system science (CSS in the following) and organisation theory is inherited from the mathematically elegant results of neoclassical economics, according to which agents acting rationally (that is having full knowledge of a problem and employing that knowledge to take a decisions which are economically optimal for themselves) under certain circumstances can reach a configuration which is also globally optimal. This result holds considerable policy appeal since it promises to align the interest of the individual with that of society in a fully decentralised fashion. Since in common parlance a well performing community is considered to be a well organised one, we can say that under this setting the agents self-organise in an optimal configuration. It is not surprising that this interpretation resonates with the CSS and organisation theory community.

Of course things are never so simple and failures of neoclassical economics theory to deliver its promises to the real world are well known (Barker, 2009; Colander, 2009; Farmer and Foley, 2009; Krugman, 2009). This is not surprising given the previously mentioned bounded rationality and the human propensity for occasionally taking decisions not for rational, but for emotional, ethical or undetermined reasons. While these may appear merely as empirical limitations, other arguments question the validity of the concept of rationality itself. A number of game theoretical examples (prisoners' dilemma and tragedy of the commons, which we describe below), show how easy it is to devise settings in which the alignment between individual and community interests is no longer achievable: in these settings agents acting rationally may not merely miss out on the optimal outcome, they may face doom. The latter are often referred to as the paradoxes of rationality (Campbell and Sowden, 1985). The issue can be made as complex as we wish: when a community comprises several inhomogeneous subgroups, the agents' defecting within their own sub-group may be necessary for global optimal outcome, which is why all free-markets economies devise laws against collusion (Bowles, 2006). As it as well argued in the same work, it is the incentive to defect which allows neoclassical economy theory to provide global optimal outcome: in other words, defecting can lead to doom or to optimal outcomes depending on the context.

So, if the issue is so complex and resistant to definition, why is the concept of rationality still so commonly used within CSS discourse? Why so many 'emergent properties' are still studied in relation to the agents' supposedly rational behaviour? The reason is that most computational work carried out in game theory and network theory, for example, requires a 'default' behaviour against which results are evaluated. Said differently, they need a null-hypothesis against which an experiment is run and such null-hypothesis, under the influence of economic tradition, is often offered by the rational neoclassical economic agent. A simple example clarifies the point: a current research topic of considerable interest explores the apparent evolutionary dilemma of how cooperation can arise from the interaction of selfish actors. However, it is the very assumption that an agent is by default (or rationally) selfish which implies that is cooperation, rather than competition, which needs explaining out.

In this paper I discuss a number of simple game theoretical scenarios which are commonly considered as generalisations of real world problems and employ these scenarios to study the relation between different views of rationality and self-organisation. The purpose is to analyse whether it is useful for CSS scientists to choose the rational agent of neoclassical economics as default behaviour and, if not, how a default rational behaviour could be chosen at all. This is important, because it influences how model results are interpreted and thus popularised outside the modelling community[1].

I suggest that, unless imposed externally by a priori ideology, a unified determination of rationality which applies uniformly to all scenarios is not obvious; rather a meaningful definition of rational behaviour depends on the context and structure of the game: when an

[1] How important such interpretation can potentially be can be judged by the impact neoclassical economics has had in shaping our current world.

agent faces a new problem, determining this rational behaviour will inevitably be part of problem solving itself.

In doing so, I assume an agent to be rational if it has a purpose and takes actions which are consistent with that purpose. Since much CSS research is carried out by computer simulation or mathematical analysis, it is natural to require that a rational behaviour be described in algorithmic terms, i.e., it can be reduced to a formula or a set of instructions[2]. Notice that according to this view random choices can be rational[3]; in other words, agents may decide to let their actions be guided by chance. To summarise, in our discussion a rational behaviour is a set of instructions an agent follows in order to achieve a purpose; the question is what the purpose should be and how the instructions should be given.

2 The basics: the prisoners' dilemma

The prisoners' dilemma (Axelrod, 1984) has been extensively studied in the literature and because of its simplicity it offers a good starting point for our discussion. Among its different possible formulations here I adopt the following. We have two agents; each owns an object that it values at $1; each values the other agent's object at $2; they agree to swap the objects; if they do, each will then own an object which it values at $2 and each will have increased its wealth. However, the transaction is carried out in such a way that each agent may cheat: it may deliver an empty box, not containing the object. If this happens the cheater will obtain the desired object and keep its own for a total value of $3, while the cheated agent will be left with nothing.

The mainstream view of the problem is that it is rational for each agent to cheat. Suppose you have to decide whether or not to place your object in the box just before the transaction; you may think that the other agent is *independently* facing the same decision; if it decides not to include its object, thereby trying to cheat you, you are better off keeping your own object to avoid being cheated; if it decides to include its object, you are still better off keeping your own, thereby cheating, and ending up with both yours and its object. In both cases, independently of what the other agent does, you appear to be better off by not fulfilling your part of the agreement. However, if both agents follow this rational choice, they both fail to carry out the swap and they both miss out on the deal: the end outcome is that they are both left with their original object, which they preferred to swap; rational choices will have led to a sub-optimal outcome.

Some authors disagree with this view and suggest that rationality should lead agents to fulfil the agreement and swap the objects. This objection is based on four arguments. The first argument is that in real world cases agents do not act in a vacuum but are part of a larger society to which moral values apply; breaking such values by cheating may carry an emotional burden or future retribution. I agree with this view but I disregard it in this discussion because it is based on considerations which are external to the game: in the setting I consider there are only two agents carrying out a one-off transaction.

A second argument is based on refusing the paradox: if cheating results in a loss for all agents, there are no grounds to consider it a rational choice.

[2] This needs not be instructions for a computer program; any list of unambiguous instructions which a human agent has to follow can still be considered an algorithm.
[3] So far as we accept an algorithmic version of randomness and we disregard deeper mathematical and information theoretical concerns.

A third argument is based on the symmetric nature of the game (for an in depth analysis of this topic see (Campbell and Sowden, 1985)); there is a perfect symmetry between the two players: the cost, potential rewards, knowledge, and opportunity for cheating are exactly the same for both; given the situation, if a choice is optimal for an agent, it must be optimal for the other agent too. From this perspective there are only two options for optimality: either they both cheat or they both fulfil the agreement; since cheating results in a sub-optimal outcome for both, the rational choice must be for both agents to adhere to the agreement.

The same conclusion can be reached via an alternative avenue which holds a different meaning in terms of rationality. Because of its symmetry, this problem has a unique optimal solution and two equally rational players will necessarily both seek such a single solution; more important, being rational, they both *know* that the other agent will seek such a solution. This sort of rationality, called super-rationality by Hofstadter (Hofstadter, 1985), will entangle the two agents, making them act as one; as a result, the asymmetric option of a single agent cheating will not just be considered sub-optimal, but will in fact not be available at all.

I subscribe to the view that rationality should lead an agent to fulfil its side of the agreement and propose a further, only slightly different interpretation: because of symmetry, each agent should expect that whatever conclusion it has reached, the other agent will likely have done the same. Consequently, hoping to get a benefit from cheating is based on a gamble: that either the other agent has made a mistake or that it has failed to act rationally; gambling on either option is hardly the hallmark of rationality.

In summary, the prisoners' dilemma offers two views of rational behaviour; one is defined a priori: an agent has to act in order to maximise its individual return in most available situations; this leads to a paradox which, for our discussion, implies that the system does not self-organise optimally. A second view of rational behaviour defies the paradox but is defined a posteriori, only after a careful analysis of the structure of the game; this view leads the system to self-organise optimally. In the next sections I consider how this analysis extends to different scenarios.

3 The tragedy of the commons and the minority game

The tragedy of the commons (Hardin, 1968) can be seen as a special case of the prisoners' dilemma (Ostrom, 1990): a number of agents has access to a common but limited resource; each agent has an incentive to exploit as much resource as possible; however, if all agents do so, it is likely that the resource will be overexploited and irreversibly crash. In this case, an agent's constraining its own resource use is analogous to fulfilling an (implicit or explicit) agreement to use the resource sustainably; an agent's attempting to maximise its own resource use is analogous to cheating.

The discussion on rational behaviour in the previous section appears to extend to this problem: following symmetry considerations, rationality should suggest that each agent limits its resource use and aims for responsible and sustainable management. However, a small modification to the problem settings can have a considerable impact on the performance of the agents. Let's suppose that agents want to exploit a resource distributed in space and can access it at different locations. If the resource is abundant, the agents can decide where to access the resource with few consequences. If the resource is scarce, agents risk to overexploit the resource wherever they access it and the tragedy of the commons described above occurs (Boschetti and Brede, 2010). From a game theoretical and CSS perspective, the interesting scenario occurs when the amount of resource available is comparable with the agents' need. In this setting the location at which agents access the resource affects how much resource can be exploited by the community, since optimal exploitation happens only when the community spreads its harvesting effort proportionally to the resource distribution (Boschetti, 2007). Interesting dynamics occur when no external coordination is available: each agent naturally strives to access the resource at the least exploited location, where it can expect to share the limited resource with the least number of competitors.

This problem has been studied extensively under the name of Minority Game or Bar Problem (Arthur, 1994; Challet and Zhang, 1997; Zhang, 1999). At each iteration, each agent chooses where to access the resource and has the same harvest potential and knowledge of the problem: as before, symmetry applies. According to our previous discussion, rational agents should then realise that the best strategy should be the same for everyone: the most symmetric option is for each agent to choose randomly at each iteration, with probability ½, which of the two areas to harvest.

Interestingly, this approach is not optimal either individually or globally (Savit *et al.*, 1999). For a resource amount smaller than the sum of the total agents' harvest potential the best harvest outcome is obtained when agents acts competitively and selfishly by trying to outsmart each other. This is the standard setting of the minority game employed in the literature.

For a resource amount slightly larger than the sum of the total agents' harvest potential the best harvest is provided by the collective intelligence approach (Wolpert *et al.*, 2000). The collective intelligence suggests that individuals should try to maximise the impact their action has on the community behaviour. This is 'measured' as the difference in harvest between what the community gains minus what the community *would* have gained *had* that specific agent *not* participated to the harvest. In other words an action which results in an individual gain but no community gain is penalised. For a resource amount larger than the sum of the total agents' harvest potential (but not so large as to make searching for optimal location unnecessary) the collective intelligence allows the agents to spread effort proportionally to the resource distribution. In this setting, the harvest will be maximised (effectively getting very close to optimality) for each agent and also for the overall community (Boschetti, 2007; Brede and De Vries, 2008).

Interestingly for our discussion, this efficient distribution of effort is obtained within a few iterations, after which agents rarely change the harvesting location: very quickly the symmetry at the individual level (what location each agent accesses) is broken and the symmetry at the community level (even split between locations) is achieved, despite the agents having neither information about the resource distribution nor external coordination.

As a result, an amount of resource roughly equal to the sum of the agents' harvest potential represents a threshold for the community behaviour, which determines whether it is convenient for the agents to choose a collective intelligence approach or a selfish one. In (Boschetti and Brede, 2009) we have proposed a method for dealing with this threshold: agents can adaptively choose whether to employ the collective intelligence or act selfishly by choosing the strategy which provides the largest amount of information about the problem.

4 Constraints on agents' behaviour

Despite their similarity, the tragedy of the commons and the prisoners' dilemma differ in an important feature. In the tragedy of the commons, if every agent cooperates 'too much' no one benefits: if every agent is so concerned about the resource sustainability not to exploit it at all, then no one accesses it and its potential healthy and sustainable productivity is wasted. This possibility is not encountered in the traditional prisoners' dilemma. Of course, not all agents are necessarily so other-regarding; some agents may decide that the best way to preserve their access to the limited resource without restraining their harvest is to deny others such access: some agents may impose constraints on other agents' behaviour by regulating their harvesting rights.

Two natural questions arise: a) how much regulation should be used to protect a resource and b) who should regulate whom. The answer to the first question has been addressed in (Brede and Boschetti, 2009): if too little regulations is used, unlimited exploitations is possible and a likely tragedy of the common occurs; if too much regulations is used, it is likely that each agent is denied exploitation and the resource flourishes without providing harvest to anyone (this outcome is often defined as tragedy of the anticommons in the literature (Heller, 1998; Heller and Eisenberg, 1998; Benkler, 2004); only for intermediate amounts of regulation sustainable exploitation is possible. The answer to the second question is also addressed in (Boschetti and Brede, 2010), the community can self-organise in a viable configuration only provided a sub-set of the agents imposes regulations on others: the subset of agents who impose the regulation will be able to exploit the resource better that the rest; should the rest decide that the unequitable arrangement is not acceptable and to exert their right to regulate, no exploitation would be possible, leading the overall community into the tragedy of the anticommons scenario. Any kind of symmetry is now effectively lost: while in principle any agents could happen to fall in the 'regulator' subset (and thus, if the game was played an infinite number of times each agent would probably obtain an comparable harvest), any actual realisation of the game is limited to a single (or an handful) of runs and lack of symmetry would inevitably arise.

What does this discussion imply in term of rational behaviour? In principle, rationality does not necessarily have to match economic performance, for example issues of resilience or alternative non-economic benefits may be accounted for; however, in the problems I discussed these alternative criteria do not apply, since they lie outside the scope of the problem itself. So, as for the prisoner's dilemma, unless we define it a priori, rationality can be evaluated only on issues of performance and self-organisation. Still, associating rationality with performance and organisation in the minority game and tragedy of the anticommons seems to be somewhat unrealistic, since in order to discriminate what the rational behaviour should be, an agent would need to carry out a number of fairly sophisticated computational steps which is hardly the intuitive understanding of rationality we are accustomed to. Finding a suitable and general determination of rational behaviour in complex problems is thus not straightforward.

5 Discussion

In Figure 1 I summarise the analysis of the scenarios in terms of balance between individual and community performance. The X axis represents the best possible performance of an individual agent and the Y axis the community return averaged over the number of individuals. Since the community return is the average of the individual ones, it is impossible for the community return to be larger than the best possible return an individual can obtain: all scenarios have to lie on the right hand side of the 'average global return= best individual return' line (that is, within in the light gray triangular sub-domain).

The best match between individual and community returns is achieved by the Collective Intelligence (Coin in the figure) in the Minority Game (MG in the figure) when the amount of available resources allows this (Boschetti, 2007). In this scenario individual harvest for both the best and worst performing agents are close to optimal and as a result so is the average community return (Boschetti, 2007). When symmetry is respected in the prisoners' dilemma ('Sym PD' in the Figure) it is impossible for both individual agents to achieve their optimal result (they can't both keep their object and obtain the other agent's one) but they both fulfil their aim by carrying out the swap successfully: both individual and community performances are good. Stepping down the imaginary ladder of optimality we find the selfish behaviour in the Minority Game (‘MG Selfish agent’ in the Figure); this is the best the community can do under severe resource constraints (Brede and De Vries, 2008; Boschetti and Brede, 2010). In this case some agents may fare quite well while others may perform quite badly (Boschetti, 2007) and as a result the average community performance is far from optimal. Similar is the behaviour in the tragedy of the anticommons setting, when some level of resource exploitation is achieved. Next we find the asymmetric case of the prisoners’ dilemma (‘Asym PD’ in the figure) in which one agent cheats and the other is cheated; these are represented by two connected ovals in the figure since the occurrence of one implies the other. Here one agent obtains the maximum possible outcome (both objects) while the other obtains nothing at all; the global performance is worse than had they both acted fairly. Next comes the symmetric prisoners’ dilemma when both agents cheat. Finally, the worst global outcome occurs in the tragedy of the anticommons when no exploitation is achieved.

From the previous discussion it may appear that our analysis of rationality reduces to finding a strategy which provides a maximum return, thereby leading us back to the original definition of economically rational agents we started from. In fact, the previous discussion addressed the question of what implication an individual decision has on the community and thus, in turns, on itself. This feedback between individual and community actions has two crucial implications: it links the dynamics across scales and generates self-referentiality. The first feature is a hallmark of complexity and, in the optimisation parlance, turns into a global problem what an agent may perceive as a local one. The second is the hallmark of many known logical paradoxes and it is not surprising that its presence renders the determination of rational behaviour so challenging.

It is important to remember that the simple games I presented involve dilemmas of an economic nature only. Also, the core of the analysis lies along a line spanning from the individual to the community; in social science parlance, this line connects terms like ‘selfish’ and ‘unethical’ to ‘generous’ and ‘morally conscious’. These terms however are not available to this discussion because the mechanistic nature of the games I analysed lacks social context and thus moral implications. In my opinion this strengthens, rather than weakens the argument: even within the vacuum and simplicity of the analysed scenarios a single criterion to define rationality is not available. Lacking social, moral and emotional implications, the culprits for the problem are easy to pinpoint: self-referentiality and the existence of different levels of analysis. The different levels of analysis are where the problem of rationality itself is framed; self-referentiality is what makes any argument bounce endlessly back and forward between the agents and between the levels themselves. It is easy to imagine that the inclusion of social, moral and emotional implications would complicate the analysis further.

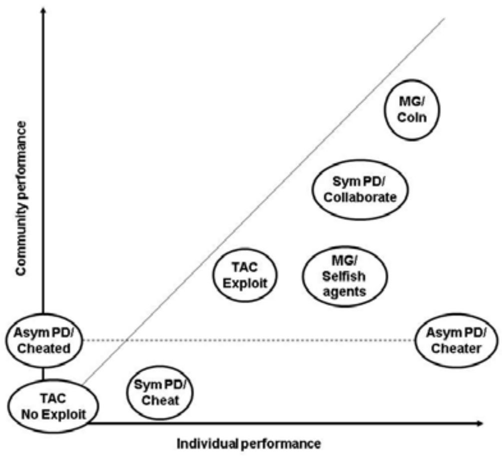

Figure 1. The scenarios discussed are plotted as a function of the individual and community performance they display**.** MG=Minority Game; Sym PD= symmetric prisoners' dilemma; Asym PD= asymmetric prisoners' dilemma; Coin=Collective Intelligence, TAC= Tragedy of the Anticommons.

While it is true that self-referentiality and the existence of different levels of analysis are hallmarks of complexity, it is also true that is it hard to imagine any 'simpler' complex problem, that is any complex problem with fewer components, fewer possible scenarios and greater ease of complete definition. The complexity in these problems is fundamental and 'structural', it does not arise from human cognitive limitation in addressing the size and dimensionality of the problem. As a result, it appears that we are left with three options for defining rational behaviour, each with its benefit and drawbacks:

1) we could rely on an a priori definition, for example based on moral or ideological motives (several authors believe the perfectly rational agent of neoclassical economic theory is the result of one such ideology). The benefit of this choice is its simplicity and its intuitive appeal. The drawbacks are logical: first, we may question how such ideology or moral argument arose in the first place; second, it may lead to the paradoxes of rationality discussed above. In the latter case, a rational behaviour may then be seen as a problem to overcome via analysis and experience, which defeats its original intuitive appeal.
2) We could work purely a posteriori, first by studying the problem empirically, then by trying to appraise its structure by induction and then by devising a workable, useful and effective definition of rational behaviour. The benefit of this approach is that we could design the definition in such a way that it avoids paradoxes and is 'useful' to the agents in terms of performance and organisation; the drawback lies in its obvious complexity.
3) Finally, we could imagine an approach in between the previous two: given a problem, a definition of rationality could be used which has proved to be useful in similar settings, where similarity could be understood, for example, in terms of system structures and possible resulting dynamics (Wolstenholme, 2003); this definition could be adopted until a difficulty is encountered, at which point the definition may be modified accordingly in an adaptive fashion.

6 Conclusions

This analysis cautions against blindly porting values or experiences from one complex problem to another, even those related to something as basic as what it means to be rational. People with extensive real world experience would probably find this unsurprising since many times our judgement is severely tested by complex problems. That no anchor to rationality is available in moving from something as simple as the prisoner's dilemma to the equally simple minority game may be frustrating, may also confirm that in many cases an empirical, rather than a conceptual approach may be needed.

References:

Arthur, W. B. (1994). "Inductive behaviour and bounded rationality," *The American Economic Review*, ISSN 0002-8282, 84: 406-411.

Axelrod, R. (1984). *The Evolution of Cooperation.*, ISBN 0-465-02121-2

Barker, T. (2009). "Towards New Thinking in Economics.pdf," *The Cambridge Trust for New Thinking in Economics*, ISSN.

Benkler, Y. (2004). "Commons-Based Strategies and the Problems of Patents," *Science*, ISSN 0036-8075, 305: 1110-1.

Boschetti, F. (2007). "Improving resource exploitation via collective intelligence by assessing agents' impact on the community outcome," *Ecological Economics*, ISSN 0921-8009, 63(2-3): 553-562.

Boschetti, F. and M. Brede (2009). "An information-based adaptive strategy for resource exploitation in competitive scenarios," *Technological Forecasting and Social Change*, ISSN 0040-1625, 76(4): 525-532.

Boschetti, F. and M. Brede (2010). "Competitive scenarios, community responses and organisational implications," *Emergence: Complexity and Organization*, ISSN 1521-3250, 11(4): 39-47.

Bowles, S. (2006). *Microeconomics: Behavior, Institutions, and Evolution*, ISBN 978-1-4008-2931-6

Brede, M. and F. Boschetti (2009). "Commons and anticommons in a simple renewable resource harvest model," *Ecological Complexity*, ISSN 1476-945X, 6(1): 56-63.

Brede, M. and H. De Vries (2008). "Harvesting Heterogeneous Renewable Resources: Uncoordinated, Selfish, Team-, and Community-Oriented Strategies," *Environmental Modelling & Software*, ISSN 1364-8152, 25(1): 117-128

Campbell, R. and L. Sowden (1985). *Paradoxes of rationality and cooperation : prisoner's dilemma and Newcomb's problem*, ISBN 0774802154

Challet, D. and Y. C. Zhang (1997). "Emergence of cooperation and organization in an evolutionary game," *Physica A*, ISSN 0378-4371, 246(3-4): 407-418.

Colander, D. (2009). "How Economists Got It Wrong: A Nuanced Account," *Middlebury College Working Paper Series* ISSN 0909.

Farmer, J. D. and D. Foley (2009). "The economy needs agent-based modelling," *Nature*, ISSN 0028-0836, 460(7256): 685-686.

Hardin, G. (1968). "The tragedy of the commons," *Science*, ISSN 0036-8075 162: 1243-1248.

Heller, M. A. (1998). "The tragedy of the Anticommons: Property in the Transition from Marx to Markets," *Harvard Law Review*, ISSN 0017-811X, 111: 622-688.

Heller, M. A. and R. A. Eisenberg (1998). "Can Patents Deter Innovation? The Anticommons in Biomedical Reserach," *Science*, ISSN 0036-8075 280: 698-701.

Hofstadter, D. (1985). *Metamagical Themas*, ISBN 0-465-04566-9

Kalberg, S. (1980). "Max Weber's Types of Rationality: Cornerstones for the Analysis of Rationalization Processes in History," *The American Journal of Sociology*, ISSN 0002-9602, 85(5): 1145-1179

Krugman, P. (2009). How Did Economists Get It So Wrong? New York Times. New York.

Ostrom, E. (1990). *Governing the Commons: The Evolution of Institutions for Collective Action* ISBN 0521405998

Savit, R., R. Manuca, et al. (1999). "Adaptive Competition, Market Efficiency, and Phase Transitions," *Physical Review Letters*, ISSN 1550-2376, 82(10): 2203-2206.
Simon, H. A. (1956). "Rational choice and the structure of the environment," *Psychological Review*, ISSN 0033-295X, 63(2): 129-138.
Stanovich, K. (1999). *Who is rational? : studies of individual differences in reasoning* ISBN 0805824723
Tversky, A. and D. Kahneman (1974). "Judgment under uncertainty: Heuristics and biases," *Science*, ISSN 0036-8075, 185: 1124-1131.
Tversky, A. and D. Kahneman (1983). "Extensional versus intuitive reasoning: The conjuctive fallacy in probability judgment," *Psychological Review*, ISSN 0033-295X, 90: 293-315.
von Neumann, J. and O. Morgenste (1944). *Theory of Games and Economic Behavior*, ISBN 0691130612
Wolpert, D., K. Wheeler, et al. (2000). "Intelligence for control of distributed dynamical systems," *Europhysics Letters*, ISSN 1286-4854, 49(6).
Wolstenholme, E. F. (2003). "Towards the definition and use of a core set of archetypal structures in system dynamics," *System Dynamics Review*, ISSN 1099-1727, 19(1): 7-26.
Zhang, Y. C. (1999). "Modeling market mechanism with evolutionary games," *Europhysics News*, ISSN 1432-1092, 29: 51-53.