

## **Infectious Ideas: Modelling the Diffusion of Ideas across Social Networks**

Mandy Haggith<sup>a</sup>, Ravi Prabhu<sup>b</sup>, Carol J. Pierce Colfer<sup>c</sup>, Bill Ritchie<sup>a</sup>,  
Alan Thomson<sup>d</sup> and Happyson Mudavanhu<sup>b</sup>

<sup>a</sup> worldforests, 95 Achmelvich, Lochinver, Sutherland, Scotland IV27 4JB

<sup>b</sup> CIFOR Southern and Eastern Africa Office, 73 Harare Drive, Mount Pleasant  
Harare, Zimbabwe

<sup>c</sup> CIFOR, PO Box 6596 JKPWB, Jakarta 10065, Indonesia

<sup>d</sup> Pacific Forestry Centre, 506 West Burnside Road, Victoria BC, Canada V8Z 1M5

Will the practice of collecting wild honey wearing no clothes become a widespread practice in Zimbabwe? Or will beekeeping take over as the main way that people acquire honey? Both practices impact on forest resources; how can the foresters influence the uptake of these ideas? This paper describes an exploratory modelling study investigating how social network patterns affect the way ideas spread around communities. It concludes that increasing the density of social networks increases the spread of successful ideas whilst speeding the loss of ideas with no competitive advantage. Some different kinds of competitive advantage are explored in the context of forest management and rural extension.<sup>1</sup>

**Keywords:** memetics, social networks, agent-based modelling, competitive advantage

### **INTRODUCTION**

The villagers of Batanai, near Mafungautsi forest in Zimbabwe, are tired of the bee stings they suffer when they collect honey in the forest. Fortunately, a young man called Richard Nyirenda<sup>2</sup> has made an exciting discovery - the bees do not sting if you raid their nest wearing no clothes!<sup>3</sup> Meanwhile, Mrs Mafa has come up with a totally different solution. She has built a hive, acquired a swarm of domesticated bees, and is producing honey right in her backyard. Will naked honey forays into the forest catch on or will Mrs Mafa's beehives be the next big thing in Batanai? How do new ideas like these spread around communities?

---

<sup>1</sup> The authors are grateful to the UK's Department for International Development (DFID) and the European Community for financial support of this project.

<sup>2</sup> All names have been changed.

<sup>3</sup> This is intended to be light-hearted example, and is not seriously advocated. However, when disturbed, bees do react to human sweat, and no clothes may well be preferable to sweaty clothes.

Richard's naked hive-raiding scheme has a lot going for it - no capital inputs or no new technology are required, and it is just a minor modification of the traditional way of smoking the bees out of their tree-top nests. Mrs Mafa's idea requires a lot of effort and investment and it also involves learning a totally new technique, but she has instigated the formation of a beekeeping group to help others get started. Being the wife of the *sabhuku* or village headman, she has a lot of influence and many opportunities for persuading people of her new idea. Which will be more successful, the intrinsically smart idea or the idea of the effective persuader? How do Richard's and Mrs Mafa's social networks influence the spread of their ideas?

Frank Matose, who works for the Forestry Commission, became interested in this situation because both options may affect the forest. The fires that people light to smoke out the bees often end up destroying the trees, so he is sceptical about the traditional technique, naked or otherwise. Building new beehives also requires timber from the forest, but may be managed more easily than burning. However, Frank knows from past experience that the villagers do not trust the Forestry Commission, so open support for the beehive idea might not be an advantage. How can the idea that is benign to the forest be promoted?

Traditional theories of innovation diffusion (Rogers 1995) have assumed that people make rational choices about adopting new ideas based on their expected utility, with the proviso that some people are inherently open to new ideas ('early adopters'), whilst others are more conservative ('laggards'). Subsequent research, as reported by Aron (1989), revealed that availability of resources to support adoption is a more critical constraint on uptake of new technology than attitudes of potential adopters. Recent research, particularly in the field of memetics, suggests that patterns of cultural change can be explained without assuming rational decisions, but rather using imitation as the basic means of transmission of cultural features such as ideas. In other words, ideas are 'contagious' (Lynch 1996) and innovation diffusion is a process of 'thought contagion', that is, the transmission, by coping, of 'infectious ideas'. This paper demonstrates that a simple agent-based model of thought contagion by imitation is an effective way of exploring questions such as those posed above, and of generating insights into how new ideas spread around a community.

This paper first summarises some of the more interesting controversies about how new ideas spread or how innovations diffuse. This sets the context for describing this study of 'infectious ideas', in which an exploratory approach to modelling was used, cycling repeatedly through a series of steps to maximise the opportunities for generating insights. The story of this process forms the bulk of this paper. Such exploratory learning generated many questions along the way and uncovered some difficult issues, which are discussed, along with a summary of insights. The paper finishes by returning to Richard, Mrs Mafa, Frank and the bees and, with tongue firmly in cheek, proposes what strategies the modelling suggests they should adopt.

### **The Context for this Study**

This paper results from a study carried out for the Center for International Forestry Research (CIFOR), within its Adaptive Collaborative Management (ACM) research programme. The purpose of the study was to demonstrate to the research team how to go about using models as tools for exploring hypotheses about social dynamics. It was primarily a methodological case study, to show how models can be used within

learning cycles. The aim chosen for the study was to model some different configurations of social networks and their influence on the spread of ideas. This was intended to complement the ACM research programme's explorations in the field of how collaboration, by building on and expanding social network configurations, can enhance people's capacity to develop and share adaptive strategies of natural resource management.

This research also contributes to the FLORES modelling project (Vanclay 1998, 2003), which is developing computer-based simulation models of the intelligence of social forest landscapes (Haggith and Prabhu 2003).

### **Hypotheses Tested**

The hypotheses explored in this study were stated initially as follows:

1. Social learning and adaptation involves 'thought contagion' which varies with social networks and kinship structures. Some ideas will be primarily passed through family relationships. Others will be primarily shared over networks of peers, neighbours or other social relationships. (Note that for the purposes of this study, broadcast media such as radio and published sources of information were not considered.)
2. Collaboration changes the configuration of social networks, thus changing the channels for infection by new ideas, and so collaboration will alter the patterns of dissemination of knowledge within a community. In particular, collaboration that strengthens a social network will enhance the spread of successful ideas and limit the spread of unsuccessful ones.

The modelling explorations described in this paper have revealed ambiguities in these hypotheses, and have clarified some of the important concepts which will enable more specific hypotheses to be framed for future empirical testing. In particular, with respect to the second hypothesis, several different interpretations are possible for 'strengthening' a social network and there are several crucially different ways of defining what counts as 'success' of an idea.

### **Theories About Diffusion of Ideas**

Mainstream theories of innovation diffusion assume that change occurs as a result of people acting as rational decision-makers (Simon 1960). When faced with a new idea, the assumption is that people more or less carefully weigh up the utility of the idea based on a range of criteria such as complexity, compatibility with existing ideas, how public the adoption will be, how easy it is to try out, and the extent to which it is better than existing alternatives. Clark (1999) proposed a decision process involving five stages: knowledge or exposure to an idea and understanding its functions; persuasion or formation of a favourable attitude; decision or commitment to its adoption; implementation (putting it to use); and confirmation or reinforcement based on positive outcomes. Most innovation diffusion theory also assumes that people all fit into one of a set of personality types, which influences the time it takes for them to be persuaded to adopt a new idea. Typical is Roger's partitioning of populations into innovators, early adopters, early majority, late majority and laggards. AtKisson (1991) noted that people may display different personality traits for different kinds of ideas, so for example a real innovator with

respect to gadgets may be an extreme laggard about eating new kinds of food. A parallel study to the one described in this paper has explored this kind of model of innovation diffusion (Thomson *et al.* 2003).

The theory of memetics (Blackmore 1999) gives an entirely different explanation of cultural change as the spread of *memes* by imitation. The idea of a meme was first introduced by Richard Dawkins, as ‘a unit of cultural transmission’ (Dawkins 1976), such as a catch-phrase, a tune, a ritual, a taboo, a fashion such as wearing one’s hat backwards, a technique such as beekeeping or an idea such as ACM, or indeed, naked honey collection. Memes are the cultural equivalents of genes and they are evolving in a process of natural selection in which only the fittest memes survive. Fitness needs to be understood from the point of view of the meme, not the point of view of the person who is the ‘carrier’ of the meme. A fit, or successful, idea is one that succeeds in being copied widely and achieving widespread use. Other ways in which memes can be seen as ‘fit’ have been explored by Lynch (1996). Memetics aims to provide evolutionary accounts of all mental and cultural processes and phenomena, without being ‘greedily reductionist’ (Dennett 1995).

An example of an evolutionary explanation of cultural change is Axelrod’s agent-based model of cultural dissemination (Axelrod 1997), which, like a memetic explanation, uses imitation as the basic transmission mechanism for new ideas. This model provided the starting point for the construction of several simple models of diffusion of ideas, and stimulated much research into the complex and adaptive nature of social systems (Axelrod and Cohen 1999).

In the complex systems field, agent-based modelling (Epstein and Axtell 1996, Ferber 1998) is extensively used. It involves simulations of societies made up of simple agents or representations of the individuals in the population, with behaviour at the system level emerging from the complex interactions between the agents.

There are many other areas of theoretical relevance to this study, including theories of communication, social learning and information networks, but these are beyond the scope of this paper.

## THE TWO-LEVEL PROCESS OF MODELLING

The ACM programme uses action research as its methodology, which involves working in learning cycles or loops, of planning, action, monitoring and reflection. This study used a two-level process of modelling for learning, with an outer and inner loop, involving a broad range of interested ACM researchers in the non-technical aspects of the modelling work in the outer loop, and a smaller set of people interested in all the technical details in the inner loop.

The outer loop involved articulating a set of hypotheses about diffusion of ideas, circulation of these hypotheses to the ACM team by email, and responding to the resulting feedback from other team members. The results of the modelling work were also presented to the team, and to a broader audience at a seminar on modelling in natural resource management, for feedback followed by reflection and insights. This led to refinement of the lessons learned and exploration of several new avenues of research and related ideas, particularly social learning. It has also led to discussion of the role of modelling and greater understanding of its potential in the ACM’s primarily action-research methodology. This outer loop allowed members of

the ACM team who have had no modelling experience to gain some exposure to what modelling involves, and also for the modelling work to be informed by their thinking, without them needing to commit the time needed to engage with the inner loop of technical modelling work.

The inner loop involved reimplementing Axelrod's model of cultural dissemination (see below) and 17 iterations of the following steps:

- modification of the model;
- testing the behaviour of the modified model; and
- analysis and reflection on the results so far.

This inner loop involved collaboration primarily with other modellers. This part of the process took a total of 10 days of work, spread over two months.

### **Basic Premises Underlying the Model**

Axelrod's agent-based model of cultural diffusion (Axelrod 1997, pp. 145-177) revealed how cultural norms can arise at the system level through local-level imitation. It is based on the premise that 'the more similar an actor is to a neighbour, the more likely that actor will adopt one of the neighbour's traits.' The model is becoming a well-deserved classic, for its simplicity, its adaptive style, its surprising emergent behaviour and the fact that it demonstrates how cultural change can come about without the individual agents making rational choices.

The model is highly abstract. It involves 100 agents, each of whom has five *features* or dimensions, the values or *traits* of which range from 0 to 9. The culture of the agents is defined as the list of their five feature values. One can think of these features as characteristics such as religion, ethnicity, sports played, crops grown, staple food, or in the case of the example, honey collection method. The degree of *cultural similarity* of two agents is the percentage of their features with identical traits. The agents are spatially distributed on a 10x10 grid, and each agent is in a neighbourhood relationship with the adjacent agents (north, south, east and west). There is no movement between grid cells. Agents at the edge have only three neighbours and agents in corners have two.

The process of social influence in the model is as follows:

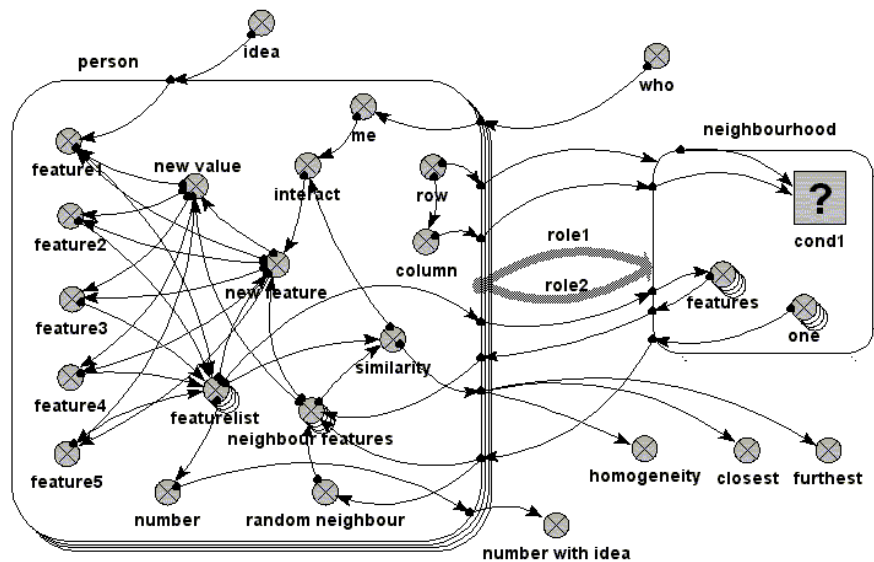
1. Pick an agent to be active at random.
2. Pick one of its neighbours at random.
3. These two agents interact with probability equal to their cultural similarity.
4. If they interact, pick at random a feature on which they differ.
5. If there is such a feature, the active agent copies the trait of her neighbour on that feature.

Therefore, if two agents have features as follows: [2,5,6,1,0] and [2,4,6,1,9] (i.e. three features are the same, two are different) then there is a 60% chance that one will copy the other, and the result will be that they will have only one feature different. By contrast, two agents with features [2,5,6,1,0] and [3,1,7,9,3] for example, have nothing in common (note the two instances of 1 represent different features, so do not count as the same) and therefore they will not interact with each

other. Similar interacting agents therefore become more similar, whilst agents who have nothing in common stay that way.

Under this model, Richard’s and Mrs Mafa’s chances of being imitated by other members of their community depends on shared experiences. If Richard’s favourite sport is football, he is Muslim, Ndebele-speaking, grows cotton and likes naked honey collection, then he needs to find people who are similar to some extent to him, who may imitate his great new idea. A football-playing, Shona-speaking Catholic who grows maize and keeps bees will be only 20% likely to copy him. (Note that this model makes no *a priori* assumptions about any particular features being more or less likely than others to generate imitation, so religious persuasion is treated as being no more or less culturally important than crop preference, for example. Like all modelling assumptions, this is a gross simplification of reality).

For the purposes of this study the model was re-implemented from Axelrod’s specification using Simile, a graphical simulation modelling environment which supports both system dynamics and object-oriented model development. The model diagram is shown in Figure 1.



**Figure 1.** Simile model diagram of the re-implementation of Axelrod’s model of cultural diffusion

On the left hand side of the diagram looking like the stack of cards are the agents, represented as a population sub-model called ‘person’ with 100 instances, each of which has their five feature variables in the column on the left. The neighbourhood relationships between the agents are represented on the right hand side of the model. At each time step an agent ‘who’ and a ‘random neighbour’ are selected, and if their cultural ‘similarity’ suffices they ‘interact’ which may result in a ‘new feature’ or ‘new value’ for the agent.

The model produces the behaviour predicted by Axelrod. Over runs lasting 100,000 time-steps the population forms a small number of discrete cultural groups, with complete homogeneity within the groups, and nothing in common between groups. At this point a stable state is reached. Often there is a single majority culture plus some minority groups of neighbours all alike but utterly unlike everyone else.

### Exploring Different Kinds of Social Relationships

In order to explore the first hypothesis, that the spread of ideas varies with social networks and kinship structures, it is necessary to model a range of patterns of social relationships, and then to explore how an introduced new idea fares in each. The new idea is added to the model as a new trait for one of the features. One can think of this as adding ‘go naked’ or ‘domesticate bees’ to the range of options for honey collection. This is implemented in the model depicted in Figure 1 by means of the variable ‘idea’ which assigns the new feature trait to one person at random on initialisation of the model, thus seeding the population with one person with the idea. Seven different social patterns were implemented by modifying the definition of the relationship on the right side of the model diagram:

1. *Neighbours*. A spatial arrangement where each person has physical neighbours in adjacent positions.
2. *Two friends*. A network of relationships such as friendship, in which each person has exactly two best friends, and the people whom you consider your best friends may have different best friends, i.e. the relationship may not be reciprocated.
3. *Ten friends*. A similar network of relationships in which each person has 10 friends.
4. *Hierarchy one-way*. Hierarchy such as kinship, in which the pattern of relations forms a family tree, in which there is definite direction in the flow of ideas (ideas flow down the tree from parents to children, i.e. from senior to junior) and there can be no circularity (one person cannot be both senior to and junior to another).
5. *Hierarchy*. Similar kinship hierarchies but in which ideas can flow from parents to children, from children to parents, and between siblings and spouses.
6. *Two large clubs*. Two closed subgroups such as clubs or churches with 50 people in each.
7. *Ten small clubs*. Ten closed subgroups with 10 people in each.

A variant of the Axelrod model was produced for each alternative using exactly the same imitation algorithm as Axelrod’s model except that instead of agents just copying their neighbours, they copy their friend, or kin, or co-member of the club, depending on the model variant. Note that the restriction still applies that copying is dependent on the similarity of cultural features, hence two people in the same club or two friends, will still be more likely to copy each other the more alike they are.

Tests were run on each model variant, one person being picked at random and ‘seeded’ with the new idea at the start of each model run. A test constituted 10 runs of 100,000 time-steps each. At the end of each run, records were made of the diffusion of the idea (i.e. the maximum number of people who ever had the idea) and

the duration of the idea (i.e. the latest time at which anyone had the idea). A summary of the results is provided in Table 1.

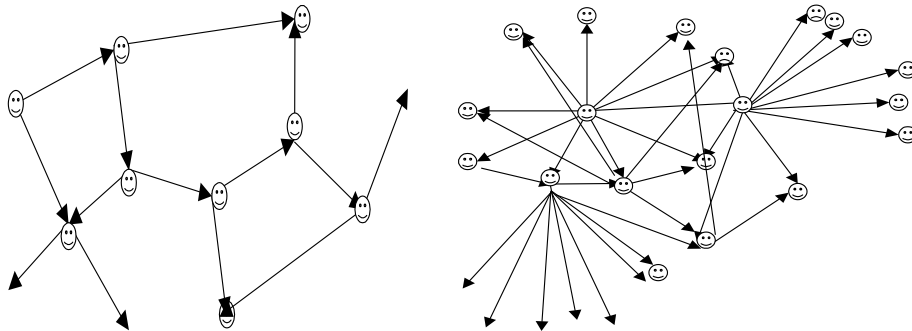
**Table 1.** Summary of average values from model tests with different social configurations

Model	Diffusion	Duration
Neighbours	6.3	25,181
2 friends	4.4	33,590
10 friends	2.2	6,185
Hierarchy one-way	3.9	33,495
Hierarchy	1.2	12,028
2 big clubs	2.4	5,949
10 small clubs	1.2	40,718

Note: Diffusion is the maximum number of people who ever get the idea. Duration is the maximum number of timesteps that the idea lasts. The seven model variants reflect the seven distinct social network patterns described above.

**Reflections on Networks**

At this point it is necessary to introduce the idea of *density* of the social network. The network of friends in which everyone has 10 friends is *more dense* than the network with only two best friends, which is a *loose* network. These are illustrated in Figure 2.



**Figure 2.** A loose social network, with at most two relationships from any one person to other people (left); and a dense social network, with up to 10 relationships from any one person to other people (right)

From Table 1, it is intriguing that increasing the density of the network of friends results in a decrease in the duration of the new idea. The idea lasts longest in the case when it is hidden away in a church of 10 people who do not interact at all with anyone else. The idea lasts second longest in the case when everyone only interacts with two people. Its durability is least when people interact with 10 other people. In addition, in the model families appear to be dreadful environments for the spreading of new ideas, and they do better if they follow a strict hierarchy with only parent-to-child transmission. This is because of the large number of people who may interact



with in large families, which is similar to having many friends. The broadest diffusion of the new idea was found to occur where interactions are limited to a few neighbours or only two friends.

The result appears to support the first hypothesis, that the spread of ideas varies with social networks and kinship structures, because different kinds of social structure are captured, in their bare bones, by these model alternatives, and they do affect the spread of ideas in different ways.

In contrast, results from the second hypothesis are more difficult to interpret. If ‘strengthening social networks’ by collaboration implies an increase in the density of the network (i.e. getting more people to interact with each other), and if ‘successful’ ideas are those that are durable, then this is a counter-intuitive result. A new idea will be less successful in a stronger social network. Collaboration that strengthens social networks thus reduces the chances of new adaptations surviving. The conclusion, deeply perplexing and controversial, is that collaboration impedes adaptation!

To understand these results, it is important to realise that when an individual interacts with other people their own ideas are subject to competition from many competing ideas. The imitation algorithm picks a trait to copy at random, so retaining one’s own original ideas is easier if there are fewer interactions with other people whose ideas are different. This can also be interpreted as a kind of ‘institutional inertia’ resulting from large groups converging on an institutional culture, which is then resistant to change and can overwhelm new ideas. Strong social networks can be powerful forces for conservatism.

These results raise important conceptual questions about the extent to which powerful groups can affect the spread of ideas, and also the way that ideas about different topics may be more subject to competition than others, and how these factors may vary in time and space. The model is of course too simple to take these complexities into account. However, there is a compelling need to explore in more depth this issue of how collaboration affects competition between ideas. Mrs Mafa needs to know whether it will be best for her idea if she keeps her beekeeping group exclusive or opens it up to a broad membership. Richard needs to know if it is worth joining the football club to widen his social circle and try to spread his idea, or if this will merely make him more likely to conform to the social norm.

A striking thing about these results is that overall, the transmission of the new idea is dismal. The greatest uptake is where the new idea is transmitted to six people out of 100. Perhaps this is not surprising because the idea has no competitive advantage over all the other ideas encountered, and is at a disadvantage because it begins with only one instance in the population, as opposed to all the other values in the feature space, which begin with an average of 10% coverage.

### **What is a Successful Idea?**

The next set of tests explored various ways of giving an idea a competitive edge over the other traits which could be imitated. Two simple-minded approaches were tried initially. First, the idea was made ‘unloseable’, so that once a person had copied the new idea, that person would never copy a different value. Not surprisingly, the idea rapidly dominates, but this is a rather uninteresting and extreme situation. Secondly, the disadvantage of comparative rarity with other traits was removed, by seeding the models with multiple copies of the new idea, to test

whether this broadens the transmission of the idea. Results indicate that an idea that begins with 10% coverage achieves, on average, about 10% diffusion, whereas an idea that begins with 40% coverage achieves, on average, about 40% diffusion.

A more interesting question is how an idea might achieve better-than-average diffusion by having an edge in competition over other ideas, of a form other than a privileged starting position or an unrealistic inherent dominance like the 'unloseable' idea. This question raises a key ambiguity in the second hypothesis of this study, concerning the notion of a 'successful' idea. This could have at least three different interpretations.

1. An idea might be successful if, as an emergent property of the system, it has achieved a high level of transmission. In this case the success of an idea is found in the extent to which it is effective at spreading itself.
2. An idea might be successful if it produces local benefits, such as health, wealth or happiness, for the person who has it. In this case the success of the idea is the extent to which its host is successful, according to a local measure.
3. An idea might be successful if it produces benefits to the system as a whole, for example, benefiting a resource such as a forest or an asset held in common by the community. In this case the success of the idea is the extent to which the system is successful, according to a global measure.

These three different interpretations of what indicates a successful idea require different approaches to modelling. In the case of the first interpretation, the new idea is given greater chance of success by making it stand out as the best idea to copy, when two people interact. This is dubbed 'copy-the-idea'. If Mrs Mafa adopts this strategy, she doesn't go about trying to talk to more people than usual, but whenever she interacts with anyone, she will tell them about beekeeping. She becomes an 'evangelist' or 'proselytist'.

In the case of the second interpretation, the idea is made successful by making its host stand out as the best person to interact with, when choosing someone from the set of friends, neighbours or relatives. Which of the host's ideas is copied remains random, so the idea's advantage here is only indirect. This is dubbed 'copy-the-carrier'. If Richard adopts this strategy he will try to ensure that he stands out from the crowd and has as many interactions with people as possible. When he talks to people he doesn't talk about naked honey collection more or less than anything else, but he hopes that the increased interest in him will lead to an increased imitation of his ideas in general and therefore naked honey collection in particular. He becomes a 'trend-setter'.

The third interpretation of 'successful' is more confusing. It is not clear how to make an idea more likely to be copied just because it leads people to act in ways which provide shared benefits. Even assuming that people will like an idea that benefits the commons, there is still the question of not knowing who or what to copy. If 100 people are interacting with a resource, how can one tell which of those people are producing a benefit and which not, and even if one could do that, how does one tell which of their ideas leads to the action producing the benefit? This problem is profoundly important for the question of promoting ideas that favour sustainability of forest resources, but it is merely raised here, not answered. The problem is summarised by the question: 'What can be done to increase the spread of

ideas that sustains the forest?’ This is the problem faced by Frank, from the Forestry Commission. Unfortunately, the model contains no representation of the forest or of a common resource, so it may not be suitable for exploring this question.

One possibility is that the answer will involve cultural identity, such that people may copy not just a particular individual, but rather those traits that define an emergent culture. This is related to the idea of community norms, which evolve through emergent processes of identification by individuals with their forest or other natural resources. It would be interesting to explore this idea, perhaps by developing feedback loops, within a multi-agent model such as this one, from the emergent culture back to the individuals within it.

The first two versions of successful ideas are explored through the models ‘copy-the-idea’ and ‘copy-the-carrier’. In each instance, the underlying social configuration is the network of friends, as this can be altered easily to produce the extremes of behaviour shown above, by changing the density of the network. In addition, the initial number of instances of the new idea seeded in the population was varied. Therefore eight model variants were tested. Each test involved 10 runs of the model, of 10,000 timesteps each. The results are reported in Table 2.

**Table 2.** Summary of average diffusion and duration observed in tests of two alternative strategies, copy-the-idea, and copy-the-carrier

Density	Seeds	<i>Copy the idea</i>		<i>Copy the carrier</i>	
		Diffusion	Duration	Diffusion	Duration
10	1	28.2	6,853	49.2	8,078
10	10	89.0	10,000	67.6	10,000
2	1	1.9	5,957	8.5	5,284
2	10	44.9	10,000	44.4	10,000

### Reflections on Transmission of Ideas

The most successful individual strategy is copy-the-idea, with high density and 10 seeds. In other words, in a densely networked community, it is best to seed the community with many instances of an idea that is itself copied in preference to other ideas. In memetics terminology this copy-the-idea strategy represents an idea that involves ‘proselytism’, in other words, the idea is like a chain-letter which contains instructions to pass it on, or like a religious creed that makes a virtue out of evangelism.

In a high density community, a proselytic idea with only one initial proponent (i.e. copy-the-idea with only one seed), will spread more slowly than an idea that promotes ‘copy-the-carrier’, which with only one seed can still be expected to diffuse to 49% of the community. Hence Mrs Mafa needs to be warned that if she is the only beekeeper around, her strategy is not likely to be successful and it is really important for her to make sure that her initial ‘seed’ group of beekeepers is a substantial proportion of the village.

The average of all four situations reveals that ‘copy-the-carrier’ is a more effective strategy than ‘copy-the-idea’. On the whole, ideas that improve the reputation of their hosts and make them stand out from the crowd appear to spread more widely than proselytic ideas. It therefore seems Richard’s strategy is likely to be effective. Note, however, that the difference between the two approaches is not

statistically significant, and doubts can be cast on the meaningfulness of taking an average across the four case studies. It would be interesting to develop a model of a community in which both types of transmission compete with each other in order to see which is the more successful.<sup>4</sup>

Note that the tests also confirm the second hypothesis, in that a more dense ('stronger') social network leads to more rapid diffusion and longer duration of both kinds of successful ideas, whilst speeding the elimination of ideas with no competitive advantage.

## LESSONS LEARNED AND QUESTIONS REMAINING

This section addresses some questions raised by the study and draws some conclusions.

### Summary of Conceptual Lessons

The behaviour of the model is such that different configurations of social networks appear to have different impacts on the diffusion and durability of new ideas. Social networks based on hierarchies (e.g. the kinship model) seem to be less effective for spreading ideas than networks of non-reciprocal relationships (e.g. the friendship model), so it seems that vertical (senior-to-junior) transmission is less effective than horizontal (peer-to-peer) transmission. This seems to echo findings (e.g., Gladwell 2002) about how new ideas can reach 'epidemic' proportions through peer-to-peer spread. Social networks involving closed subgroups appear to be less effective in achieving diffusion but more effective in achieving durability of ideas than all other configurations.

New ideas, with no inherent advantages, seem to be easily drowned out by competition. The greater the density of the social network, the greater is the competition, and thus the more rapidly are new ideas drowned out. If all ideas are equal, a new idea may have less likelihood of diffusion in a highly interconnected social network, because better ideas, or better established ideas, will dominate. If collaboration increases these social interconnections, then it can be expected to reduce the life-expectancy of ideas with no inherent advantages.

There appear to be various ways of giving a new idea a competitive advantage:

1. One way to get a competitive advantage is to be 'unloseable'. Unloseable ideas spread extremely well. (Although examples from the real world of unloseable ideas are difficult to find it has been suggested that one might be a language, which once understood is difficult to stop understanding.)
2. The more instances of a new idea seeded into a community, the better the expected survival of the idea. (This is why advertisers use so many instances of their images.)
3. If an idea encourages its own imitation, or the imitation of its carrier, then it is more likely to diffuse and be durable. (This is why evangelism works and why advertisers use popular personalities to project images.)

---

<sup>4</sup> Thanks to Jerry Vanclay for this idea.

4. Dense social networks appear to reinforce the spread of such successful ideas more effectively than less dense networks. The better ideas can exploit the increased competition by wiping out inferior ideas more rapidly. (Bearing in mind the result above that peer-to-peer imitation is most effective, and the fact that young people's social networks are usually dense due to membership of institutions such as school, this explains why advertisers try to sell to young people.)
5. Encouraging the imitation of the carrier of a new idea is a slightly more effective strategy for diffusion than encouraging the imitation of the idea itself. (There may be implications here for agricultural extension, for example, in terms of choice of people in a community with which to work. In particular, this result questions the appropriateness of using 'innovators' or 'evangelists' as community gatekeepers. A better criterion may be how popular someone seems to be, or how much effort other community members make to interact with them.)

### **Unanswered Questions Arising from this Study**

This modelling study has revealed a number of questions that the model cannot help to answer in its present form, and it seems only fair to emphasise these limits of the model.

First, the model uses a highly idealised and simple form of imitation as its core method of transmitting ideas. The choice of person to copy is based only on cultural similarity (except in copy-the-carrier) and the choice of feature is random (except in copy-the-idea). Copying is always exact, and complete, and involves no decision-making. There is a growing body of evidence (e.g. see Blakemore 1999), that people do use automatic copying behaviour to a huge extent. However, one can still wonder whether people sometimes choose what to imitate, and if so, how do they choose? What kind of a judgement is made, and what values are used? Some critics of memetics particularly dislike its lack of use of rational decision-making, and it may be that they can be satisfied by an explanation involving rational decisions about what ideas or what parts of ideas to adopt, even if the adoption is imitative.

Another key question to do with the imitation issue is how does the level of influence of some people affect whether they are copied, or whether they can prevent copying by others? It would be interesting to explore the role of power dynamics, and the potential of empowerment processes, in facilitating the spread of ideas.

It is also important to recognise the limitations of representation of the various social network patterns in the model. In the real world, all of the various kinds of relationships operate simultaneously, but in the model they have been explored only one at a time. It would be interesting to explore the behaviour of a model with all of the relations from the various models superimposed, and even to model competition between the copy-the-idea and copy-the-carrier strategies. In addition, the relationships and feature values are randomly applied to the population at model initialisation, and it would be instructive to attempt to initialise the models to reflect a real community such as Batanai. Such empirical testing of the model would certainly be possible and this would require concrete data on the configurations of real social networks, such as those collected in studies such as Standa-Gunda and Prabhu (2001). It would be particularly fascinating to explore empirically how actual

collaborations reconfigure social networks and compare this with model behaviour to elucidate further how collaboration changes the social fabric. This needs to take into account the fact that collaboration may occur without the creation of new relationships, by changing the nature or intensity of existing social networks.

Another simplification of the model is that there is only a single new idea, and it arrives by fiat at the start of the model. It is also necessary to enquire what are the processes by which new ideas are generated and where do they arise in a population. How do different idea-origination processes influence their spread? Do different sources of ideas (e.g. local versus outsider) have inherent advantages? Historical studies into the source of innovations and their patterns of spread are needed, and it is possible that the records of extension agencies and development organisations might reveal answers to some of these questions. An empirical study is needed into communities which have made the shift from wild honey collection to beekeeping, particularly exploring the path of adoption of current beekeeping technology.

Finally, the most profound question remains. How can an idea that is beneficial for everyone (in this case, good for the forest), and not just for the individual, gain a competitive advantage sufficient to succeed? To address this question with a model will require representing in the model the emergent culture of the population as a whole and addition of a feedback loop from this emergent culture back to the individuals within it. This would enable exploration of models in which the agents imitate the traits of an emerging 'cultural identity' that supports and is supported by the health of a shared resource such as the forest.

### **Adoption Theories and the Batanai Situation**

With tongue in cheek, let's return to Batanai village and the questions posed at the beginning of the paper. The model suggests advice for Mrs Mafa. If her only strategy is to use her existing social networks to evangelise about beekeeping then she may be fighting a losing battle against all the other ideas around the village. However, establishing the beekeeping group was a good move as it has seeded the idea through a good proportion of the community and that will give the idea more sticking power.

It seems there is good news for Richard. As long as he maintains his strategy of attempting to stand out from the crowd, being seen as a trend-setter, and making the most of peer-based social networks such as the football team, it is likely that naked honey collection could take off in Batanai. He needs to ensure that he has something in common with as many people as possible so that they have some reason to interact with him. If this works, then many people will want to copy him, and naked honey collection thing is bound to attract quite a few followers, who will in turn gain in notoriety. It could spread like wildfire.

Speaking of which, that's just what Frank's worrying about. And unfortunately this paper must leave him, and you, pondering how an idea that's good for the forest can become as infectious as naked honey collection.

## **REFERENCES**

- Arnon, I. (1989), *Agricultural Research and Technological Transfer*, Elsevier Applied Science, New York.

- AtKisson, A. (1991), 'The innovation diffusion game: A tool for encouraging participation in positive cultural change – or for doing something different at your next party', *In Context*, 28: 58, <http://www.context.org/ICLIB/IC28/AtKisson.htm>, accessed 3 February 2003.
- Axelrod, R. (1997), *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration*, Princeton University Press, Princeton.
- Axelrod, R. and Cohen, M. D. (1999), *Harnessing Complexity: Organizational Implications of a Scientific Frontier*, The Free Press, New York.
- Blackmore, S. (1999), *The Meme Machine*, Oxford University Press, Oxford.
- Clarke, R. (1999), 'A primer in diffusion of innovations theory', Xamax Consultancy Pty Ltd, Canberra, <http://www.anu.edu.au/people/Roger.Clarke/SOS/InnDiff.html>, accessed 5 February 2003.
- Dawkins, R. (1976), *The Selfish Gene*. Oxford University Press, Oxford.
- Dennett, D. (1995), *Darwin's Dangerous Idea*, Penguin Books, London.
- Epstein, J.M. and Axtell, R. (1996), *Growing Artificial Societies - Social Science from the Bottom Up*, MIT Press, Cambridge.
- Ferber, J., (1998), *Multi-Agent Systems*, Addison-Wesley, Reading.
- Gladwell, M. (2000), *The Tipping Point: How Little Things Can Make a Big Difference*, Abacus, London.
- Haggith, M. and Prabhu, R. (2003), 'Unlocking complexity: The importance of idealisation in simulation modelling', *Small-scale Forest Economics, Management and Policy*, 2(2): 293-312.
- Lynch, A. (1996), *Thought Contagion: How Belief Spreads Through Society*, Basic Books, New York.
- Rogers, E.M. (1995), *Diffusion of Innovations*, The Free Press, New York.
- Simon, H.A. (1960), *The New Science of Management Decisions*, Harper & Row, New York.
- Standa-Gunda, W. and Prabhu, R. (2001), 'Spatial lags hypothesis: Discerning conceptual frameworks for communication networks', Unpublished paper, CIFOR, Harare.
- Thomson, A., Haggith, M. and Prabhu, R. (2003), 'Innovation diffusion: Modelling the rate of adoption of new technologies and policies', *Journal of Technological forecasting and social change*, in process.
- Vanclay, J.K. (1998), 'FLORES: For exploring land-use options in forested landscapes', *Agroforestry Forum*, 9(1): 47-52.
- Vanclay, J.K. (2003), 'Why model landscapes at the level of households and fields?', *Small-scale Forest Economics, Management and Policy*, 2(2): 121-134.