

Agents, intelligence, and social atoms

R. Alexander Bentley and Paul Ormerod

'Rational models are psychologically unrealistic..... the central characteristic of agents is not that they reason poorly, but that they often act intuitively. And the behavior of these agents is not guided by what they are able to compute, but by what they happen to see at a given moment'

-- Daniel Kahneman, Nobel economics lecture, 2003

1. Introduction

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* (Buchanan 2007).

In the last 15 years, the science of interacting particles 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., Ackland et al. 2007; Barabási 2005; Buchanan 2007; Farkas et al. 2002; Gabaix et al. 2006; Helbing et al. 2000; Newman et al. 2006).

There is of course a long history of work in the social science which effectively applies agent-based modelling and which precedes the modern, computer-aided ‘complexity science’. A good example is Ijiri and Simon (1964). Even earlier, social scientists were discussing the implications of emergent – which Durkheim (1912) called “effervescent” – phenomena resulting from human social interactions. Despite this, awareness of recent work in ‘social physics’ is still low amongst social scientists themselves. There is a growing list of exceptions to this, such as an increase in the number of complexity science centres, some enthusiastic reviews (Hanson 2004; Valverde 2004; Borgotti et al. 2009), journals such as *Quality and Quantity* and *Journal of Artificial Societies and Social Simulation*, and a few outstanding collaborations (Knappet et al. 2008)

Two major insights have arisen from this work, both from classic studies and the more recent, computer-aided revival. The first is the recognition that human societies are generally open systems, under constant flux, and not the kind of closed system that leads to equilibrium. The second is that many of the emergent, often complex, patterns in society need not require complex behavior on the part of individuals.

As a corollary of the second point, we suggest that the most appropriate “null model” of

individual behavior in larger societies is in fact what we term the “zero-intelligence” model (e.g. Farmer et.al. 2005) with its roots in statistical physics, rather than that of the behavioral model from economics of fully rational agents. In other words, decisions are generally 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.

Particles, of course, cannot act with purpose and intent, cannot learn, and hence cannot adapt their actions on the basis of the outcomes of previous actions by themselves or by other particles. Social physics is therefore quite radical – it posits that we might usefully see how far we can get by modelling human agents *as if* they behaved with ‘zero intelligence’). It is rather literally a ‘null model’ – assuming as little as possible, in order to identify the most general characteristics of collective human behavior.

Economics has of course made some attempt to move in this direction, following the development of the concept of bounded rationality (Akerlof 1970). Individuals may have incomplete information, and the amount of information may vary amongst individuals. So agents are taking decisions with only a portion of the information assumed to be available to them under the assumption of full rationality. However, agents under bounded rationality are still assumed to make optimal decisions subject to the amount of information which they have. So in general, this concept is closer to that of full rationality than it is to that of zero intelligence

We expand on these points in the following way. First, we make some remarks on the concept of non-equilibrium. We then set out a rationale underlying the idea that in general agents behave as if they had zero, or close to zero, intelligence

We go on in the bulk of the chapter to discuss the various ways in which the zero-intelligence null model may need to be modified in a social science context, and set out ways in which this has been done, with particular reference to anthropology.

The modifications may appear extensive at first sight, but in any given context only one or two will usually be required in order to develop a satisfactory model rather than the full set. Importantly, these modifications still leave the behavior of agents closer to that of zero intelligence than that of rationality. Finally, we draw some brief conclusions.

2. Non-equilibrium

In anthropology, assumptions of equilibrium, optimization and rational cost-benefit analysis by individuals is prominent in optimal foraging theory and human behavioral ecology (Winterhalder and Smith 2000), Such models often assume that human decisions have an

optimal value, in the sense of a fitness-enhancing payoff. To fit various case studies, these payoffs can be adjusted to account for behaviors that are highly culturally-dependent or individually-variable, as within small-scale societies living in similar environments (e.g., Cronk and Gerkey 2007).

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 (Winterhalder and Smith 2000).

Many of these models use differential equations to model “typical” behavior, and often use the assumption of equilibrium to solve them (e.g. Gintis 2007). Equilibrium models often assume that something, such as energy or wealth, is conserved in exchange processes. This is often not the case with human society and economics, which are characterized by open, non-equilibrium systems (Farmer and Geanakoplos 2009).

In modern capitalist economies, for example, “income is most definitely not conserved.” (Gallegati et al. 2006). For example, the level of income per head in the United States in 2009, even allowing for inflation, is over six times higher than it was in 1909. Wealth and income both follow highly skewed, non-equilibrium distributions (Pareto 1907), in a vast range of economies. Even in small pastoralist societies, where livestock are the principal form of wealth and transactions (Salzman 1999), wealth is not in equilibrium, because it flows “uphill” rather than “downhill”—that is, the rich-get-richer through competition (Salzman 1999; Hayden 2001). Even a setback, like major drought, can make the rich richer, as wealthy pastoralists absorb the livestock of smaller holders who are driven out of the pastoral economy (Fratkin 1989:46; Hayden 2001). As a result, the distributions of wealth in pastoralist societies follow right-skewed distributions, some of which are close to Pareto form (Bentley 2003).

Another aspect of non-equilibrium systems is renewal, that is, a continual flux of entities. In these open systems, new agents are continually added through birth, inauguration, immigration, while others are eliminated through death, modification, and extinction. These agents may be on any human scale, from ideas in the mind to individuals to groups, corporations and so on.

3. The zero intelligence agent

In addition to its assumption of equilibrium, the standard economic approach ascribes considerable intelligence to agents in the decision making process, both in terms of the

information they gather and the rules they use to process it. Even in models in which agents have incomplete information (e.g. Akerlof 1970), the decision making rule is still based upon the principle of maximizing, i.e. taking the optimal decision on the basis of the information available.

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 unpredictable, as has been demonstrated in ecological (Melbourne and Hastings 2009) and human settings (Salganik et al. 2009). Such scenarios are where zero-intelligence models do better at understanding emergent patterns in collective behavior. However, despite their empirical success (e.g., Ball 2004; Buchanan 2007; Newman et al. 2006: 415-551), they have met with resistance amongst social scientists.

This lack of interest is understandable, given the apparent incongruity between the anthropological focus on cultural, thinking people—in all of their complexity and variety—versus the simple particles of physics. Humans report acting with purpose and intent, and we do not assume that they literally act with zero intelligence, whether as individuals or as part of collective decision-making groups (e.g. kinship groups, economic firms, governments). However, the dimension of the problem which often faces human decision makers, particularly in the modern world, is so large that it is *as if* agents were operating with approximately zero intelligence. In addition, the environment in which agents operate changes frequently, so that it is extremely hard or even impossible for agents to learn systematically and evolve rules of behavior which are in some sense “optimal”.

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 by Mirowski (2002), 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.

The game of chess offers a further illustration of limits to the ability of agents to process information in an optimal way. Chess has a relatively small number of unequivocal rules and these rules do not change over any relevant time scale. The agents, the players in this case, have full information about the rules. And the position in the game is completely transparent at any point in time, so that a player knows for certain the moves that his or her opponent has already made. In addition, the player knows all the legitimate moves that his/her opponent could make when it is his/her turn to play. So each player has a large amount of precise information about the game.

Yet in most situations in a chess game, there is no known best move. Chess grandmasters such as Gary Kasparov describe how they use their skill and experience to make what they consider to be a reasonable move instead of engaging in a futile search for the best possible move. Computers can now beat the best human chess players, but the fraction of situations in which they can find the optimal move is tiny, and such optimal solutions are believed even by the very strongest players to be completely beyond the capacity of humans to solve. The information is readily available, but the constraint to finding the best moves is the ability to process it.

The practical limitations on agents' cognitive abilities were in fact discovered almost at the outset of the development of game theory. In addition to the experiment with RAND scientists, Flood carried out a range of practical tests, and indeed shortly afterwards abandoned game theory altogether because of the disjuncture between how people "ought" to behave and how they actually do behave (Mirowski 2002).

For example, Flood offered RAND secretaries a choice. One of them was given the option of either receiving a fixed sum of money (\$10, say), or receiving a total of \$15 provided that agreement could be reached with another secretary as to how this money was to be divided between them. One Nash solution is that the two split the marginal difference. In other words, they divide the extra \$5 between them so that they get \$12.50 and \$2.50 respectively. In practice most secretaries appealed not to the new idea of the Nash equilibrium, but to the concept of fairness: they divided the total amount exactly equally, \$7.50 each. The

development of experimental and behavioral economics in the past 20 years or so effectively confirms the findings of Flood's initial work (e.g. Smith 2003).

While some concept of fairness may be as old as humanity itself (Ofek 2001; Hardy 2009) or even earlier (de Waal 2009), norms of fairness vary among cultures. When the Ultimatum Game (where one player chooses how much of the money pot to share with a second player) has been played in cross-cultural settings – ranging from the Hadza hunter-gatherers of Tanzania to the Machiguenga horticulturalists of the Amazon, to the Lamalera whale hunters of Indonesia – the amount offered varies substantially from culture to culture (Henrich et al. 2005). Societies that survive by more cooperative endeavours, such as whale hunting, generally exhibited more generosity in economic games (> 50%) than those based on individual activities (< 50%), such as gathering (Henrich et al. 2005).

This further confirms that human decisions are social, with very little “intelligence” in the sense of optimal cost-benefit calculation. There may be a cultural, local logic to these decision tendencies, but even if they were optimally tailored to each particular environment, we see the cost-benefit ‘calculator’ in this case as evolutionary selection over many human generations (O'Brien 1996), rather than on the part of the individual actors (Boone and Smith 1998).

4. Adding intelligence (a little goes a long way)

Consider the extinction of firms in modern economies, which is currently being made vivid in the economic crisis of 2008/09. Two key stylised facts have been established about the extinction patterns of firms. First, the probability of extinction is highest at the start of the firm's existence, but soon becomes more or less invariant to the age of the firm. A second finding—a much more recent one of which most economists are unaware—is that the relationship between the size and frequency of firm extinctions is closely approximated by a power law (Cook and Ormerod, 2003 and Di Guilmi et.al. 2004). This statistically self-similar relationship is very similar to that which exists in the fossil record for the extinction of biological species (Solé et al. 1997; Newman 1996; Newman and Eble 1999).

Economic theory has a great deal to say about many aspects of firm behavior, but offers relatively little on the deaths of firms. Building on previously successful zero-intelligence models (e.g., Amaral et al. 1998; Stanley et al. 1996), Ormerod and Rosewell (2003) developed an evolutionary, agent based model of firm evolution and extinction, which, following other such evolutionary models (e.g., Sneppen et al. 1995), yielded properties that conform closely to the aforementioned stylised facts.

The base version of the model was a “zero intelligence” version, in that each model firm was unable to acquire knowledge about either the true impact of other firms’ strategies on its own fitness, or the true impact of changes to its own strategies on its fitness. Specifically, N firms were connected to each other by a matrix of uniformly distributed interconnections $J_{ij} \in [-1, 1]$. The initial connections J_{ij} were drawn at random, and in each step each agent i had one of its J_{ji} updated, i.e. assigned with a new value chosen at random in the interval $[-1, 1]$. The fitness $F_i(t)$ of each agent was calculated, where $F_i(t) = \sum_{j=1}^N J_{ji}(t)$. If $F_i(t) < 0$ then the agent was deemed extinct. An extinction event of size M is defined as a period in which M agents become extinct.

Despite the fact that agents are acting in a purely random way, the model was nevertheless capable of generating the two stylised facts on firm extinctions: high extinction probability initially and overall power law distribution of extinction event sizes.

The effect of providing even a small amount of intelligence to the agents was dramatic. Starting with the ‘zero-intelligence’ base, Ormerod and Rosewell then assigned the model firms different amounts of knowledge about the effects of strategy. Even a small amount of knowledge produced a sharp increase in the mean firm age at extinction, which was a departure from the real-world data. As both the amount of knowledge available and the number of knowledgeable firms were increased, the extinction patterns of firms departed dramatically from the real-world evidence. Strikingly, as the behavior of firms in the model approached the level of rationality of economic theory, the extinction patterns became completely different from the empirical evidence. Needless to say, this implies that the standard assumption of economic theory that firms act in a rational way – whether full or bounded rationality is wholly incompatible with the empirical evidence on patterns of firm extinction

5. Adding copying

Ormerod and Rosewell’s (2003) study indicates that (a) economic agents have very limited capacities to acquire knowledge about the true impact of their strategies; and (b) if those capabilities were enhanced among firms generally, the effect would be significant even for small increases in knowledge. This suggests that one of the key parameters of a complex interactive system is the degree to which agents are acting upon objective information versus simply responding to neighbouring agents or even just copying them.

This question has wide relevance, well beyond the case study of economic firms (e.g. Surowiecki 2004). It is well known, in situations from responding to an emergency to preparing for climate change, that when conditions change rapidly, people do not readily “re-

calculate” an optimal response, but rather they often persist in traditional behavior, doing as others do and/or as they are socially expected to do. In fact, in situations where information is incomplete and/or time to decide is short, copying the behavior of others can be seen as cost-beneficial decision in itself:

Once some cultural transmission capacities exist, natural selection favors improved learning efficiencies, such as abilities to identify and preferentially copy models who are likely to possess better-than-average information... Copiers thus evolve to provide all sorts of benefits (i.e., “deference”) to targeted models in order to induce preferred models to grant greater access and cooperation. (Henrich and Gil-White 2001: 167)

As Henrich and Gil-White (2001: note 2) point out, these benefits are self-serving, in that “as a species' reliance on true imitation increases, the relative benefit from improved copying abilities increases. But, if a species only rarely uses its imitative abilities ... the benefits of improved imitation may not exceed the costs.”

Assessing independence of behavior from aggregate statistics is a challenge. The paradox is that that the behavior of the majority could arise because that behavior is optimal, and every individual has made the same decision independently, or it could be a rather arbitrary behavior that has spread through the population through copying.

Perhaps the best way to approach this is to assume a continuum: in some cases people make decisions fairly independently, based on the inherent qualities of the choice, whereas in others they are prone to be highly socially influenced. In fact, this distinction can become the subject of empirical testing, rather than an assumption we must put *a priori* into a model.

The following classic behavioral spread model (cf. Rogers 1964; Henrich 2001) models the probability of each individual adopting a certain new idea/behavior, at time t as:

$$p(t) = (\mu + qF(t))(1-F(t)), \quad (1)$$

where the parameters q and μ represent the degree of imitation and individual selection, respectively. Hence by setting the parameters q and μ , this model can represent two evolutionary modes of behavior spread: pure social influence versus purposeful selection.

These different models predict identifiably different patterns in the adoption and subsequent abandonment of ideas or products within a population. The pure social influence (high q , low μ) produces a symmetric pattern through time: that which is adopted rapidly in the population

does not accumulate, and consequently declines rapidly from peak frequency, whereas that which is slower to move to its peak accumulates more, and declines more slowly (Henrich 2001; Berger and Le Mens 2009; Bentley and Ormerod 2009). Deliberate behavioral selection (high μ , low q), in contrast, produces an asymmetric pattern, whereby the decline from the peak is considerably slower than is the rise to the peak, and vice versa (Bentley and Ormerod 2009).

A modern example is searches on Google of the phrase ‘swine flu’ in late April/early May 2009 (Bentley and Ormerod 2009). We can ask the question, was the interest in ‘swine flu’ generated from independent selection – perhaps people genuinely individually scared by the media hype – or was it purely the socially-mediated spread of an idea, no different than any other flash news item?

Model (1) above can be used to characterize internet searches for the phrase ‘swine flu’ in a wide range of countries across the world. We find that the phenomenon was somewhere between these extremes, with selection $\mu = 7.4\%$ and imitation $q = 46\%$. This might seem strongly in favour of social imitation, but the selection parameter was actually much higher for ‘swine flu’ searches than for more solidly socially-mediated spread, as with baby names (Hahn and Bentley 2003), where $\mu = 0.06\%$ might be typical. The selection parameter μ is powerful and just a small increase can lead to a large boost in the rate of popularity increase. Hence, overall, we found that genuine concern (selection) for ‘swine flu’ was significant in the rise of the mass interest around the world, as opposed to simply the social spread of a panic.

Optimality models would struggle to account for the average degree of social influence in the collective behavior. Of course, our ‘swine’ flu study, for example, merely characterizes the parameter values to fit a model to a population, but not the variability in individual for being interested in ‘swine flu’ among the millions of Googlers. Optimality models are even more insufficient to explain such variability of human behavior (Nettle 2009), and yet the heterogeneity of human behavior is so basic that it may actually have evolved to enable groups to function through specialization and exchange (Ofek 2001; Hardy 2009).

One of the simplest ways to study the effect of this heterogeneity is to start with a zero-intelligence model and then apply simple *thresholds* to individual behaviors. This can add a rich dimension of realism and complexity to a model. For example, an agent might copy other agents only if the fraction of those agents with a different opinion exceeds a certain discrete threshold.

Watts (2002) did this when incorporating the classic psychological results of Solomon Asch (1955) into a model of information cascades across a network of interconnected agents. The

agents were varied according to the fraction of neighbours needed to exhibit a behavior before the agent itself would adopt it. Watts found a rich variety of cascade behavior, based not only on the mean threshold and network structure, but also on the heterogeneity of thresholds among the agents. Ormerod and Colbaugh (2006) extended the model into a dynamic context, in which the topology connecting the agents is not fixed, but evolves over time as agents seek alliances with the aim of increasing their overall fitness.

In general, social science requires some addition to the assumption of purely random interactions between agents (the particles of the social physics models). Our brains are wired to be social (Dunbar and Shultz 2007; Hrdy 2009), which clearly biases the interactions (Henrich and Gil-White 2001). The specific nature of the addition will vary according to the particular circumstances, and we suggest a classification of such potential additions.

6. Networks: structured copying

Many studies of non-equilibrium social phenomena focus on the origin of long-tail (power-law) distributions (Newman 2006) as emergent properties. One famous hypothesis for their origin has been the “scale-free network” (Albert and Barabási 2002; Barabási and Bonabeau 2003), which is essentially a zero-intelligence model. In what have become seminal studies of the network of Internet web pages, the number of links to each site was found to be power-law distributed (Huberman et al. 1998; Albert and Barabási 2002). The title of Buchanan’s (2007) book does not really capture the explosion of *network* approaches where the “social atoms” are essentially linked to each other and to human collective behavior (e.g. Buchanan 2007; Albert and Barabási 2002; Hufnagel et al. 2006; Newman et al. 2006; Solé et al. 2005).

Despite grandiose claims of re-inventing social science (Barabási 2005), these models as developed in physics often fail to capture essential elements of human behavior. Almost all models of scale-free network growth are ultimately grounded upon the rule of “preferential attachment” (Albert and Barabási 2002), in that new connections continually made within the network are preferentially attached to already well-connected agents.

The Barabási-Albert (1999) model of preferential attachment was 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 (Anderson 2006). A variety of preferential attachment models ensued, most in some way modelling the number of connections an agent adds in time $t + 1$ as proportional to the connections it already had at time t . With continual growth in the number of agents, this leads to a power-law degree distribution (Adamic and Huberman 2000), as many similar models of cumulative advantage have demonstrated (e.g., Ijiri and Simon 1964; Price 1965;

West and Deerting 1995; Newman et al. 2006: 335-348). Basically, the larger the variance in the growth rate, the more that variance contributes multiplicatively to inequality.

At first glance, preferential attachment may seem to have numerous parallels with anthropology, for example the classic “Big Man” societies of Melanesia (e.g., Sahlins 1963). As a Big Man’s coalition increased in size, so too did his personal prestige, enabling him to attract even more supporters (e.g. Sahlins 1963; Henrich and Gil-White 2001).

As an *explanation* for long-tail phenomena, however, the preferential attachment models are inherently unsatisfying, because preferential attachment is precisely what we would want to explain; with it already built into the model, the long-tail, rich-get-richer result is no surprise. It does not *explain* wealth inequality, for example, to assume the existence of an interest rate (e.g., Bouchaud and Mézard 2000; Burda et al. 2002). Some modellers have produced long-tailed wealth distributions by imposing the interactions onto a power-law network (e.g., Wright 2005), or using a branching “family tree” with inheritance (Coelho et al. 2005), but effectively this presents the same problem of pre-designing the long-tailed distribution into the results.

A key problem with preferential attachment models is the constraint on how change occurs over time: in essential form, growth by proportionate advantage means those with less can never overtake those with more. In real cultural arenas, the rising and falling of fortunes is the norm, even where there exists strong positive feedback for success. In the case of cities ranked by population size, which follow a heavily right-skewed distribution, Batty (2006) recently devised “rank clocks” to demonstrate their continual turnover in the rankings, such that while power law distributions remain stable in form over time, they nonetheless undergo continual turnover in composition of individual components.

This turnover is part of the essence of an open, non-equilibrium system, but most preferential attachment network models have difficulty accounting for such turnover. This is due to the restrictiveness of the network analogy for non-network phenomena compared with the majority of real-world cases, where the idea and the agent are separable. In preferential attachment network models with aging (e.g. Dorogovtsev et al. 2002), the probability of the choice itself (node) diminishes with its age. This is inappropriate for the adoption of ideas, where an idea does not go extinct because it is old, but because no one uses it anymore. It is more appropriate to model the limited memory of the choosers, rather than the aging of the choices. A recent network model (Hajara and Sen 2006) does, in fact, briefly explore limited memory, but still equates nodes with ideas.

7. Random copying with innovation

Many of the conceptual problems that arise from social physics approaches, such as preferential attachment, can be overcome by models based on evolutionary principles, in which the agents are still operating much closer to the zero intelligence null model than they are to the model of full rationality.

These are particularly useful in situations where each individual has a very large variety of options. A system of agents with many options presents a different category of problem not suited well to either network models or binary decision models. An evolutionary model can generate a wide range of the right-skew distributions observed in cultural, economic and social situations from different combinations of its very few parameters (Bentley et al. 2009).

Fashions and popular culture – often behaviors that are not essential to survival -- are often best considered as “neutral” traits, in that what is chosen has no *inherent* value relative to other options (Gillespie 1998; Hahn and Bentley 2003; Lipo et al. 1997; Shennan & Wilkinson 2001). This has been formalised in a model, akin to the neutral-trait model of population genetics, for popular culture change (e.g., Koerper & Stickel 1980; Neiman 1995; Shennan & Wilkinson 2001; Bentley & Shennan 2003). The model simply allows us to ask, what if everyone simply copied each other, with occasional innovation?

The model assumes that, for example, the decision of whether to choose a pug, a terrier or a poodle as a pet primarily depends on the relative popularities of the breeds, rather than the inherent properties of the dogs themselves (size, temperament, etc.). As with Watts’ (2002) threshold model, there is therefore no inherent “fitness” assigned to any of the choices, because actions depend entirely on what *others* are doing. The difference is that the network is never static – since any agent can copy any other, the “network” of who copied whom changes at every time step of the model.

Derived from evolutionary theory – designed specifically to understand change through time – random drift is a much more powerful analogue than preferential attachment for a range of collective behaviors. As Kimura and Crow (1964) showed decades ago, the simple evolutionary model of random drift produces power law frequency distributions. This has also been demonstrated analytically (Evans 2007) and through computer simulation (Hahn and Bentley 2003; Bentley et al. 2004). Generally, the more popular a variant is, the more likely it will be copied again. Unlike a pre-defined rule, this proportionate advantage emerges from the model, and as a general tendency rather than a fixed behavior. As a result, the constituents of the resulting power law distribution are in continual flux, in just the way that Batty (2006) observed, for the model as well as in ranked lists of pop-culture elements in the real world (Bentley et al. 2007).

The turnover results from a balance of innovation, which introduces new ideas, and random drift, by which variation is lost through sampling. New ideas become highly popular by chance alone, and then over time become replaced by others, all through drift. The random-copying model predicts turnover will increase with the fraction of innovators in the population but (somewhat counter-intuitively) not on population size itself. Both these predictions are reflected in real-world data (Bentley et al. 2007).

Hence, without any pre-defined rules for such effects, random copying can account for (a) power law distributions and the “rich get richer” effect of the most popular being the most likely to be copied; (b) continual turnover, with any new variant having a finite (if usually small) chance of becoming highly popular and very popular variants eventually falling in popularity; and (c) network dynamics, in that the network is never the same from one time step to the next. With the simple parameters of population size and fraction of innovators, the model can thus be used to predict change rates as well as potentially distinguish copying from other forms of collective behavior.

8. The units of agency

Because the explanation of what is going on depends critically on the scale of the observation, the units of variation and evolution are a critical variable in evolutionary studies applied to human society (O’Brien and Lyman 2002).

Change through **time** is of course the essence of evolution. Human activity and cultural evolution occur on different temporal scales, including *multi-generational* (such as evolving languages, cultures, built environments, or settlement patterns), *inter-generational* (such as wealth inheritance, education, or the migration upon reaching adulthood), *monthly/annual* (such as many economic and political processes), *daily* (e.g. traffic congestion, peak electricity use) and *momentary* (e.g. the spread of a rumor, web-based interactions, or human aggression).

The temporal scale very strongly determines the appropriate rules to assign to agents. For example, particle-like diffusion (random walks) might be appropriate for the spread of migrants across a continent over multi- or inter-generational time scales (e.g. Ackland et al. 2007), but we would not use diffusion to model, say, an individual’s deliberate migration to a more desirable community (Boyd and Richerson 2009). Many data now exist to characterize people’s movements on these different time scales, from historical/census data to mobile phone records.

Evolution also occurs in our second dimension, **space**, but this does not refer merely to physical space. Sewall Wright (1932) proposed the concept of fitness landscapes, which Kauffman (1995) parlayed into agent-based models of “NK” fitness landscapes, where the space was effectively a medium for interactions. Space may also refer to information, and Daniel Dennett (1996)

popularized the concept of a “design space” in which evolution operates. Speaking more specifically of human actions, we can see these taking place in different spatial categories, such as *geographic space* (e.g., commuters, migrants); *network space* (e.g., friendships, hyperlinks, hierarchical relations, telephone calls); *relational space* (to represent non-discrete entities that may overlap); and *abstract space*, for intractable entities without real-world coordinates (e.g., popular ideas in a population).

A third crucial dimension is **demography**. Human interactions are fundamentally dependent on the number of actors and their heterogeneity. Basic scales of difference can include the *neurological* (where multiple intentions may compete within one mind (Daw et al. 2006; Kable et al. 2007)); the *isolated individual*, where a 'rational' action is selected in the absence of social influence (Gintis 2007); the *small group* (few enough to be known individually (Dunbar and Shultz 2007; Winterhalder and Smith 2000)); and the *population*, where there are enough people so that individuals are not recognized (Bentley et al. 2007; Powell et al. 2009, Watts 2002).

In demographic terms, there may even be a crucial ‘tipping point’ in the effect of group size on aggregate behavior. Robin Dunbar (1993) famously proposed that the number of social relations our brains have evolved to handle is on the order of about 150. This number is roughly equal to the typical size of contemporary hunter-gather communities, and so probably reflects the demands of the social environment in which our ancestors evolved. Above this order of magnitude group size, a qualitative transition in social dynamics takes place above it (e.g. Chagnon 1975; *The Economist* 2009), such that we can expect societies of many people to behave differently than groups or small communities of a few people.

9. Discussion and conclusions

The above may seem a formidable and extensive list of ways in which the simple null model of zero intelligence agents may need to be modified in order to be really useful in the social sciences, and specifically within anthropology. However, as we noted in the introduction, the particular modifications will be context dependent. We are essentially setting out a tool kit from which modelers can select to custom-build their model to address a specific problem. None of the potential modifications in themselves are placing the model of behavior closer to that of full rationality rather than that of zero intelligence.

However, it does seem clear that there is much to be gained by making such modifications. Direct analogies between people and particles (or network nodes) often stray too far from reality to justify the convenience for the models. This is not to lament that “people are more complex than particles,” but rather to highlight occasions where humans interactions are different. Many physicists would rightfully fear that too much detail compromises the

simplifying utility of a *model* – like comedian Stephen Wright’s joke about buying an “actual size” map of the United States.

But the rationale for better collaboration between social physics and anthropology is compelling. On one hand, physicists may characterize a certain category of collective action, but use a flawed assumption about the individual human behavior. On the other hand, anthropologists have such a deep record of individual behavior and its seemingly infinite variability that understanding collective effects can seem impossible.

Aside from refining the basics of human behavior, collaboration with anthropologists would also broaden the new paradigm toward applicability across multiple diverse disciplines, with models based more on real-world human behavior, as opposed to physicists’ own academic society -- citations, collaborations, etc. (e.g., Redner 1998; Newman 2001; Guimera et al. 2005; Wuchty *et al.* 2007). Unique insight was achieved, for example, when Watts (2002, also Dodds et al. 2005) incorporated the essence of classic psychology of conformity thresholds (Asch 1955) into his network models of information cascades, or when Newman (2002; Newman and Leicht 2007) invoked the sociological concept of assortative mixing into his social network models.

In return, culture evolution studies could gain substantially from more detailed collaboration with physicists, in at least these three ways:

- 1) A sophisticated means of dealing with complexity and heterogeneity, without resorting to simplified assumptions of equilibrium or optimality.
- 2) A view of large scale emergent effects (social physics) through individual-scale behavior (anthropology)
- 3) An insight into “tipping points,” or abrupt transitions at the collective level resulting from nondescript individual interactions

A proper collaboration between social scientists and physicists would use existing social science data to provide the minimum of behavioral rules necessary to do complexity modeling for real understanding of community issues and policy. The basic action rules for a model must be appropriate to the scale and the dynamics of the problem, yet not be too complex, lest the model's results have little predictive/explanatory meaning.

References

Adamic, A.L. and B.A. Huberman (2000). The nature of markets on the World Wide Web. *Quarterly Journal of Electronic Commerce* 1: 5-12.

- Ackland, G.J. M. Signitzer, K. Stratford, and M.H. Cohen (2007). Cultural hitchhiking on the wave of advance of beneficial technologies. *Proceedings of the National Academy of Sciences USA* 104: 8714-8719.
- Akerlof, G.A. (1970) The market for lemons: quality uncertainty and the market mechanism, *Quarterly Journal of Economics*, 84, 488-500
- Albert R. and A-L Barabási (2002). Statistical mechanics of complex networks. *Reviews of Modern Physics* 74:47-97.
- Amaral, L.A.N., S.V. Buldyrev, S. Havlin, M.A. Salinger and H.E. Stanley (1998). Power law scaling for a system of interacting units with complex internal structure. *Physical Review Letters* 80: 1385-1388.
- Anderson, P.W. (1972) More is different. *Science* 177: 393-396.
- Anderson, C. (2006) *The Long Tail :Why the Future of Business Is Selling Less of More*. New York: Hyperion.
- Asch, S.E. 1955. Opinions and social pressure. *Scientific American* 193: 31-5.
- Ball, P. 2004. *Critical Mass: How One Thing Leads to Another*. William Heinemann.
- Barabási, A.-L. (2005). Network theory—the emergence of the creative enterprise. *Science* 308: 639-641.
- Barabási, A.-L. and E. Bonabeau (2003). Scale free networks. *Scientific American* 288:50-59.
- Batty, M. (2006). Rank clocks. *Nature* 444: 592-596.
- Bentley, R.A. (2003). Scale-free network growth and social inequality. In (R.A. Bentley & H.D.G. Maschner, eds.) *Complex Systems and Archaeology*: 27-42. Salt Lake City: University of Utah Press.
- Bentley, R.A. & P. Ormerod (2009). Tradition and fashion in consumer choice. *Scottish Journal of Political Economy* 56: 371-381.
- Bentley, R.A., P. Ormerod and M. Batty (2009). An evolutionary model of long tailed distributions in the social sciences. arXiv preprints 0903.2533.
- Bentley, R.A, Lipo, C.P., Herzog, H.A. & Hahn, M.W. 2007. Regular rates of popular culture change reflect random copying. *Evolution and Human Behavior* 28: 151 – 158.
- Bentley, R.A. & S.J. Shennan 2003. Cultural evolution and stochastic network growth. *American Antiquity* 68: 459-485.
- Bentley, R.A., Hahn, M.W. & Shennan, S.J. 2004. Random drift and culture change. *Proceedings of the Royal Society B* 271: 1443-1450.
- Bentley, R.A. and S.J. Shennan 2005. Random copying and cultural evolution (Comment). *Science* 309: 877-879.
- Bentley R.A. and H.D.G. Maschner 2001. Stylistic evolution as a self-organized critical phenomenon? *Journal of Archaeological Method and Theory* 8: 35-66
- Berger, J. and G. Le Mens (2009). How adoption speed affects the abandonment of cultural tastes. *Proceedings of the National Academy of Sciences USA* 106: 8146-8150.
- Boone, J. L. and E. A. Smith. (1998). Is it evolution yet? A critique of Evolutionary Archaeology. *Current Anthropology* 39: S141-S173.
- Borgatti, S.P., A. Mehra, D.J. Brass and G. Labianca (2009). Network analysis in the social sciences. *Science* 323: 892-895.
- Bouchaud J.P., and M. Mezard (2000). Wealth condensation in a simple model of economy. *Physica A* 282: 536-545.

- Boyd, R., and P.J. Richerson. 1985. *Culture and the Evolutionary Process*. University of Chicago Press.
- Boyd, R. and P. J. Richerson (2009). Voting with your feet: Payoff biased migration and the evolution of group beneficial behavior. *Journal of Theoretical Biology* 257: 331-339.
- Buchanan, M. (2007). *The Social Atom*. London: Bloomsbury.
- Burda, Z., D. Johnston, J. Jurkiewicz, M. Kaminski, M.A. Nowak, G. Papp and I. Zahed (2002). Wealth condensation in pareto macroeconomies. *Physical Review E* 65: 02610.
- Chagnon, N. (1975). Genealogy, solidarity, and relatedness: Limits to local group size and patterns of fissioning in an expanding population. *Yearbook of Physical Anthropology* 19: 95-110.
- Coelho, R., Z. Neda, J.J. Ramasco and M.A. Santos (2005). A family-network model for wealth distribution in societies. *Physica A* 353: 515-528.
- Collard M., S.J. Shennan and J.J. Tehrani. 2006. Branching, blending and the evolution of cultural similarities and differences among human populations. *Evolution and Human Behavior* 27: 169–184.
- Cook, W., Ormerod, P., 2003. Power law distribution of the frequency of demises of US firms. *Physica A* 324: 207–212
- Cronk, L. and A. Gerkey. 2007. Kinship and descent. In *The Oxford Handbook of Evolutionary Psychology* (R. Dunbar and L. Barrett, eds), pp. 463-478. Oxford University Press.
- Daw, N.D. et al. 2006. Cortical substrates for exploratory decisions in humans. *Nature* 41: 876-9.
- Di Guilmi, C., M. Gallegati and P. Ormerod (2004). Scaling invariant distributions of firms' exit in OECD countries. *Physica A* 334: 267-273.
- Dennett, D. 1996. *Darwin's Dangerous Idea*. New York: Simon and Schuster.
- de Waal, F. (2009). *The Age of Empathy: Nature's Lessons for a Kinder Society*. New York: Random House.
- Dodds, P.S. & Watts, D.J. (2005). A generalized model of social and biological contagion. *Journal of Theoretical Biology* 232, 587-604.
- Dorogovtsev, S.N. and J.F.F. Mendes (2000). Evolution of networks with aging of sites. *Physical Review E* 62: 1842.
- Dunbar, R.I.M. (1993) Coevolution of neocortex size, group size and language in humans. *Behavioral and Brain Sciences* 16:681--735.
- Dunbar, R.I.M. and S. Shultz (2007). Evolution in the social brain. *Science* 317: 1344-1347.
- Dunnell, R.C. 1978. Style and function, A fundamental dichotomy. *American Antiquity* 43: 192-202.
- Durkheim, É. 1912 *The Elementary Forms of the Religious Life*. New York: Free Press.
- Dyson-Hudson, N. 1966. *Karimojong Politics*. Oxford: Clarendon Press.
- Economist, The* (2009) Primates on Facebook: Even online, the neocortex is the limit. 26 Feb.
- Eerkens, J., Lipo, C.P. (2005). Cultural transmission, copying errors, and the generation of variation in material culture and the archaeological record. *Journal of Anthropological Archaeology* 24: 316-334.
- Evans, TS (2007). Exact solutions for network rewiring models. *European Physical Journal*

- B* 56: 65-69.
- Farkas, I., D. Helbing and T. Vicsek. 2002. Mexican waves in an excitable medium. *Nature* 419: 131-132.
- Farmer, J. Doyne and John Geanakoplos (2009). The virtues and vices of equilibrium and the future of financial economics. *Complexity* 14: S11-S38
- Farmer, J.D., P. Patelli and I.I. Zovko (2005). The predictive power of zero intelligence in financial markets. *Proceedings of the National Academy of Sciences USA* 102: 2254-2259.
- Fortunato, S. (2005). Damage spreading and opinion dynamics on scale-free networks. *Physica A* 348: 683-690.
- Fratkin, E.M. 1989. Household variation and gender inequality in Ariaal Rendille pastoral production: results of a stratified time allocation survey. *American Anthropologist* 91:45-55.
- Gabaix, X., P. Gopikrishnan, V. Plerou, and H.E. Stanley (2006) Institutional investors and stock market volatility, *Quarterly Journal of Economics* 121: 461-504.
- Gallegati, M. et al. 2006. Worrying trends in econophysics. *Physica A* 370: 1-6.
- Gillespie, J. H. (1998). *Population Genetics: A Concise Guide*. Baltimore (MD): Johns Hopkins University Press.
- Gintis, H. 2007 A Framework for the unification of the behavioral sciences. *Behavioral and Brain Sciences* 30:1-61.
- Guimerà, R., B. Uzzi, J. Spiro & L.A.N. Amaral (2005). Team assembly mechanisms determine collaboration network structure and team performance. *Science* 308: 697-702.
- Hahn, M.W. & R. A. Bentley 2003. Drift as a mechanism for cultural change: An example from baby names. *Proceedings of the Royal Society B* 270: S1-S4.
- Hajra, B. & P. Sen (2006). Modelling aging characteristics in citation networks. *Physica A* 368: 575-582.
- Hanson, F.A. (2004). The new Superorganic. *Current Anthropology* 45: 467-482.
- Hayden, B. (2001). The dynamics of wealth and poverty in the transegalitarian societies of Southeast Asia. *Antiquity* 75: 571-581.
- Helbing, Dirk, Illes Farkas, and Tamás Vicsek. 2000. Simulating dynamical features of escape panic. *Nature* 407: 487-490.
- Henrich, J., R. Boyd, S. Bowles, H. Gintis, E. Fehr, C. Camerer, R. McElreath, M. Gurven, K. Hill, A. Barr, J. Ensminger, D. Tracer, F. Marlow, J. Patton, M. Alvard, F. Gil-White and N. Henrich (2005). 'Economic man' in cross-cultural perspective: Ethnography and experiments from 15 small-scale societies. *Behavioral and Brain Sciences* 28: 795-855.
- Henrich J. and R. Boyd (2008). Division of labor, economic specialization and the evolution of social stratification. *Current Anthropology* 49: 715-724.
- Henrich, J. and F.J. Gil-White (2001). The evolution of prestige: freely conferred deference as a mechanism for enhancing the benefits of cultural transmission. *Evolution and Human Behavior* 22: 165-196.
- Herzog, H.A., R. A. Bentley & M. W. Hahn (2004). Random drift and large shifts in popularity of dog breeds. *Proceedings of the Royal Society B* 271: S353-S356.

- Hrdy, S. B. (2009). *Mothers and Others: The Evolutionary Origins of Mutual Understanding*. Harvard University Press.
- Huberman, B.A., P.L.T. Pirolli, J.E. Pitkow and R.M. Lukose (1998). Strong regularities in World Wide Web surfing. *Science* 280: 95-97.
- Hufnagel L., D. Brockmann and T. Geisel (2006) The scaling laws of human travel. *Nature* 439:462– 465.
- Kable, J.W., P.W. Glimcher 2007. The neural correlates of subjective value during intertemporal choice. *Nature Neuroscience* 10: 1625-33.
- Kauffman, S. (1995). *At Home in the Universe: The Search for Laws of Self-Organization and Complexity*. Oxford University Press.
- Kimura, M., & Crow, J. F. (1964) The number of alleles that can be maintained in a finite population. *Genetics* 49:725-738.
- Knappett, C., T. Evans and R. Rivers (2008). Modelling maritime interaction in the Aegean Bronze Age. *Antiquity* 82: 1009–1024
- Koerper, H.C. and E.G. Stickel 1980. Cultural Drift: A primary process of culture change. *Journal of Anthropological Research* 36: 463-469.
- Ijiri, Y. and H. Simon (1964). Business firm growth and size. *American Economic Review* 54: 77-89.
- Lipo, C.P., M.E. Madsen, R.C. Dunnell and T. Hunt (1997). Population Structure, Cultural Transmission and Frequency Seriation. *Journal of Anthropological Archaeology* 16:301-333.
- Melbourne, B.A. and A. Hastings (2009). Highly variable spread rates in replicated Biological invasions: Fundamental limits to predictability. *Science* 325: 1536-1539.
- Mirowski, P (2002). *Machine Dreams: Economics Becomes a Cyborg Science*. CUP.
- Neiman, F.D. (1995) Stylistic variation in evolutionary perspective. *American Antiquity* 60:7-36.
- Nettle, D. (2009). Beyond nature versus culture: Cultural variation as an evolved characteristic. *Journal of the Royal Anthropological Institute* 15: 223-40
- Newman, M.E.J. and E. A. Leicht (2007). Mixture models and exploratory analysis in networks. *Proceedings of the National Academy of Sciences USA* 104: 9564-9
- Newman M.E.J. (1996). Self-organized criticality, evolution and the fossil extinction record. *Proceedings of the Royal Society B* 263: 1605-1610.
- Newman, M.E.J. and G.J. Eble (1999). Power spectra of extinction in the fossil record. *Proceedings of the Royal Society B* 266: 1267-1270
- Newman, M.E.J. (2001). The structure of scientific collaboration networks. *Proceedings of the National Academy of Sciences USA* 98: 404-409.
- Newman M.E.J. (2002). Assortative mixing in networks. *Physical Review Letters* 89: 208701
- Newman MEJ, Barabási A-L, Watts DJ (2006) *The Structure and Dynamics of Networks*. Princeton University Press.
- Newman, M.E.J. and E. A. Leicht (2007). Mixture models and exploratory analysis in networks. *Proceedings of the National Academy of Sciences USA* 104: 9564 –9569.
- O'Brien, M.J. and R.L. Lyman, eds. 1996. *Evolutionary Archaeology: Theory and Application*. Salt Lake City: University of Utah Press.
- O'Brien, M.J., and R.L. Lyman (2002). The epistemological nature of archaeological units.

- Anthropological Theory* 11: 26–36.
- Ofek, Haim (2001). *Second Nature: Economic Origins of Human Evolution*. Cambridge University Press.
- Ormerod, P. and R. Colbaugh (2006). Cascades of failure and extinction in evolving social networks. *Journal of Artificial Societies and Social Simulation* 9, 4
- Ormerod, P. and B. Rosewell (2003). What can firms know? plenary talk, *Proc. NAACSOS Conference*, Pittsburgh
- Pareto, V. (1907). *Manuel d'Économie Politique*. Paris: Giard et Brière.
- Powell, A., S. Shennan, M.G. Thomas 2009. Late pleistocene demography and the appearance of modern human behavior. *Science* 324: 1298-1301.
- Price, D. de la S. 1965. Networks of scientific papers. *Science* 149: 510-515.
- Redner, S. (1998). How popular is your paper? An empirical study of the citation distribution. *The European Physical Journal B* 4: 131-134.
- Sahlins, M. (1963). Poor man, rich man, big man, chief: Political types in Melanesia and Polynesia. *Comparative Studies in Society and History* 5: 285-303.
- Salganik, M.J., P.S. Dodds and D.J. Watts, 2006. Experimental study of inequality and unpredictability in an artificial cultural market *Science* 311: 854-856.
- Salzman, P.C. 1999. Is inequality universal? *Current Anthropology* 40:31-61.
- Shennan S. J. and J.R. Wilkinson (2001) Ceramic style change and neutral evolution: A case study from Neolithic Europe. *American Antiquity* 66:577-594
- Smith, V. (2003). Constructivist and Ecological Rationality in Economics *American Economic Review*, 93, 465-508
- Sneppen, K., P. Bak, H. Flyvbjerg, and M.H. Jensen (1995). Evolution as a self-organized critical phenomenon. *Proceedings of the National Academy of Sciences USA* 92: 5209-5213.
- Solé, R.V., S.C. Manrubia, M. Benton and P. Bak (1997). Self-similarity of extinction statistics in the fossil record. *Nature* 388: 764-767.
- Solé, R.V., R. V., Corominas, B., Valverde, S., and Steels, L. (2005) Language Networks: their structure, function and evolution. Santa Fe Institute Working Paper 05-12-042.
- Stanley, M.H.R., L.A.N. Amaral, S.V. Buldyrev, S. Halvin, H. Leschhorn, P. Maass, M.A. Salinger and H.E. Stanley (1996). Scaling behaviour in the growth of companies. *Nature* 379: 804-806.
- Surowiecki, J (2004). *The Wisdom of Crowds: Why the Many Are Smarter than the Few*. London: Abacus.
- Valverde, Sergi (2004). A new science for a connected world. *Current Anthropology* 45: 565-566.
- Watts, D.J. 2002. A simple model of global cascades on random networks. *Proceedings of the National Academy USA* 99: 5766-71.
- West, B. J. and B. Deering (1995). *The Lure of Modern Science: Fractal Thinking*. Singapore: World Scientific.
- Winterhalder, B. and E. A. Smith. 1999. Analyzing adaptive strategies: Human behavioral ecology at twenty-five. *Evolutionary Anthropology* 9:51-72.
- Wright, S. (1932). The roles of mutation, inbreeding, crossbreeding, and selection in

- evolution. *Proceedings of the Sixth International Congress on Genetics* 1: 355-366.
- Wright, I. (2005). The social architecture of capitalism. *Physica A* 346: 589–620.
- Wuchty, S., B.F. Jones and B. Uzzi (2007). The increasing dominance of teams in production of knowledge. *Science* 316: 1036-1039.