Modelling Creative Innovation

Paul Ormerod and R. Alexander Bentley

Abstract

The economic concept of rationality seems inappropriate in the context of creative innovation, because of its assumption that the tastes and preferences of agents are fixed. The concept of copying, of imitating the behaviour of others, has equal claim to the description 'rational' in an innovative context. Models of 'binary choices with externalities' are predicated on copying and potentially show us not only why most innovations fail, but also why big social changes do not necessarily require big causes. In the 'Long-tail' world of a huge range of choice, however, many choices are not 'binary, either-or'. In the long-tail world, popular choices tend to become more popular, but not forever, as innovation drives a constant turnover in the popularity rankings. A very simple model of 'neutral' copying with occasional originality of choice can explain real-world patterns of long-tail distributions under continual turnover.

1. Introduction

Cultural science should seek an evolutionary understanding of a knowledge-based society past and present, with one of its goals being to map possible future scenarios to which public policy and businesses must adapt. Identifying what proceeds in predictable directions, as opposed to drifting upon the tides of fashion, would be of great utility in understanding the evolution of creative innovation.

Seemingly an oxymoron, 'cultural science' is inherently interdisciplinary. An important issue in cultural science remains an old one: the relationship of formal quantitative modelling to more descriptive, domain-specific knowledge. Alfred Marshall, who helped develop the Department of Economics at Cambridge in the 19[th] century (and had obtained the second best maths result at Cambridge in his final exams) advised us to "Use mathematics as shorthand language, rather than as an engine of inquiry," followed by these steps

- Keep to them till you have done.

- Translate into English.

- Then illustrate by examples that are important in real life

- Burn the mathematics.

- If you can't succeed in 4, burn 3. This I do often.'

Marshall's views can be applied more generally to maths and the social sciences as a whole; if formal modelling is used to complement domain-specific knowledge, then the two together can

gain more insight than either taken on its own. Modelling is helpful in describing the key features of a problem. In the social sciences, however, the parameters and features of the model, once identified, then need to be capable of being accounted for by plausible arguments.

As an example of this interdisciplinarity, we focus on describing a particular category of modelling, namely representing the diffusion of behaviour across social networks. Our use of the term 'behaviour' refers to a wide range, from choice in consumer goods markets to the spread of much more abstract concepts such as cultural norms and trust. Although it has a long history (e.g. Tarde 1895), the metaphor of ideas spreading through 'contagion' or diffusion has become commonplace in the modern world of viral marketing and rapid spread of global trends and fashions. Malcolm Gladwell (2000) published a bestseller based on this premise, and fittingly, the metaphor itself has spread and become popular in academia. Richardson (2000), for example, significantly influenced public policy literature by describing ideas spreading as a 'policy virus'.

2. How agents choose

2.1   Rational choice in economic theory

Since Marshall's time, most quantitative social science has used a standard assumption that human decisions are routinely made by individual, rational economic agents. Agents are postulated to operate as autonomous individuals, each with his or her fixed set of tastes and preferences. The agent gathers all available information relevant to a particular decision, and then process it so as to make the best possible choice – the 'optimal' one – within its fixed set of tastes and preferences.

In certain specific contexts, the pure rational agent model gives a good understanding of the world. For example, electronic supermarket sales data show that relative prices are a very important determinant of brand choice, as the rational choice model suggests. Here the assumption of fixed tastes and preferences (for price) seems a reasonable approximation to reality. The consumer has already decided to spend rather than save, has already decided to buy food in a supermarket rather than, say, a meal in a restaurant, preferences on different brands of breakfast cereal or whatever are already formed. Further, the information to be processed in buying cereal or soap is relatively constrained, and easily rectified if the wrong decision is made.

The concept of bounded rationality (Akerlof 1970) extends the empirical validity of this model considerably. In relaxing the assumption of complete information, bounded rationality can lead to different outcomes, even as agents are still assumed to make the best possible decision (given their fixed tastes and preferences) from limited information, with different groups of agents possibly having different information. An example that Akerlof gave is the market for second hand cars, where sellers typically possesses more information than buyers about the vehicles.

2.2 Copying

The economic concept of rationality, however, even bounded rationality, seems ill-suited to

issues of creative innovation. Fixed tastes and preferences is not the obvious model for situations in which tastes and preferences necessarily evolve, as agents are confronted with innovations. Culture, by definition as a shared set of ideas and practices, cannot arise through autonomous, rational agents. In other words, the assumed world of autonomous individuals, each with his or her fixed set of tastes and preferences, could be seen as quite contrary to common experience and even the very nature of culture itself.

In popular culture, the thought of ideas spreading through social diffusion is entirely commonplace, and modern English is filled with phrases describing ideas that "spread like wildfire" or rumours getting round, or the changes to a story as it spreads via "Chinese whispers'. Copying is so normal that great 19[th] novels were often set amidst 'society' so stifling that a protagonist's struggle for originality was sufficient to drive the narrative, for Nietzsche and Dostoyevsky, for example. Even before writing existed, human knowledge was copied; through oral traditions, parenting, and craft apprenticeships.

An alternative is a completely different 'null model' of rationality, where agents rely primarily on copying other agents. In the simplest form of a copying model, agents pay no attention to the objective properties of the various choices available. Instead they simply copy the behaviour of others, according to a variety of possible rules (Laland 2004). A classic illustration of the operation of this postulate in practice is given by the conformity experiments of Solomon Asch (1953). The behaviour of an agent tends to become more similar to the group of which he or she is a member. This could happen, for example, because the agent believes the group to have better information than he or she does, or from a simple desire to conform to group norms (Boyd et al. 2010, 2011). In other contexts, peer acceptance can be the behavioural motivation by which agents become more similar to their peers. Obesity, for example, has spread not just due to modern fatty diets but through social diffusion as well (Christakis and Fowler 2007), and binge drinking similarly has a strong social network component (Ormerod and Wiltshire 2009).

We agree with Rendell et al. (2010: 208) that copying – or 'social learning' – 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'. They organised a computer tournament in which entrants submitted strategies specifying how to use social learning and other alternatives to acquire adaptive behaviour in a complex environment. While many predicted the winning strategy to be some ideal combination of in individual and social learning, it turned out that the 'winning strategy relied nearly exclusively on social learning'.

We should be careful, however, not to shift to the opposite pole of thinking, and automatically attribute all outcomes in modern communications contexts to the process of copying. For example, in analysing a global instant messaging network of 27.4 million users for longitudinal, demographic, and geographic patterns, Aral et al. (2009) argued that community similarity of behaviour can be attributed more to homophily -- the tendency for people with similar interests to wind up in the same group or space -- than to contagion.

Nevertheless, models for creative innovation that emphasize copying are likely to prove more powerful than those which rely on agent rationality. Rational behaviour postulate that tastes and preferences are fixed which cannot explain cultural differences, i.e., endogenous tastes and preferences that contrast between groups (e.g., Potts et. al. 2008).

3. Binary choice with externalities

Thomas Schelling (1973) proposed a classic model in which dramatic segregation of groups could follow from agents faced a binary choice with only very weak biases of their own. Since then, binary choice models have been given wide application. Binary choice simply means an agent adopts a behaviour or not, as in voting for candidate A or someone else, choosing a particular restaurant or rejecting it; believing or not that stock prices will go up. Schelling's crucial advance was to show that the decision of any given agent may alter the choice made by other agents – the decisions have 'externalities'. All subsequent models of 'binary choice with externalities' invoke some process of copying. Agents are often connected on some form of network, and any given agent is aware of the choices of those agents to which he or she is connected (the relevant peer group).

Duncan Watts (2002), who is now Director of the Human Social Dynamics group at yahoo.com, made elegant use of this model category. Watts (2002) reasonably assumed that agents will differ in their willingness to be persuaded to adopt (copy) a different mode of behaviour from their neighbours. Each agent is allocated a 'threshold' that represents the proportion of other agent it pays attention to who must adopt the behaviour before it does as well. So if an agent pays attention to five others, and has a threshold of 10%, then it only needs one of the five (= 20%) to choose the alternative for the agent to copy this behaviour. If the agent had a threshold above 80%, however, it would not adopt the behaviour until all five of the others had (Watts 2002). Watts (2002) assigned these thresholds to agents randomly, from a uniform distribution ranging from 0 to 100% (In a practical situation, of course, we might estimate this range more precisely, through survey evidence on their persuadability, for example).

Despite its apparent simplicity, the Watts (2002) model lends two powerful insights into how creative innovation spreads across any given population.

- **Most innovations fail.** Watts (2002) found that most agent decisions did not cascade very far among the other agents. This fits real-world evidence (Ormerod 2005).

- **Big effects do not require big causes.** Small initial disturbances of identical size can have dramatically different consequences.

We illustrate the second point as follows. We populate the Watts (2002) model with 1,000 agents, connected on a 'small world' network (Watts and Strogatz 1998) – although similar results random or scale-free networks yield similar results. Each agent can be in one of two states of the world, A or B. Initially all agents are in state A, with their thresholds drawn uniformly in the range 0 to 100%. A small number of agents, in this case 20, is selected at random to choose B not A. We then allow the model to run and observe the total number of agents which eventually switch to B. We call this number as a proportion of the total number of agents (1,000) the 'size of the cascade'. By running the model repeatedly many times, and each time choose the seeds at random, and plot the resulting distribution of the size of the cascade. The result in Figure 1 shows that the percentage that switches to B from A is most of the time considerably less than 100%.
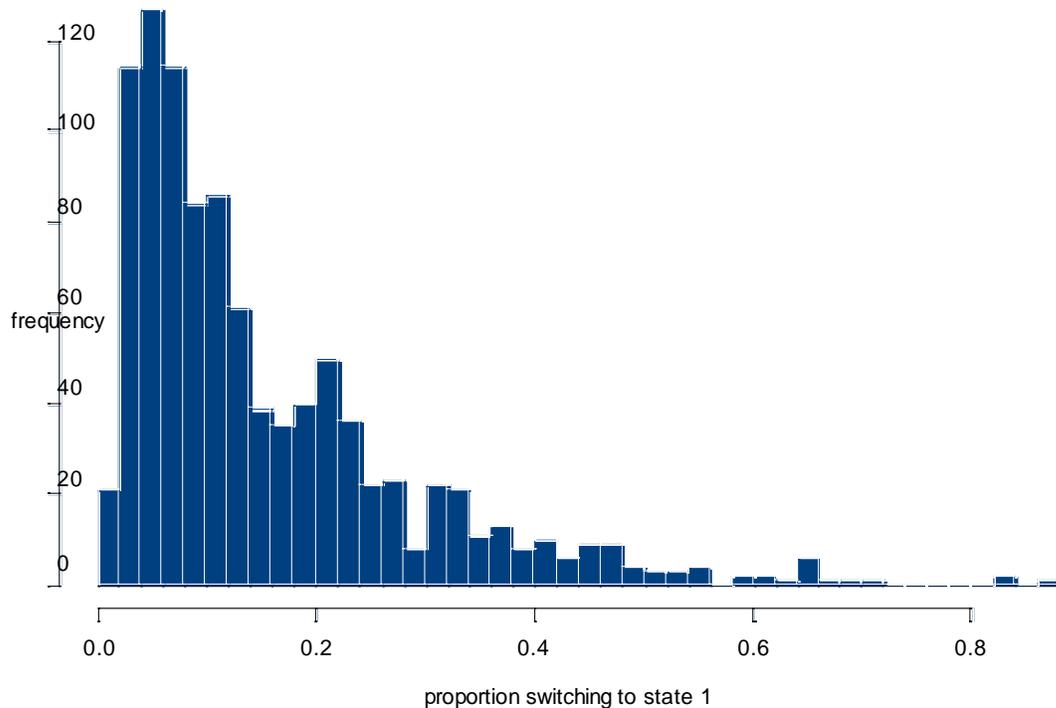
Figure 1   *Cascade size in Watts (2002) model with 1,000 agents with heterogeneous thresholds, 20 agents selected at random as the initial seeds*

In cases involving 'either/or' decisions, such models of 'binary choice with externalities' yield many insights, in cases such as voting (Galam 2007). Voting in a referendum, for example, typically involves a huge range of factors to be considered, often enormously complicated, but the choice in a referendum is nevertheless usually a simple binary one: yes or no.

4.   The Long Tail world

Despite Danish philosopher Søren Kierkegaard's best efforts to suggest so, not all decisions in modern life are *Either-Or*. We now consider cases where agents choose between many alternatives. In the modern, 'long-tail' world (Anderson 2006), people in market economies are faced with of orders of magnitude more choice than previous centuries and millennia (Beinhocker 2006; Ridley 2010). There can be such an overwhelming number of similar consumer items, expert opinions, popular ideas, and other choices that a well-informed, 'rational' decision becomes impossible. Even a particular field of academic research may be characterised by thousands or even tens of thousands of relevant articles. Many of modern choices come from long tail distributions (Anderson 2006) – with some extremely popular choices of 'blockbuster'

proportions and the majority of other ideas relatively unpopular, out in the long tail.

In these cases, there is a particular way of characterising the degree of choice, which Shannon (1948), among the founders of Information Theory, proposed to measure information content. Shannon's measure is a way of measuring the average information content of each different item of choice, where information is measured in bits, as for a computer. Shannon (1951) estimated, for example, that the English language has an average information content of just over 9 bits per word. The idea is that each word helps clarify the sentence by this amount of information. Of course, as George Orwell (1946) pointed out adamantly, people tend to copy "long strips of words" from each other as ready-made clichés, and so the information content of each word gets reduced because so often the missing word is quite predictable. For example, if someone is heard to say "we need a … nuanced understanding of…" it is about 99% certain that the missing word was "more" not "less" (compare Google search results for these two different exact phrases). Hence, Shannon (1951) estimated that that the entropy of English was reduced to only 4.5 bits per word when the 100 letters of the proceeding message was already known.

Shannon's entropy measure gives us some quantitative measure of how drastically choice has increased in the last several decades. By applying this measure to baby names, for example, we find that in 1960 the information entropy for the range of first names given to baby girls was about 8.5, whereas in 2009 it was over 10. This may not seem like much, but the difference is actually logarithmic – about a threefold increase in this dimension of choice. For another example, we can consider case studies of word usage in particular niches of social science and physical science (Bentley 2008) and find that the physical science journal niche (all articles who have cited Barabási and Albert 1999) has an information entropy of about 8.8, whereas in the social science niche (all articles that have cited Bordieu 1977) it is 11.3. This implies that the usage of words in the physical sciences is more predictable – to the social scientists, who use language more liberally, this might mean 'dull', whereas the physicists might say that they are using words to refer consistently to specific phenomena rather than invent new words to say the same thing. In any case, those working in the social sciences could be said to have a fivefold greater range of terminology (or jargon) to choose from.

5. Many alternatives and copying

In such cases of 'choice overload' a good strategy is often to copy others. Among a range of animal groups, from humans to other primates and even fish, copying is a highly effective, time-saving strategy for learning the best way to behave in a given situation (Laland 2004). Conformity, which Solomon Asch (op.cit.), Stanley Milgram (1963) and other social scientists demonstrated so effectively, is copying biased towards the most popular behaviour. Copying what others do works for individuals, because it gives a good chance for group acceptance, and it also works for groups, because it yields group cohesion and coordinated activity (Couzin et al. 2005; Boyd et al. 2010, 2011).

The advantages of copying have been largely missed by game theory, which has most often considered agents as being paired off in 'prisonners dilemma' games. For a large interactive populations, however, copying is often the most effective strategy. This was shown in a recent high-profile tournament of simulated strategies in competition, agents could act independently or

copy other agents after observing their levels of success (Rendell et al. 2010). The tounament demonstrated how copying recent, successful strategies of others is highly adaptive in an environment of a large range of choice that is continually changing.

Of course, there will be a huge range of idiosyncratic, personal reasons people might have for copying. The marketing scientist Andrew Ehrenberg and colleagues (Goodhardt et al. 1984) developed a classic model — termed a "Dirichlet" in honour of the nineteenth-century German mathematician — in which it was assumed that consumers had no inherent brand preference and made choices based on chance and availability. It is thus not a great stretch to assume that people might copy in the same undirected way.

Copying others yields a natural 'rich-get-richer' effect, because the more common something it is, the more likely it is to be copied – even more so if conformity is a social norm. Modelling this 'rich-get-richer' effect can be traced back to Yule (1925), who applied it to the number of species per genus of flowering plants. Simon (1955) applied the approach in the social sciences, examining for example the size distribution of cities or of corporations.

The key concept underlying this approach - 'preferential attachment' is the modern description - is that an agent copies the choices previously made by other agents with a probability which is proportionate to the number of times each of these choices has been selected. So, purely for example, suppose A has been chosen by 50 agents, B by 30 agents and C by 20 agents, then the next agent making a choice will select A with a probability of 50%, B with one of 30%, and C with one of 20%.

As a behavioural rule it has intuitive plausibility. If choice is a difficult problem (too many alternatives or just inherently complicated, etc.), it seems reasonable that an agent will select in proportion to the relative popularity of the alternatives. Ormerod and Roach (2008) give the example of religious choice in England in the middle of the 16th century, showing how this simple model gives results which are compatible with the contemporary historical documentary evidence.

The 'preferential attachment' concept was essentially rediscovered by Barabási and Albert (1999) in a paper which has become widely cited (over 5,000 citations already), especially by physicists, who have taken an interest in problems in the social sciences. When the model is populated by a large number of agents, the statistical distribution of the relative popularity of the various choices converges on a long-tail distribution, which is highly non-Gaussian. A small number of alternatives is selected by many agents, whilst most alternatives attract few choices. There is a substantial and arcane literature on whether empirically observed distributions are best approximated by a particular long-tailed function, a power law, or by some other non-Gaussian distribution (e.g. Laherrère and Sornette 1998; Newman 2005; Mitzenmacher 2008). But for social scientists, it is much more important to know that a distribution is non-Gaussian than it is to make a fetish of the power law, a distribution which in the social sciences has no special significance.

6. The 'neutral' model of cultural evolution

Creative expression relies on social contexts, and symbolic information depends not just on what a style is, but also who adopts it. In the Internet age, this is increasingly envisaged as networks, with the creative outputs (Web pages, videos, academic publications, etc.) as 'nodes', and their interactive influences (cited references, related Web pages) as 'links' (e.g., 15, 24) . Since the late 1990s, researchers, particularly in physics, have applied generalised network analysis toward a range of creative media (Guimerà et al. 2005; Newman 2010; Christakis and Fowler 2009).

Most of these network models invoke some form of a preferential attachment process. Preferential attachment models can fit long-tailed distributions of popularity, but the problem is that these models do not naturally generate any turnover in the rankings of popularity which emerge amongst the various choices. Once a sufficient number of agents are in the model, the most popular choice remains the most popular, the second remains number two, and so on. There are various developments of the basic model to try to rectify this problem, but they essentially involve rather artificial add-ons, such as the additional assumption that the probability of the choice itself diminishes with its age (e.g., Dorogovtsev and Mendes 2000; Hajra and Sen 2006).

In a good model, change should be part of the essence of the process, rather than just a modification. Change, in fact, is central to evolutionary theory, which is strictly the study of how attributes are passed on and modified through time (22). Evolutionary approaches to culture change (20, 22) have included tools from epidemiology (11, 27, 31), network theory (15, 23, 24) and population genetics (1-9, 13).

The key, however, to modelling cultural innovation, lies in the transmission process being essentially one of copying what others do (Mesoudi and Whiten 2008), with creative modifications contributing new ideas that eventually replace old ones through being copied. For this reason a remarkably powerful evolutionary approach allows agents to either copy in relative proportion to popularity, but also allows each agent the option to invent something unique, or to select something which no previous agent has chosen. Formally, an agent innovates with probability μ and chooses according to preferential attachment with probability $(1-\mu)$.

This apparently simple addition makes a dramatic difference to the properties of the model, and in particular turnover in rankings is an inherent and natural feature of the approach. Borrowing from evolutionary biology, it is also known as the 'neutral' model (Neiman 1995; Shennan and Wilkinson 2001) because when operating according to preferential attachment, an agent is indifferent – 'neutral' - to the inherent qualities of the choice which he or she is copying. All that matters is the relative popularity of the choice.

With the neutral model, a rich-get-richer effect occurs, but is not a fixed rule under the undirected copying model, which is characterised by unpredictability and continual flux in the most popular ideas. This turnover and flux reflects real world data, such as the constant rising and falling of cites ranked according to population size (Batty 2006) academic buzzword use (Bentley 2008), dog breed popularity (Herzog et al. 2004), and even birdsongs (Byers et al. 2010). This turnover is often remarkably regular and the model predicts it will be proportional to the square root of the modelled fraction of inventors, but it does not correlate strongly, as one might expect, with population size (Bentley et al. 2007).

A subtle assumption underlying this model, and that of that of simple preferential attachment, is that *all* previous choices which have been made are taken into account when any given agent makes a choice. Whilst this may be a reasonable approximation to reality in situations where long memory is important, such as in decisions by firms to locate in a particular city, it is clearly much less valid in, say, popular culture such as the downloads on YouTube or Flickr.

Bentley et al. (2011) generalise this neutral copying model by introducing a second parameter, that of 'memory'. This specifies how far back previous choices are taken into account at any point in time. It explicitly allows for the extinction of any given choice if no agent has selected it during the number of previous steps specified by the memory parameter. Under this generalisation of the copying model, turnover results from a balance of the introduction of new ideas by the specified minority fraction of innovators, and extinction, through failing to be copied by any agents among the copying majority.

This generalised neutral copying model, despite its parsimonious structure with only two parameters, is surprisingly powerful. It can replicates population-scale patterns one finds in real data — long tailed distributions that undergo continual flux, and unpredictability in what particular idea becomes highly popular . Through the process, most newly invented ideas fail, and only a few lucky ones become inordinately popular. The model can fit almost perfectly the long tail of baby-name popularity, for example.

The extra memory parameter provides flexibility in the results, such that the model can replicate the long tailed distributions of a range of forms to match a variety of real-world data sets (Bentley et al. 2011). Figure 2 illustrates the calibration of the model to a range of data sets, in how the neutral model can fit both the distributions of popularity (Figure 2a), as well as the finite lifespans of the choices among the 'most popular' rankings (Figure 2b).
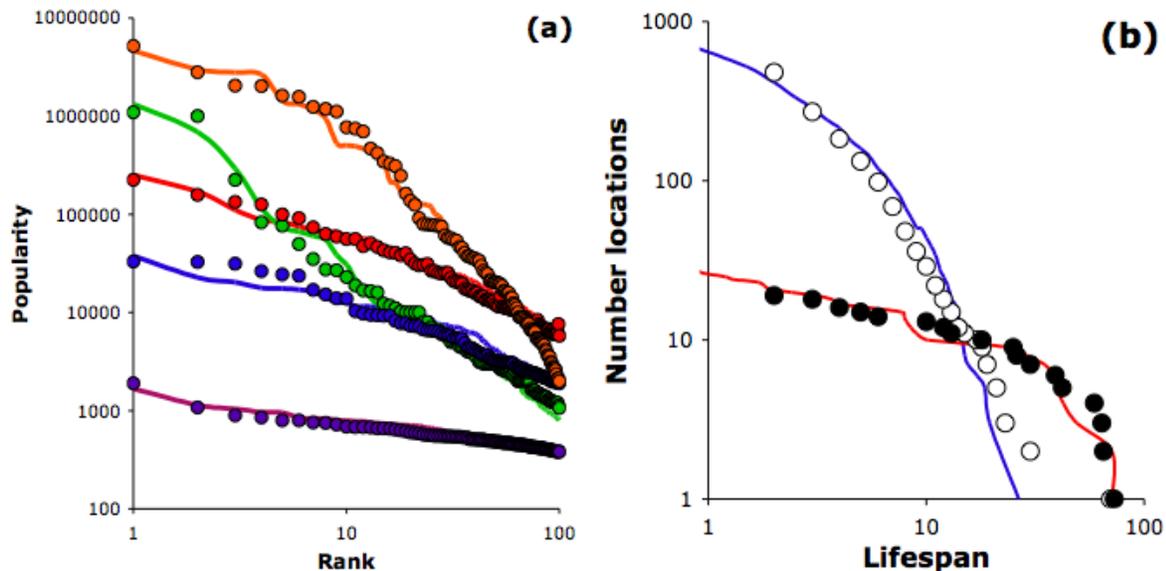
Figure 2. (a) Rank versus popularity for real-world top 100 ranked lists (dots) versus neutral model results (lines). Top 100 lists include: male baby name frequency (per million) in the 1990 US census (blue), RSS feed subscriptions 2001-2008 (orange), English words (red), cited economists 1993-2003 (purple), and religions in thousands of adherents (green). **(b)** Life-spans

of UK Number One Hits (www.theofficialcharts.com) for 1956-2007 (open circles), versus the neutral model (m = 1, μ = 0.1) shown by the blue line), and also years in the Top 5 US boys' names (www.ssa.gov/OACT/babynames), 1907-2006 (filled circles) versus the neutral model (m = 10, μ = 0.001), shown by the red line. Adapted from Bentley et al. 2011: Figs 2b and 4b.

The modified neutral model (Bentley et al. 2011) is specifically designed to model the effect of different innovation rates, in ways that are testable against real-world data. The results shed light on the role of innovation in the acceleration of cultural change and open the door of the humanities to the dynamic science of evolution. The data needed are all realistically obtainable, as they consist of the frequencies (relative popularities) across the range of choices - be they fashion designs, trendy buzzwords, music choices, or other ideas - and how those frequencies changed through time for each of the possible choices.

Specifically, the data needed to apply the model include the (a) frequency distributions (how popularity is distributed among a variety of choices), (b) record of popularity of each individual choice through time and (c) record of the turnover in the most popular variants over time. Each of these measures carries a different prediction under different levels of copying versus independent choice-making (or original invention). As a result, the neutral model can help to characterize different realms of creative media along a 'spectrum' between general and selective copying, as well as provide estimates of innovation rates in each medium (Bentley et al. 2011).

The scale of analysis is a key variable (O'Brien and Lyman 2002). For example, while choices of baby names at the scale of the entire United States are indistinguishable from neutral copying (Hahn and Bentley 2003), different ethnic groups select from different pools of names (Leiberson 2000; Freyer and Levitt 2004), yet within each group, random drift could predominate again (which remains to be studied). Similarly, prehistoric pottery designs may constitute a selected range of variation acceptable to a group, yet characterised by neutral copying within this range (Lipo et al. 1997). In this sense, the neutral-copying model helps us objectively identify the groupings of creative expression, without resorting to our own subjective opinions as to what "meant" what (Neiman 1995; Lipo et al. 1997).

If the copying is neutral, we can predict a lifespan distribution as in Figure 2b, but we cannot predicted exactly *which* new ideas or behaviours will replace the old ones (e.g. Salganik et al. 2006). In other words, we can predict the steady production of new winners, but the randomness means we can't forecast the particular winners themselves. New ideas can become highly popular by random copying alone, and be replaced over time as the next generation of innovations are copied.

The steady turnover under the random copying model (Bentley et al. 2007) could be used to predict turnover rates on bestseller lists, for example. How quickly a list will change depends on the size of the list – a Top 100 changes proportionally faster than a Top 40 – but, surprisingly, the size of the population does not have an impact. Although a larger population means more new ideas, it also means more competition to reach the top, and the two balance each other out: the turnover on bestseller lists remains steady as population size changes. Instead, what actually drives fashion change is innovation – the more innovators per capita, the faster the turnover. Innovators are those who 'pump' new fashions into our world -- most are ignored, but some get

copied. Viral marketing professionals grasp this, identifying a tiny minority of true innovators among a vast majority of copiers.

Looking forward, ours aims for the neutral model are to refine the means of predicting change rates, and ways of distinguishing copying from independent decision-making in collective behaviour. This evolutionary approach models culture evolution as a process of people mainly copying each other, with occasional original invention. For simplicity, this copying process can be envisaged along a spectrum ranging from copying others completely at random, to selective copying in which the attributes of the behaviours are carefully considered. We can then further explore the neutral-copying model with tests of the effects of varying generation time, population size, 'memory', invention rate, and network structure. We would then be in a position to develop this model through varying levels and kinds of independent decision-makers who weigh costs/benefits of their options, subject to biases like novelty, validity, or conformity. Ultimately, we seek to explore how much we can explain through neutral copying processes before we need to resort to 'reasons' such as individual selection for one thing or another. The analysis could then be used to help identify the scale at which selection is exerted, which could help to define the group themselves (Lipo et al. 1997).

7.   Conclusion

Creative expression involves the transmission of information between and among individuals, with the continual production of new ideas, a minority of which give rise to prominent genres or paradigms. As a science for dynamic systems, cultural evolutionary theory (Mesoudi et al. 2006) is ever more useful for a world where this transmission is increasingly mediated by online technologies. Chance interactions, novel opportunities, and unforeseen consequences are becoming normal. The pace of adaptation is accelerating, as human interaction, which underlies creative productivity, becomes compressed though time and space by online technologies that store and disseminate past events. Conversations, ideas and relationships, which once were ephemeral, are now recorded indefinitely and available to a global audience.

For this reason, it is useful to approach changing cultures not as fixed entities, but as "social network markets" (Potts et al. op.cit.) where the division between producers and consumers of culture has dissolved. As Potts and his co-authors described, cultural science has gained renewed relevance, particularly in new conditions of affluence, democratisation, and emancipated consumers amid online networks where identity itself is exchangeable. The interaction of people within this "social network economy" creates a continual flux of ephemeral communities and novel entrepreneurial opportunities, with unforeseen consequences being the norm rather than the exception.

'Social network markets' and other copying models are much more useful than the traditional rational actor model of economics to understanding these phenomena. Indeed, in such contexts this mode of behaviour can be seen as the appropriate model of rationality for agents to adopt. The economic concept of rationality is not general, it is merely one possible way of defining rationality. Alternatives are needed in cultural science, such as the modes of behaviour described in this paper.

All this motivates a study of creative media past and present, to arrive at an idea of their future. Collaboration between cultural studies practitioners and formal modellers is a key part of building cultural science. It is easy to fall into the trap of believing that because formal modelling requires mathematical skills it is some way superior to more descriptive, domain-specific expertise. However, it is not only essential that any mathematical models have overall properties which make sense to domain experts, but that any parameters or key assumptions made in the models can also be justified in this way. Modelling is a tool to take cultural studies forward, and evolutionary theory -- the study of change through time through variation, interaction, selection and drift – is an excellent resource for such models.

References

George A. Akerlof (1970), The market for "lemons"*:* quality, uncertainty and the market mechanism. *Quarterly Journal of Economics*, 84, 488-500.

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

Sinan Aral, Lev Muchnik and Arun Sunarararjan (2009). Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks. *Proceedings of the National Academy of Sciences*, 106, 21544-9.

Solomon E. Asch (1953) 'Effects of group pressure upon the modification and distortion of judgement', in Guetzkow, H (ed.), *Groups, leadership and men: Research in human relations*. New York: Russell and Russel.

Albert-László Barabási and Reka Albert (1999). Emergence of Scaling in Random Networks. Science, 286, 509–512.

Michael Batty (2006). Rank clocks. *Nature*, 444, 592-596.

Eric Beinhocker (2006) The origin of wealth. Harvard Business School Press, Cambridge, MA.

R. Alexander Bentley (2008) Random drift versus selection in academic vocabulary. *PLoS ONE* 3(8):e3057.

R. Alexander Bentley, Carl P. Lipo, Matthew W. Hahn and Harold A. Herzog (2007). Regular rates of popular culture change reflect random copying. *Evolution and Human Behavior* 28, 151-158.

R. Alexander Bentley, Paul Ormerod and Michael Batty (2011). Evolving social influence in large populations. *Behavioral Ecology and Sociobiology*, in press.

Jonah Berger and Gael Le Mens (2009) How adoption speed affects the abandonment of cultural tastes. *Proceedings of the National Academy of Sciences* 106, 8146–8150.

Pierre Bordieu (1977) *Outline of a Theory of Practice*. Cambridge University Press, Cambridge.

Robert Boyd, Peter J. Richerson and Joseph Henrich (2011) Rapid cultural adaptation can facilitate the evolution of large-scale cooperation. *Behavioral Ecology and Sociobiology*, in press.

Robert Boyd, Herbert Gintis and Samuel Bowles (2010) Coordinated punishment of defectors sustains cooperation and can proliferate when rare. *Science* 328:617–620

Bruce E. Byers, Kara L. Belinsky and R. Alexander Bentley RA (2010) Independent cultural evolution of two song traditions in the chestnut-sided warbler. *American Naturalist* 176, 476–489

Nicholas A. Christakis and James H. Fowler (2007), 'The spread of obesity in a large social network over 32 years', *New England Journal of Medicine*, 357, 370–379

Nicholas A. Christakis and James H. Fowler (2009). *Connected*. New York: Little & Brown.

Iain D. Couzin, Jens Krause, Nigel R. Franks and Simon A. Levin (2005). Effective leadership and decision-making in animal groups on the move. *Nature*, 433, 513-516.

Sergey N. Dorogovtsev and J.F.F. Mendes (2000) Evolution of networks with aging of sites. *Physical Review E* 62, 1842.

Roland G. Fryer and Steven D. Levitt (2004). The causes and consequences of distinctively black names. *Quarterly Journal of Economics* 119, 767-805.

Serge Galam (2007). From 2000 Bush-Gore to 2006 Italian elections: voting at fifty-fifty and the contrarian effect. *Quality and Quantity* 41, 579-589.

Malcolm Gladwell (2000) *The Tipping Point*. New York: Little, Brown.

Gerald J. Goodhardt, Andrew S.C. Ehrenberg and Christopher Chatfield (1984) The Dirichlet: a comprehensive model of buying behaviour. *Journal of the Royal Statistical Society A* 147, 621–655

Roger Guimerà, Brian Uzzi, Jarrett Spiro and Luis A. Nunes Amaral (2005) Team assembly mechanisms determine collaboration network structure and team performance. *Science* 308, 697–702.

Matthew W. Hahn, and R. Alexander Bentley (2003). Drift as a mechanism for cultural change: an example from baby names. *Proceedings of the Royal Society B* 270, S1–S4.

Kamalika Basu Hajra and Parongama Sen (2006) Modelling aging characteristics in citation networks. *Physica A* 368, 575–582.

Harold A. Herzog, R. Alexander Bentley and Matthew W. Hahn (2004). Random drift and large shifts in popularity of dog breeds. *Proceedings of the Royal Society B* 271, S353–S356

Jean Laherrère and Didier Sornette (1998). Stretched exponential distributions in nature and economy: "fat tails" with characteristic scales. *European Physics Journal B* 2, 525–539 (1998).

Kevin N. Laland (2004). Social learning strategies. *Learning and Behavior* 32, 4–14.

Stanley Lieberson (2000). *A Matter of Taste: How Names, Fashions, and Culture Change*. Yale University Press, New Haven, CT.

Carl P. Lipo, Mark E. Madsen, Robert C. Dunnell and Tim Hunt (1997) Population structure, cultural transmission and frequency seriation. *Journal of Anthropological Archaeology* 16, 301–333.

Alex Mesoudi and Andrew Whiten (2008). The multiple uses of cultural transmission experiments in understanding cultural evolution. *Philosophical Transactions of the Royal Society B* 363, 3489–3501.

Alex Mesoudi, Andrew Whiten and Kevin N. Laland (2006). Towards a unified science of cultural evolution. *Behavioral and Brain Sciences* 29, 329–383.

Stanley Milgram (1963) Behavioral study of obedience. *Journal of Abnormal Social Psychology* 67, 371–378.

Michael Mitzenmacher (2008). A brief history of generative models for power law and lognormal distributions. *Internet Mathematics* 1: 226-251.

Fraser D. Neiman (1995) Stylistic variation in evolutionary perspective. *American Antiquity* 60:7–36.

Mark E.J. Newman (2005). Power laws, Pareto distributions and Zipf's law. *Contemporary Physics* 46, 323-351.

Mark E.J. Newman (2010). *Networks: An Introduction*. Oxford University Press.

Michael J. O'Brien and R. Lee Lyman (2002). The epistemological nature of archaeological units. *Anthropological Theory* 11: 26–36.

Paul Ormerod (2005), *Why Most Things Fail: Extinction, Evolution and Economics*, Faber and Faber, London

Paul Ormerod and Andrew P. Roach (2008), Emergent scale-free social networks in history: burning and the rise of English Protestantism, *Cultural Science* 1:1.

Paul Ormerod and Greg Wiltshire (2009). Binge drinking in the UK: a social network phenomenon. *Mind and Society*. DOI 10.1007/s11299-009-0058-1

George Orwell (1946). Politics and the English language. *The Penguin Essays of George Orwell*: 348-60. London: Penguin.

Jason Potts, Stuart Cunningham, John Hartley and Paul Ormerod, (2008). Social network markets: A new definition of the creative industries. *Journal of Cultural Economics* 32, 167-185.

Luke Rendell, Robert Boyd, D. Cownden, Marquist Enquist, Kimmo Eriksson Mark W. Fledman, L. Fogarty, Stephano Ghirlanda, T. Lillicrap and Kevin N. Laland (2010). Why copy others? Insights from the social learning strategies tournament. *Science* 328, 208-213

Jeremy Richardson (2000) Government, interest groups and policy change. *Political Studies*, 48:5, 1006–1025.

Matt Ridley (2010). *The Rational Optimist: How Prosperity Evolves*. London: Fourth Estate.

Matthew J. Salganik, Peter S. Dodds and Duncan J. Watts (2006). Experimental study of inequality and unpredictability in an artificial cultural market *Science* 311, 854-856.

Thomas C. Schelling (1973), Hockey helmets, concealed weapons, and daylight saving: a study of binary choices with externalities. *Journal of Conflict Resolution*, 13, 381-428.

Claude E. Shannon (1948). A mathematical theory of communication. *Bell System Technical Journal* 27, 379-423.

Claude E. Shannon (1951). Prediction and entropy of printed English. *Bell System Technical Journal* 30, 50-64.

Stephen J. Shennan and J. Richard Wilkinson (2001) Ceramic style change and neutral evolution: a case study from Neolithic Europe. *American Antiquity* 66, 577–594.

Herbert A. Simon (1955) On a class of skew distribution functions. *Biometrika* 42, 425-440.

Gabriel Tarde (1895) *Les Lois de l'Imitation*. Paris, Alcan.

Duncan J. Watts (2002) A simple model of global cascades on random networks. *Proceedings of the National Academy of Sciences* 99, 5766–5771.

Duncan J. Watts and Steven H. Strogatz (1998). Collective dynamics of 'small-world' networks. *Nature* 393: 440-2.

George U. Yule (1925) A mathematical theory of evolution, based on the conclusions of Dr. J. C. Willis, F.R.S. *Philosophical Transactions of the Royal Society of London Series B* 213, 21–87.