Terrorist networks and the lethality of attacks: an illustrative agent based model on evolutionary principles

Paul Ormerod

Abstract

A data base developed from the Memorial Institute for the Prevention of Terrorism's (MIPT) Terrorism Knowledge Base for the years 1998–2005 was provided to participants in the workshop. The distribution of fatalities in terrorist attacks is, like many outcomes of human social and economic processes, heavily right-skewed. We propose an agent based model to analyse this, and to enable generalisations to be made from the historical data set. The model is inspired by modelling developments in cultural evolutionary theory. We argue that a more appropriate 'null' model of behaviour in the social sciences is on based upon the principle of copying, rather than the economic assumption of rationality in the standard social science model.

Keywords: Agent based model, Terrorist fatalities, Cultural evolution, Mutations

Introduction

Asal and Rethemeyer [1] report an econometric analysis of a data base developed from the Memorial Institute for the Prevention of Terrorism's (MIPT) Terrorism Knowledge Base. The dataset is complete for the years 1998–2005. Of the 395 clearly identified terrorist organizations operating throughout the world over this period, only 68 killed 10 or more people during that period. Indeed, only 28 killed more than 100 people. The econometric study analyses the factors which can account for this dramatic difference in organizational lethality.

They conclude that "(1) large organizations, (2) organizations that address supernatural audiences through religious ideologies, (3) organizations with religious-ethnonationalist ideologies— ideologies that define an other and play to the supernatural, (4) organizations that build and maintain extensive alliance connections with peers, and (5) organizations that maintain control over territory are the primary actors in this story. Though much of the organizational and social movements literature suggest that new organizations are less effective and able, our data was unable to find evidence that newness matters. Some widely held theories about the correlates of lethality—including the belief

that state sponsorship and "homebase" regime-type would affect organizational lethality—could not be substantiated with our data. In fact, there is equivocal evidence that state sponsorship tend to restrain killing by client organizations. Size coupled with religious and ethnonationalist ideology generates the capability needed to pursue deadly ends".

This paper considers a potential generalisation of the econometric approach using the methodology of agent based modelling (ABM). Section 2 briefly considers methodological aspects of the issue, and section 3 discusses some principles of agent behaviour. Section 4 sets out the model and section 5 discusses some illustrative results.

Some methodological reflections

The Asal and Rethemeyer paper is detailed and thorough, and takes proper account of the heavily right-skewed nature of the dependent variable, the number of people killed in each incident. The dependent variable ranges from 0 to 3505 with a median of 0, a mean of 31.36 and a standard deviation of 202.04. Of the 395 organizations for which there are data, 240 of those organizations perpetrated one or more incidents that resulted in no fatalities.

However, econometric analysis of data, no matter how sophisticated, essentially involves fitting a plane through

Correspondence: pormerod@volterra.co.uk
Volterra Partners LLP, London, UK

the n dimensions of the explanatory data. Even when the regression technique is based upon the principle of maximum likelihood, the result can still be given this geometric interpretation. Helbing [2] offers a more detailed description of the potential restrictions of this approach, but one of the essential problems can be summarised as follows.

A restriction on the ability to generalise from econometric results is that a small number of data points may exercise a strong influence on the fit of the n-dimensional plane. Consider, for example, the standard linear regression model $y = X\beta + \varepsilon$, where the vector of estimated parameters is $b = (X^T X)^{-1} X^T y$ and the fitted values are $y^* = Xb = X(X^T X)^{-1} X^T y$.

The hat matrix, H, maps the vector of fitted values to the vector of observed values, and describes the influence each observed value has on each fitted value [3], where $H = X(X^T X)^{-1} X^T$. It is the orthogonal projection onto the column space of the matrix of explanatory factors, X.

This matrix can be used to identify observations which have a large influence on the results of a regression. If such observations exist, ideally we would like to have more data from this part of the observation space. So if we have a small number of observations in the tail of the distributions of both an explanatory and the dependent variable and these are correlated, such observations will inevitably exercise a strong influence on the results.

Perhaps not surprisingly, therefore, large terrorist organisations, and particularly those which maintain and build extensive alliances, are identified in econometric analysis as being two of the key factors which are linked to effective attacks leading to high levels of fatality.

However, more generally, we are interested not so much in projecting the single path of history which we actually observe into the future, but in considering the evolutionary potential of the groups. The policy message from the econometrics is to focus on large, well connected groups in the prevention of attacks. But, for example, what is the potential for small, less connected groups to acquire the ability to develop the capacity to carry out highly effective attacks?

Agent based modelling is a way of trying to examine the evolutionary potential of any given system. Helbing and Barietti [4], for example, state that 'computer simulation can be seen as experimental technique for hypothesis testing and scenario analysis'. They go on to argue that 'agent-based simulations are suited for detailed hypothesis-testing, i.e. for the study of the consequences of ex-ante hypotheses regarding the interactions of agents. One could say that they can serve as a sort of magnifying glass or telescope ("socioscope"), which may be used to understand our reality better... by modeling the relationships on the level of individuals in a rule-based way, agent-based simulations allow one to produce characteristic features of the system as emergent phenomena without having to make a priori assumptions regarding the aggregate ("macroscopic") system properties'.

Approaches to agent behaviour

The standard socio-economic science model, SSSM [5], postulates a high level of cognitive ability on the part of agents. Agents are assumed to be able to gather all relevant information, and to process it in such a way as to arrive at an optimal decision, given their (fixed) tastes and preferences. Even with the relaxation of the assumption of complete information [6], agents are still presumed to have formidable cognitive powers. This, the core model of economics, has a certain amount of explanatory power, though as [5] goes on to state; 'Within economics there is essentially only one model to be adapted to every application: optimization subject to constraints due to resource limitations, institutional rules and/or the behavior of others, as in Cournot-Nash equilibria. The economic literature is not the best place to find new inspiration beyond these traditional technical methods of modeling'.

Perhaps the most important challenge to this approach comes when decisions do not depend not on omniscient cost-benefit analysis of isolated agents with fixed tastes and preferences, but when the decision of any given agent depends in part directly on what other actors are doing. In such situations, which are probably the norm rather than the exception in social settings, not only do choices involve many options for which costs and benefits would be impossible to calculate (e.g., what friends to keep, what job to pursue, what game to play, etc.), but the preferences of agents themselves evolve over time in the light of what others do.

Complex choices can be fundamentally different from simple two-choice scenarios, such that the problem becomes very difficult to predict, as has been demonstrated in ecological [7] and human settings [8]. Such scenarios are where so-called zero-intelligence models [9] do better at understanding emergent patterns in collective behaviour. However, despite their empirical success e.g. [10-12], they have met with resistance amongst social scientists.

The zero intelligence model is based upon the particle model of physics. Indeed, one of the fastest-rising keywords in the physics literature is 'social'. Literally thousands of papers in physics are devoted to modeling social systems, and indeed regular sections of leading journals such as *Physica A* and *Physical Review E* are devoted to this topic. The analogy between people and particles has been so consistent that a recent popular review was appropriately titled *The Social Atom* [11].

This approach has provided significant insights into modeling collective interactions in social systems, from Internet communities to pedestrian and vehicular traffic, economic markets and even prehistoric human migrations e.g. [13-15].

The idea that there are serious limits to human cognitive powers in complex systems is one which has strong empirical support. Kahneman, for example [16] argues that 'humans reason poorly and act intuitively'. Decisions may often be taken in circumstances in which the assumption that individuals have no knowledge of the situation is a better approximation to reality than is the assumption that they possess complete information and have the capacity to process this information to make optimal decisions.

An illustration of the limits to human awareness and social calculation is the well known Prisoner's Dilemma game, invented by Drescher and Flood in 1950. The optimal strategy or "Nash equilibrium" for the one period game was discovered very quickly. However, as documented in detail in [17]. Flood recruited distinguished RAND analysts John Williams and Armen Alchian, a mathematician and economist respectively, to play 100 repetitions of the game. The Nash equilibrium strategy ought to have been played by completely rational individuals 100 times. It might, of course, have taken a few plays for these high-powered academics to learn the strategy. But Alchian chose co-operation rather than the Nash strategy of defection 68 times, and Williams no fewer than 78 times. Their recorded comments are fascinating in themselves. Williams, the mathematician, began by expecting both players to co-operate, whereas Alchian the economist expected defection, but as the game progressed, co-operation became the dominant choice of both players.

Even now, after almost 60 years of analysis and literally thousands of scientific papers on the subject, when sufficient uncertainty is introduced into the game, the optimal strategy remains unknown. Certainly, some strategies do better than most in many circumstances, but no one has yet discovered the optimal strategy even for a game that is as simple to describe as the Prisoner's Dilemma.

Assumptions of optimality and rationality can be certainly be useful when payoffs are predictable from one event to the next – hunting and gathering in a consistent environment, for example, or even modern situations where the complexity of choices is low [18]. But more generally, the zero intelligence model may be more useful as the 'null'.

However, as is argued in [19], human agents are fundamentally different from particles in physics in that, however imperfectly, they can act with purpose and intent. So the basic null model of zero intelligence needs to be modified to incorporate some aspect of this functionality [20].

Simon in his seminal paper on behavioural economics [21] argued that the fundamental issue which all sentient beings have to take into account when taking decisions is to reduce the massive dimensionality of the choice set which they face: 'Broadly stated, the task is to replace the global rationality of economic man with a kind of rational behavior that is compatible with the access to information, and the computational capacities that are actually possessed by organisms, including man, in the kinds of environments in which such organisms exist'.

An important way of coping with an evolving, complex environment is to follow a strategy of copying, or social learning as it is described in cultural evolution. In social and economic systems, decision makers often pay attention to each other either because they have limited information about the problem itself or limited ability to process even the information that is available [22].

A striking example of the power of simple copying strategies in an evolving environment is provided by the computer tournament described by [23]). As the abstract states: 'Social learning (learning through observation or interaction with other individuals) is widespread in nature and is central to the remarkable success of humanity, yet it remains unclear why copying is profitable and how to copy most effectively. To address these questions, we organized a computer tournament in which entrants submitted strategies specifying how to use social learning and its asocial alternative (for example, trial-and-error learning) to acquire adaptive behavior in a complex environment. Most current theory predicts the emergence of mixed strategies that rely on some combination of the two types of learning. In the tournament, however, strategies that relied heavily on social learning were found to be remarkably successful, even when asocial information was no more costly than social information. Social learning proved advantageous because individuals frequently demonstrated the highest-payoff behavior in their repertoire, inadvertently filtering information for copiers. The winning strategy relied nearly exclusively on social learning and weighted information according to the time since acquisition'.

In other words, in a complex environment in which the pay-off to various strategies was, by design, constantly evolving, simple copying proved a very effective strategy.

Copying, of course, is the essence of the principle of preferential attachment, initially formulated in a general way by Simon [24] and rediscovered by Barabasi and Albert [25]. In its more recent incarnation, the model has been hugely influential because it resulted in a power law (or at least long-tailed) degree distribution (connections per network node) – a kind of distribution so intriguing to many that the editor of *Wired* magazine wrote an entire book,

ten years later, about its significance to modern online economies [26].

However, a key drawback of the approach is that it is ultimately static. The rankings in the distribution, in other words, gradually become fixed, and attempts to modify the basic model are rather artificial [27]. Turnover in rankings is not just a feature of modern markets of popular culture, but operates on the time scale required for cities to evolve [28].

Preferential attachment is a special case of a model developed in cultural evolutionary theory e. g, [29,30], in which it remains the basic principle of decision making, but an agent is also able (with a small probability) to make an innovative choice. The model is developed from the concept of genetic evolution, which is based on the principles of copying and mutation (innovation).

In its most recent and most general formulation [31] the model of social learning in cultural evolution is, with just two parameters, capable of replicating *any* right skewed distribution and of accounting for the turnover which is observed in relative standings in any evolutionary environment. One parameter describes the relative frequencies with which preferential attachment (copying) and innovation are used to make decisions, and the other describes the span of historical time used to observe the decisions which other agents have made. Clearly, this latter is different for a firm choosing its location to a teenager choosing which video to download from YouTube.

The model

In order to try to generalise from the data set used by Asal and Rethemeyer, we retain copying as one of the basic principles used by terrorist organisations. They can observe and copy, exactly as in the evolutionary learning computer tournament, tactics and strategies used by other organisations. In a situation in which the environment in which they operate evolves, the pay off (i.e. number of fatalities inflicted) may of course differ from previously enacted versions of the same tactic. But copying is one of the building blocks of the model.

We augment the model with further behavioural principles which seem appropriate in this context. By its very nature, the terrorist world is clandestine, and organisations will differ in their propensity to share information about tactical capabilities with other such organisations. In addition, organisations will differ in their willingness or ability to absorb and execute a tactic they have not previously used, even if it is explained to them by another organisation.

We consider these factors in the context of the ability of the organisations to acquire the capacity to carry out an innovation, or tactical mutation, developed by a terrorist group. An overview of the model is shown in

Figure 1, which is similar to the model developed to account for the diffusion of technological innovations across companies in four industries in the Greater Manchester region of the UK [32]: The model takes an initial mutation to be exogenous and it is taken up by one agent/organisation at the outset. The characteristics of the agents are governed by their willingness to innovate, their desire to keep innovation to themselves, and their willingness to communicate with others. The innovating agent will be connected to other agents via the network structure, and at the next step of the model the innovation will be passed on according to the extent agents discover the mutation and their own willingness to take it up. At further steps of the model further agents may be able to discover and take up the mutation, until eventually no further take up occurs.

We define two different methods by which tactical mutations may be passed on via the network linkages. The first is a direct relationship between two partners, while the second is a group relationship.

First, an organisation with an innovation will provide it to another only if its level of secrecy, or the propensity of a group to try to retain the benefits of its innovations, is less than the absorptive capacity, or the degree to
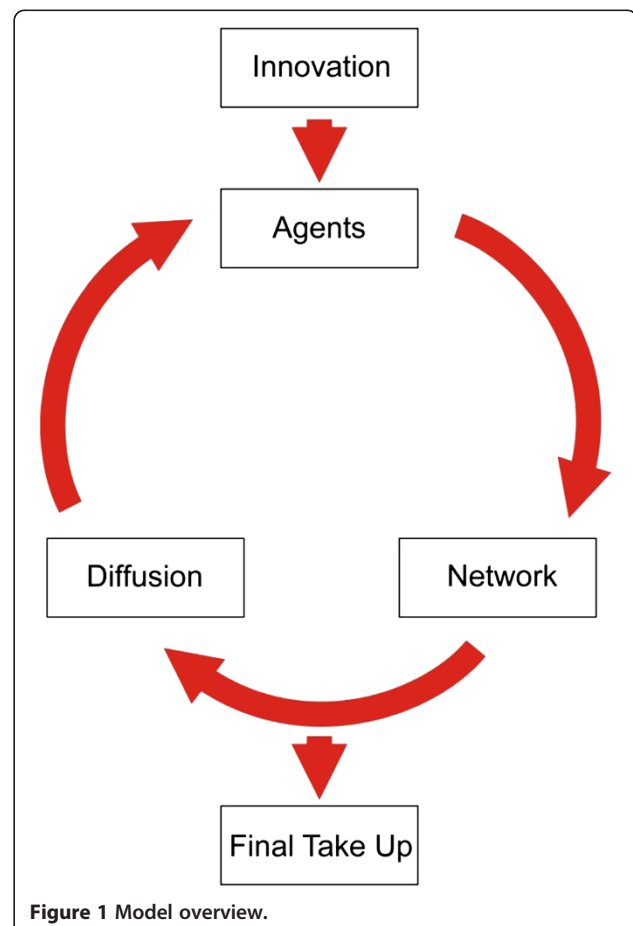

Figure 1 Model overview.

which a terrorist actively engages in activities which enable it to identify and adopt new innovations, of the organisations it is linked with. This method of adopting an mutation represents a mutual relationship or exchange between terrorist organisations and implies a degree of trust or collaboration.

The second method for spreading a mutation is based on the principle of copying. Here if a group looks at the spectrum of organisations to which it is linked and finds that the proportion that have adopted an innovation is higher than their own personal threshold, they will mimic their behaviour and adopt the mutation. In some circumstances this threshold may be very high and only when all or nearly all of the other groups an organisation has relationships with have taken up an innovation will they be persuaded to do the same. For other organisations relatively few outfits may have to have the same innovation before they adopt it. This mechanism represents a copying behaviour. This may occur even when a terrorist group may not fully understand the reasons and benefits of a mutation but relies on observing that others have adopted it. This behaviour is more likely to be a response to competitor behaviour.

Figure 2 plots the distribution of the numbers of connections between organisations in the data set.

In each separate solution of the model, a network is generated which has the same distribution as that of the data. More precisely, the null hypothesis that the model network has the same degree distribution as the data is not rejected on a Kolmogorov-Smirnov test at a p-value <0.01, far below the conventional level of statistical significance. In other words, we can regard the model networks as being identical in degree distribution to the actual data.

In this illustrative application of the model:

- Each agent is allocated at random a willingness to seek and adopt an innovation/mutation, α_i, drawn from a uniform on [0,1].

Figure 2 Degree distribution of connections between terrorist organisations.
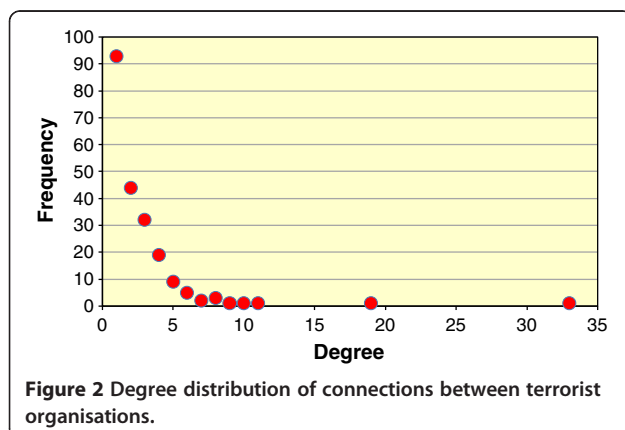
- Each agent is allocated at random a willingness to share an innovation, σ_j, drawn from a uniform on [0,1].
- An agent adopts the innovation of another if $\alpha_i > \sigma_j$.
- Each agent has a threshold for imitating the innovations of its neighbours, τ_i drawn from uniform on [0.5,1].

This latter rule is the principle of binary choice (adopt or not adopt) with externalities [21,33].

The model proceeds in a series of steps, in each of which the following procedure is operated. First, all agents not in possession of the mutation who are not connected to any other organisation which does have the innovation are identified. These agents play no further part in this particular step of the model. Of course, in subsequent steps, agents to which they are connected may by then have acquired the mutation. So whilst they may not acquire it in this period, they are not precluded from so doing in future steps within the same solution of the model.

If the absorptive capacity, α_i of the agent without the innovation is greater than the willingness to share (or secrecy), σ_j, of the agent with the innovation, the agent is assumed to adopt the innovation. If not, the copying rule is then invoked, and the agent adopts it if its threshold is lower than the proportion of agents to which it is connected which has the innovation.

Any particular solution of the model ends when no agent adopts the mutation in a given step of the model. We measure and record the proportion of agents which have adopted the innovation.

This process is the basic building block of the illustrative approach. We can readily imagine, however, that certain capabilities are inherently harder to acquire than others. Shooting a civilian (e.g. a member of the IRA shooting a Protestant at random in Northern Ireland) requires much lower levels of skill and expertise than a co-ordinated bombing. We illustrate the effects of introducing this into the model in the following heuristic way.

We populate the model with 1000 agents and solve it 1000 times. At the end of each solution, the organisations which have acquired the innovation are deemed capable of carrying out a relatively low level attack, which we describe as Level 1 capability. The model is reinitialised, a new agent is selected at random to acquire an innovation, and the process is repeated. At the end of this, more agents will have acquired Level capability. And those which acquired it in the first solution and in addition acquire it in this new one are deemed to have acquired Level 2 capability. In other words, they have acquired sufficient technical skill to mount more serious attacks, with presumed higher levels of fatalities.
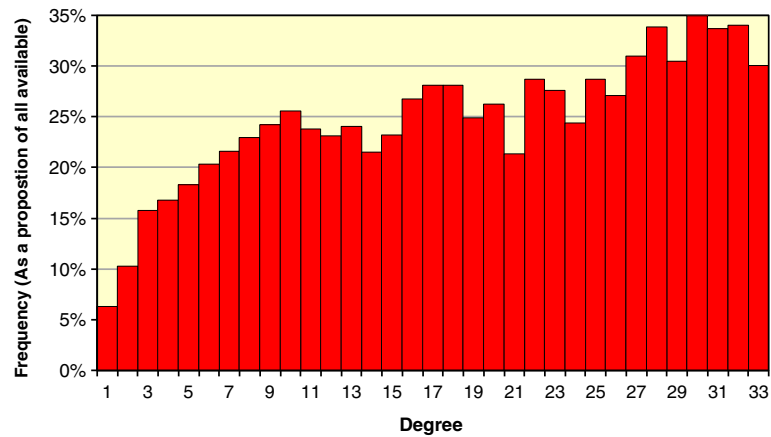
Figure 3 Proportion of terrorist organisations for each value of degree of connections which acquire the capability to carry out Level 1 attacks, average across 1,000 solutions.

Finally, we repeat the process again, and those agents which acquire the innovation on all three occasions are deemed capable of Level 3 attacks.

Discussion and results

Some properties of the model are illustrated in the three graphs below. These are the results averaged across 1,000 separate solutions of the model, and show the total percentage of organisations with a given degree (number) of links which acquire the capabilities to acquire the ability to carry out Levels 1, 2 and 3-type attacks.

First, Figure 3 shows the results for Level 1 capabilities.

Quite rapidly, the connection between the acquisition of this level of capability and the degree of an organisation's links with other terrorist groups falls away. Only a very small percentage of groups with a very low number of connections are able to acquire even this level of

capability. But the proportion of organisations which acquire it and which have, say, just 6 links is not very different from the proportion of those with 26 links which acquire. Only for those with a very high number of connections does the proportion rise.

Figures 4 and 5 show the results for the proportions which acquire Levels 2 and 3 capabilities.

The precise topology of the network and the parameters allocated to each agent differ across each of the 1,000 solutions of the model, but the same network and parameter values are retained within each solution for each of the steps of acquisition capability. So, for example, the structure may be such that an agent is connected to another which is very likely to acquire the mutation, and the parameters are such (in particular the absorption and secrecy parameters) that it, too, is therefore likely to acquire it. So the proportion of agents of any given degree acquiring Level 3 capability is not
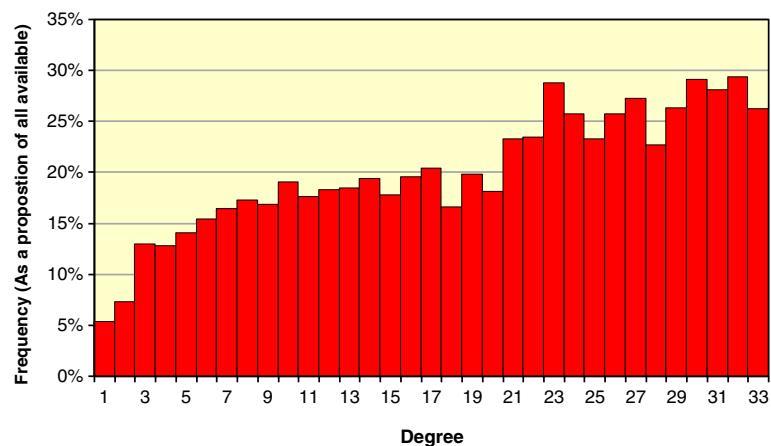

Figure 4 Proportion of terrorist organisations for each value of degree of connections which acquire the capability to carry out Level 2 attacks, average across 1,000 solutions.
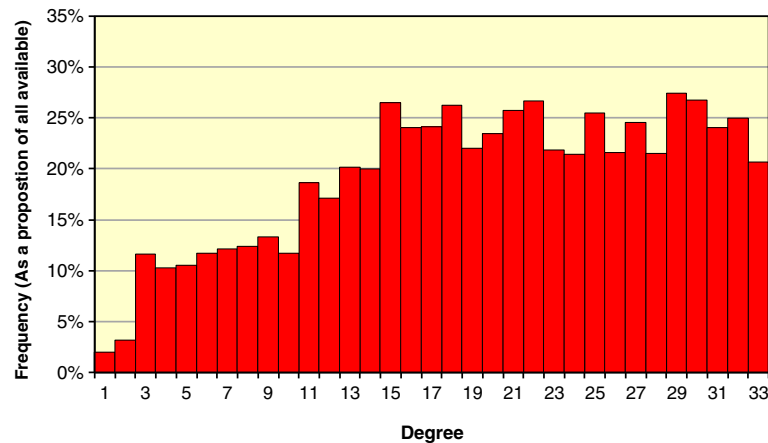
Figure 5 Proportion of terrorist organisations for each value of degree of connections which acquire the capability to carry out Level 3 attacks, average across 1,000 solutions.

simply (as an approximation) one third the value of those acquiring Level 1, but is much higher.

These results are illustrative of a principle of how to try to generalize from the data set. Indeed, we see that the proportion of organizations with only weak connections to others which acquire level 3 capabilities is, apart from those with virtually none, non-trivial. So although in the one actual history we are able to observe, organizations with low levels of connections tended not to carry out attacks involving high levels of fatalities, this does not mean that in future they will be unable to acquire such characteristics.

We are not claiming that this is a definite model with which to inform future policy, though it does offer some potential guidelines. Rather, it illustrates how an agent based model constructed on evolutionary principles in a complex environment can be used to extract more information from a data set than is possible using conventional analytical methods such as econometrics. One potential extensions is to endogenise the network, such that agents have incentives to develop new links, but also introducing extinctions of agents, which may be enhanced (made easier for the authorities) as the number of connections of an organization increase ([34] offers a general model of extinctions in an evolving network). Another is to endogenise the parameters, so that agent behavior becomes reinforced both by reference to their own previous decisions and by reference to the properties of their neighbours.

Competing interests

The author declares that he has no competing interests.

References
1. V Asal, KR Rethemeyer, The nature of the beast: organizational structure and the lethality of terrorist attacks. J. Polit. **70**, 437–449 (2008)
2. D Helbing, *Pluralistic Modeling of Complex Systems* (2010). arXiv:1007.2818v1. http://arxiv.org/, Cornell University Library
3. DC Hoaglin, RE Welsch, The hat matrix in regression and ANOVA. Am. Stat. **32**, 17–22 (1978)
4. D Helbing, S Balietti, *Agent Based Modeling,'* FuturICT meeting, Zurich, June 2011 (2011). http://dl.dropbox.com/u/6002187/16June_AgentBasedModelling.pdf
5. V Smith, Constructivist and ecological rationality in economics. Am. Econ. Rev. **93**, 465–508 (2003)
6. GA Akerlof, The market for lemons: quality uncertainty and the market mechanism. Q. J. Econ. **84**, 488–500 (1970)
7. BA Melbourne, A Hastings, Highly variable spread rates in replicated biological invasions: fundamental limits to predictability. Science **325**, 1536–1539 (2009)
8. MJ Salganik, PS Dodds, DJ Watts, Experimental study of inequality and unpredictability in an artificial cultural market. Science **311**, 854–856 (2006)
9. JD Farmer, P Patelli, I Zovko, The predictive power of zero intelligence in financial markets. Proc. Natl. Acad. Sci. **102**, 2254–2259 (2005)
10. P Ball, *Critical Mass: How One Thing Leads to Another* (Heinemann, London, 2004)
11. M Buchanan, *The Social Atom* (Bloomsbury, London, 2007)
12. MJ Newman, A-L Barabási, DJ Watts, *The Structure and Dynamics of Networks* (Princeton University Press, Princeton, 2006)
13. GJ Ackland, M Signitzer, K Stratford, MH Cohen, Cultural hitchhiking on the wave of advance of beneficial technologies. Proc. Natl. Acad. Sci. **104**, 8714–8719 (2007)
14. I Farkas, D Helbing, T Vicsek, Mexican waves in an excitable medium. Nature **419**, 131–132 (2002)
15. X Gabaix, P Gopikrishnan, V Plerou, HE Stanley, Institutional investors and stock market volatility. Q. J. Econ. **121**, 461–504 (2006)
16. D Kahneman, Maps of bounded rationality: psychology for behavioral economics. Am. Econ. Rev. **93**, 1449–1475 (2003)
17. P Mirowski, *Machine Dreams: Economics Becomes a Cyborg Science* (CUP, Cambridge UK, 2002)
18. B Winterhalder, EA Smith, Analyzing adaptive trategies: human behavioral ecology at twenty-five. Evol. Anthropology **9**, 51–72 (2000)
19. RA Bentley, P Ormerod, Agents, Intelligence and Social Atoms, in *Integrating Science and the Humanities*, ed. by M. Collard, E. Slingerland (Oxford University Press, Oxford, 2011)
20. P Ormerod, M Trabbati, K Glass, R Colbaugh, Explaining Social and Economic Phenomena by Models with Low or Zero Cognition Agents, in

Complexity Hints for Economic Policy: New Economic Windows, Part IV, ed. by (Springer, Milan, 2007), pp. 201–210. doi:10.1007/978-88-470-0534-1_10

21. HA Simon, A behavioral model of rational choice. Q J Econ **69**, 99–118 (1955)

22. TC Schelling, Hockey helmets, concealed weapons, and daylight saving: a study of binary choices with externalities. J Confl. Resolut. **17**, 381–428 (1973)

23. L Rendell, R Boyd, D Cownden, M Enquist, K Eriksson, MW Feldman, L Fogarty, S Ghirlanda, T Lillicrap, KN Laland, Why copy others? Insights from the social learning strategies tournament. Science **328**, 208–213 (2010)

24. HA Simon, On a class of skew distribution functions. Biometrika **42**, 425–440 (1955)

25. R Albert, A-L Barabási, Statistical mechanics of complex networks. Rev. Mod. Phys. **74**, 47–97 (2002)

26. C Anderson, *The Long Tail: Why the Future of Business Is Selling Less of More* (Hyperion, New York, 2006)

27. SN Dorogovtsev, JFF Mendes, Evolution of networks with ageing of sites. Physical Review E. **62**, 1842 (2000)

28. M Batty, Rank Clocks. Nature **444**, 592–596 (2006)

29. SJ Shennan, JR Wilkinson, Ceramic style change and neutral evolution: a case study from Neolithic Europe. Am. Antiq. **66**, 577–594 (2001)

30. MW Hahn, RA Bentley, Drift as a mechanism for cultural change: an example from baby names. Proc. R. Soc. B **270**, S1–S4 (2003)

31. RA Bentley, P Ormerod, M Batty, Evolving social influence in large populations. Behav. Ecol. Sociobiol. **65**, 537–546 (2011)

32. P Ormerod, B Rosewell, G Wiltshire, Network Models of Innovation Processes and the Policy Implications, in *Handbook on the Economic Complexity of Technological Change*, ed. by C. Antonelli (Edward Elgar, Cheltenham UK, 2010)

33. DJ Watts, A simple model of global cascades on random networks. Proc. Natl. Acad. **99**, 5766–5771 (2002)

34. P Ormerod, R Colbaugh, Cascades of failure and extinction in evolving social networks. J. Artif. Soc. and Soc. Simul. **9**, 4 (2006)