Simulating Information Sharing in Crisis Response Coalitions as a Minority Game

Ariën J. van der Wal
Netherlands Defence Academy
aj.vd.wal@nlda.nl

Tim Grant
Retired But Active Researchers
tim.grant.iscram@gmail.com

ABSTRACT

A major crisis or disaster attracts a response from multiple organizations. These organizations need to work together as a coalition. To do so effectively, they must share information. Differences in organizational culture give information a scarcity value, leading to the emergence of information markets. As the crisis progresses, organizations learn to work together, building up trust and lowering the "price" for information shared. Organizations that have worked together in previous crises have a price advantage. In short, the information sharing changes dynamically. Evolutionary game theory – and the minority game in particular - has been recently applied to markets in which physical goods are exchanged. This paper presents the first results from simulating an information sharing market in crisis response coalitions as a minority game.

Keywords

Information market, evolutionary game, minority game, trust.

INTRODUCTION

A major crisis or disaster event, such as a famine, an earthquake, flooding, or a train crash, inevitably attracts tens or hundreds of responding organizations. These organizations may range from UN agencies and government departments, through national and international non-governmental organizations (NGOs), to commercial suppliers, the media, and individuals. Since most organizations want to retain sovereignty over their operations, they organize themselves as a coalition. Each coalition partner must find out what is happening, negotiate its place in the coalition, synchronize and deconflict its activities with other partners, evaluate its achievements, and eventually withdraw from the coalition. To do so they must share information.

The environment that crisis response and management organizations must contend with is highly dynamic, and can be regarded as multiple mini-crises that emerge over time (Hiltz, Van de Walle & Turoff, 2010). At first, nobody has the "big picture". The physical, governmental, and telecommunications infrastructure may have failed. Responding organizations may have to reach the crisis location before they can begin their work. In the first few days or weeks, the organizations involved must discover who their fellow-partners are, learn what they know and what they can do, and build mutual trust. This must be done "on-the-job". Organizations that have previously worked together, either in exercises or during previous crises, have a lesser hurdle to overcome.

Our research focuses on information sharing in crisis response coalitions. We define information sharing as the process of making information available to other individuals, teams, or organizations in the coalition. Case studies (Grant & Van den Heuvel, 2010) show that information sharing in coalitions has the characteristics of an information market. An information market is a mechanism for allocating and distributing information resources (Stewart, 1996). Organizations barter for information, trade it for information of greater value, or use it as an instrument of power (McGee & Prusak, 1993), indicating that information is a scarce resource. In a crisis, information sharing must constantly adapt as the crisis evolves.

Our ultimate aim is to develop a tool that will enable crisis response organizations to increase their effectiveness by speeding up the on-the-job learning process they must undergo immediately after a crisis or disaster strikes. Our research uses evolutionary games (Weibull, 1995), and the minority game (MG) in particular, to simulate information sharing between crisis response and management organizations. In this paper we present the simulation results from information sharing in a coalition consisting of 100 organizations that have worked together in previous crises and 900 organizations that are responding to a crisis for the first time using a MG.

RELEVANT THEORY

Coalition Information Sharing as a Market

As for markets in physical goods, an information market has buyers and sellers (Stewart, 1996). A buyer is an organization that is seeking information, and sellers are organizations that have information available for distribution. The market provides the mechanisms for buyers to discover sellers (and vice versa), to match the respective needs of buyers and sellers, and to make the exchange possible. A given organization may at times be a buyer or a seller, and must be both simultaneously when information is traded for other information. The only difference is the nature of the goods that are traded. Unlike physical goods, information can be readily copied, it can be in several places at once, selling it does not diminish the supply, buyers generally purchase a piece of information only once, and, once sold, information cannot be recalled (McGee & Prusak, 1993).

In an earlier paper, we constructed a taxonomy of possible mechanisms for information markets (Grant & van der Wal, 2012). In terms of this taxonomy, an information market in a crisis response coalition is a digital, multi-sector market with a closed boundary. Where there is a quid pro quo for a piece of information, then it is a natural market, with the quid pro quo being determined either by supply versus demand means or by barter. If the quid pro quo is deferred, then it is a gift market. It is unlikely that information sharing in crisis response and management is planned or monitored (either centrally or decentrally). However, there are organizations whose raison d'être is to facilitate information sharing by providing a networking infrastructure (e.g. the UN Office for the Coordination of Humanitarian Affairs).

Evolutionary Games

Game theory was conceived by John von Neumann to determine mathematically the optimal strategies in a competition between adversaries. Classical game theory requires that the participants make rational choices, i.e. based on what is best for each individual participant (rather than on the coalition as a whole). The thrust of research into game theory has been to study equilibria, of which the Nash equilibrium is the most well-known.

Evolutionary game theory relaxes the requirement for rationality and tests alternative strategies for their ability to survive (Weibull, 1995). This involves a learning process, not at the level of individual participants but at the level of the population from which the individuals are drawn. Typically, learning is modelled as a Darwinian process, involving selection, mutation and replication. Evolutionary games that have been studied in detail include hawk-dove, war of attrition, stag hunt, producer-scrounger, the tragedy of the commons, and prisoner's dilemma. The research thrust has been to find evolutionary stable strategies (ESSs), such as the assessor strategy in the hawk-dove game and the tit-for-tat algorithm in the prisoner's dilemma.

The Minority Game

In the 1990s, physicists began to be attracted to problems in economics and finance, bringing in mathematical techniques from research areas such as thermodynamics and statistical mechanics. In particular, they were interested in non-equilibrium processes (Anderson, Arrow & Pines, 1988). Challet, Zhang and their co-workers introduced evolutionary game theory to modelling economic markets in the late 1990s (Challet & Zhang, 1997; 1998) (Zhang, 1998). Their key ingredients for modelling markets were: 1. A large number of independent agents participate in a market, 2. Each agent has a set of alternative strategies for making decisions, 3. The aggregate activity of the agents results in a market outcome, which is known to all, 4. Agents use the history of outcomes to make their decisions, and 5. Agents learn from their own experience, and have no theory about the market.

Challet and Zhang (1997) drew their inspiration from Arthur's (1994) El Farol Bar problem. This problem is now generically known as the minority game (MG), i.e. one in which, at each time step, participants choose to be on side A or on side B. The winners are those who are in the minority. There are many variants on this game. In the simplest version (Zhang, 1998), the winners collect a point. Each participant has a finite number of strategies, and participants make their decisions based on their memory of the history of outcomes. Their memory may be limited in capacity. The history only records which side wins, without revealing the numbers of winners. Hence, the time series can be represented as a binary sequence.

Applying the MG to Information Markets

We have not found a reference in the literature to a prior application of the MG to information markets or information sharing. Despite this, it would seem that the MG is as applicable to information markets as it is to physical goods markets. The differences between information and physical goods – information can be readily copied, in several places at once, selling does not diminish supply, information is purchased once, and sold information cannot be recalled (McGee & Prusak, 1993) – do not invalidate the MG. The winning side represents the set of participants that share a piece of information. This does not need to be a singleton set or to exclude the original seller. The MG does not represent either the purchase or recall process.

By comparison with traditional economic models of markets, the big advantage of using the MG is that it can represent dynamic behaviour. This makes it suited to simulating the highly dynamic environment in crises. By contrast, traditional economic models are aimed at finding an equilibrium, and research into evolutionary game theory seeks ESSs. It is questionable whether equilibria or stable strategies exist in crisis response and management. The MG is also suited to the coalition organisational form. Since the participants in a coalition wish to retain a degree of sovereignty, there is no central controlling organization. The MG does not require arbitration or coordination of the market, as in planned economies. The participants are intrinsically selfish and do not need to be rational decision-makers.

Compared to more elaborate models based on intelligent agent technology, the MG is much simpler to model, faster to run, and better suited to modelling markets in which thousands or even millions of agents take part. Each participant in an MG is equivalent to an agent with a reasoning capability restricted to the rule level (Rasmussen, 1983). Moreover, the pre- and post-conditions in each rule are binary. Adaptation in the MG is equivalent to reinforcement learning. However, the simplicity of the MG is also its source of limitations. If more elaborate modelling of each participant's decision-making process, of the communication protocol or negotiation process between participants, or of knowledge-level processes such as planning, sense-making or concept learning is needed, then intelligent agent technology would be unavoidable.

SIMULATION OF INFORMATION SHARING

Implementation

We assume that the number of agents is of the order of 1000. A memory of M bits is assigned to each player.

The game is started by randomly drawing S=5 strategies for each player from the pool of 2^{2^M} strategies. All the S strategies of a player can collect credits if they would win or not given the previous M bits of the minority decision time series and the actual outcome. The real payoff of an agent is for the actual chosen strategy if it predicts the outcome correctly (i.e. wins the next play). Players adapt to the market by using the highest ranking strategy. The performance of a player (its "success rate") is calculated as the credits accumulated over time divided by the time period.

Over time, information exchange as modelled by the MG improves the effectiveness of information markets, due to a number of effects. First the participants can adjust ("evolve to better strategies") their set of strategies. Even though the basic agent attitude is selfish, the memory size of each agent has a positive effect on the exchange of information as demonstrated by the decreasing standard deviation of the size of the (winning) minority group. This means that despite a selfish attitude of the individual agent, they cooperate by studying the market performance. Our preliminary results indicate that successful information markets indeed occur in these circumstances. First of all we have verified that the simulation correctly reproduced the results known from literature for the binary MG in case of a large number of agents (N =1001) and a relatively small number of strategies per agent. For N=1001 agents the attendance A varies between 0 and N=1001, implying that group G_1 was the winning group if $A_1 \leq 500$. In the case $501 \leq A_1 \leq 1001$ group G_1 loses. The attendance number of the other group, A_0, is of course the complement $A_0 = N - A_1$, since $G_0 \cup G_1$ represent all agents. As expected, the average of A_1 converges to N/2 and since there is no preference for either group G_0 or G_1, the distribution of A is symmetrical around N/2. From these results it is obvious that the asymptotic standard deviation σ for larger memory size M becomes smaller, indicative of less waste of resources (i.e. a better motivated choice per agent). It should be noted that although the memory of each agent is only M bits, i.e. only based upon the outcome of the M previous binary outcomes (win (0) or lose (1)), the performance of each individual agent initially improves by the use of a scoring system: At each time step each agent evaluates the virtual score for each strategy in its bag of strategies for the total time period that the game is played. This means that apart from the short-time memory of size M each agent uses a same time horizon that increases as time evolves. We view this as a long-time memory. We have calculated the standard deviation in the attendance number A_1, averaged over the last 800 epochs, resulting from a MG with a homogeneous population of N = 1001 agents. From these results it is clear that a larger memory size M improves the performance of an agent, giving him a lead over agents that are more poorly equipped. Another way to express this is by evaluating the average success rate of an agent, $\overline{s}$, defined by The total score of all agents divided by the number of agents N and the number of epochs. If the decision process of the agents were purely random, (corresponding with the case M = 1), one would expect $\overline{s}$ to asymptotically converge to 0.5.

Coalition-specific modelling results

We now make the link between the general MG where agents ('players') make a common decision (the minority outcome) and the cooperation between coalition partners. The previous definitions of the parameters describing the outcome of the game, viz. the attendance number A_i, success rate s and memory M and the time horizon (the number of the past simulation epochs) can in market terms be translated into: The outcome of the binary game is 0 or 1, meaning that a information transfer between the two parties is successful (1) or not (0). Note that in a further study we can refine the model to include the *direction* of the information transfer, by defining a ternary game with outcomes -1 , 0 , and 1. In the present model the standard deviation of the attendance A is a measure for the "correctness" of the price of the information in the market, indicating that if the volatility is high there is not much agreement on the "right" price, whereas if the volatility is low, parties are nearer to a correct price. The interesting feature in the model is that we do *not* define a price to model the market, but only look at the volatility (standard deviation) of the price. We define the characteristic features of our model as follows:

- The average success rate of an agent in the MG can be defined as a fraction: The information transfer rate (number of information transfers per epoch) is defined as the ratio of successful information exchanges and the total number of attempted exchanges.
- The memory M_i is identified as the trust of a coalition partner i in another partner: It is defined as the trust of agents in group i in agents of another coalition partner $j \neq i$ based on recent information transactions.
- The time horizon T is defined as the time interval (number of epochs) during which trust between agents of different coalition partners has been developed.

It is to be expected that the more trust exists between coalition partners, the easier it is to establish a successful exchange of information between them. Furthermore, we expect that if the level of trust is highly asymmetrical (one coalition partner has much more trust than the other) the transfer of information will be hindered. In order to analyze the effect of trust on the success of a subpopulation of agents within a coalition, we made a MG simulation of a mixed population of N = 1001 agents, consisting of N_1= 101 agents with a variable memory size ("trust") M_1 and a majority of N_2 = 900 agents with a fixed trust M_2 = 2, reflecting less experience in the information exchange process in coalitions. In Fig. 1 the success rates of the two groups of agents are plotted as a function of the trust level M_1. As expected the success rates for M_1 = 2 are degenerate, since in that case the group of 1001 agents is homogeneous. Interestingly the success rate s_1 increases rapidly with M_1 and for $M_1 > 3$ even surpasses 0.5, which is the expectation value if the decision to exchange information would have been made purely on a statistical basis. So it appears that even a little more experience in working together in a coalition (i.e. trust) brings a big advantage compared to the other agents. This is an interesting observation that warrants further research.

CONCLUSIONS AND FURTHER WORK

This paper presents the first results from simulating information sharing in crisis response coalitions as a Minority Game. The MG is a particular form of evolutionary game that has been successfully applied to simulating the dynamic behaviour of physical-goods markets. Here it is applied to information markets. Since crises are highly dynamic situations, traditional market models based on equilibria or evolutionary stable strategies are inappropriate. The MG is essentially an agent-based simulation with a simple, rule--based model of the agents' decision-making process and a simple model of their interactions. In this paper we illustrate the use of the MG by investigating information sharing in a coalition of 1001 organizations, of which 10% have experience from previous crises (such as UN agencies, the ICRC, US AID, and military forces) and 90% are involved in a crisis for the first time (such as local NGOs and host-nation government departments). The latter organizations must learn "on-the-job". We expected that the more experienced organizations would

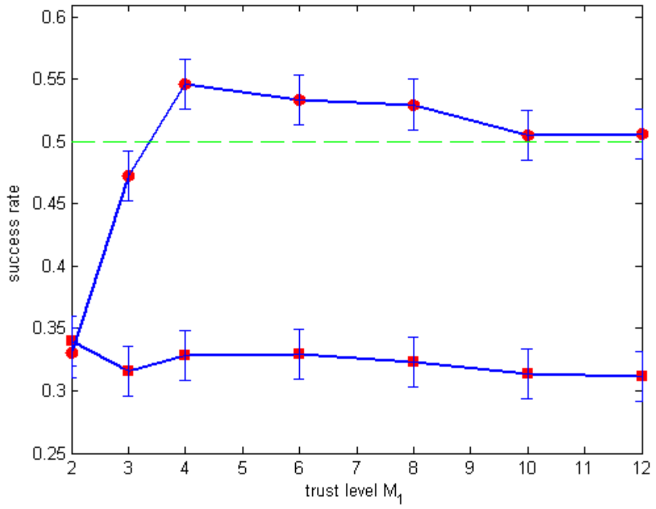

Figure 1 The success rates of population 1 (N_1 = 101, dots) and population 2 (N_2 = 900, M_2 = 2, squares) as a function of the memory size of population 1. It is clear from these graphs that even a modest increase in memory size (> 3) is sufficient for a relatively small group of agents to outsmart a majority of agents with smaller memory (M_2 = 2). For $M_1 \geq$ 4 the rate of successful information exchanges rises above the random maximum value of 0.5 (dashed line).

preferentially share information with their less-experienced partners. However, we found that more-experienced organizations preferred to share information with other more-experienced organizations, and less-experienced organizations with other less-experienced organizations. This finding needs to be verified against coalition behaviour in real crises. We have just touched on a very small part of a wide range of possible future research into simulating information sharing within crisis response coalitions. Moreover, there are several areas where related empirical research is needed. For example, little is known about how crisis managers value information and whether this is time- and/or situation-dependent. While the MG approach is superior to traditional market models, it has also limitations , such as binary payoff of the agents, their finite memory and time horizon and the absence of negotiation between individual agents.

REFERENCES

1. Anderson, P.W., Arrow, K. & Pines, D. (eds). (1988) The Economy as an Evolving Complex System. Addison-Wesley Co., Redwood City.
2. Arthur, W.B. (1994) Inductive Reasoning and Bounded Rationality. *Am. Econ. Assoc .Papers and Proc*. 84, 406-411.
3. Challet, D. & Zhang, Y.-C. (1997) Emergence of Cooperation and Organization in an Evolutionary Game. *Physica A*, 246, 407.
4. Challet, D. & Zhang, Y.-C. (1998) On the Minority Game: Analytical and Numerical studies. *Physica A*, 256, 514.
5. Grant, T.J. & van den Heuvel, G.G.A. (2010) Modelling the Information Sharing Process in Military Coalitions: A work in progress. *Proceedings of the 7th International ISCRAM Conference*, Seattle, WA.
6. Grant, T.J. & Van der Wal, A.J. (2012) A Taxonomy of Market Mechanisms for Information Sharing in Crisis Response Coalitions. *Proceedings of the 9th International ISCRAM Conference*, Vancouver, Canada.
7. Hiltz, S.R., Van de Walle, B. & Turoff, M. (2010) The Domain of Emergency Management Information. In Van de Walle, B., Turoff, M. & Hiltz, S.R. (eds). Information Systems for Emergency Management. M.E.Sharpe, Armonk, NY.
8. McGee, J., & Prusak, L. (1993) *Managing Information Strategically*. John Wiley & Sons, New York.
9. Rasmussen, J. (1983) Skills, Rules, and Knowledge: Signals, signs and symbols, and other distinctions in human performance modelling. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-13, 257-266.
10. Stewart, T.A. (1996) *Intellectual Capital - The New Wealth of Organizations*. McGraw-Hill.
11. Weibull, J.W. (1995) Evolutionary Game Theory, MIT Press, Cambridge, MA.
12. Zhang, Y.-C. (1998) Modeling Market Mechanism with Evolutionary Games. Europhysics News, 29, 51.