

'Binge' drinking in the UK: a social network phenomenon

Paul Ormerod¹ and Greg Wiltshire²

July 2008

1. Volterra Consulting UK and Institute of Advanced Study, University of Durham; corresponding author; pormerod@volterra.co.uk
2. Volterra Consulting UK gwiltshire@volterra.co.uk

We are grateful for comments from anonymous referees for the European Social Simulation Association conference, Brescia, Italy, September 2008

Abstract

In this paper, we analyse the recent rapid growth of 'binge' drinking in the UK. This means the rapid consumption of large amounts of alcohol, especially by young people, leading to serious anti-social and criminal behaviour in urban centres.

British soccer fans have often exhibited this kind of behaviour abroad, but it has become widespread amongst young people within Britain itself. Vomiting, collapsing in the street, shouting and chanting loudly, intimidating passers-by and fighting are now regular night-time features of many British towns and cities. A particularly disturbing aspect is the huge rise in drunken and anti-social behaviour amongst young females.

Increasingly, policy makers in the West are concerned about how not just to regulate but to alter social behaviour. Smoking and obesity are obvious examples, and in the UK 'binge' drinking has become a focus of acute policy concern.

We show how a simple agent based model approach, combined with a limited amount of easily acquired information, can provide useful insights for policy makers in the context of behavioural regulation.

We show that the hypothesis that the rise in binge drinking is a fashion-related phenomenon, with imitative behaviour spreading across social networks, is sufficient to account for the empirically observed patterns of binge drinking behaviour.

The results show that a small world network, rather than a scale-free or random one, offers the best description of the data.

1. Introduction

In this paper, we analyse the recent growth of ‘binge’ drinking in the UK. By this, we mean the rapid consumption of large amounts of alcohol, especially by young people, leading to anti-social behaviour in urban centres. British soccer fans have often exhibited this kind of behaviour abroad, but it has become widespread amongst young people within Britain itself. Vomiting, collapsing in the street, shouting and chanting loudly, intimidating passers-by and fighting are now regular night-time features of many British towns and cities. A particularly disturbing aspect is the huge rise in drunken and anti-social behaviour amongst young females.

The phenomenon is of serious concern to the British government, not merely for the anti-social behaviour related to it, but because of the longer term health implications for young people of massive intakes of alcohol in very short periods of time.

Ref [1] suggests that the experience of the UK is part of a more general pattern of heavy, rapid drinking which is emerging in a number of countries of the world, and uses the phrase ‘extreme drinking’ rather than that of ‘binge’ drinking to characterise it.

There is a growing literature which demonstrates the importance of social networks for consumer choice in what might be termed ‘regular’ consumer markets. A popular reference, for example, on this is [2]. The concept of the ‘tipping point’ is used to explain on why some books, films and music emerge out of obscurity with small marketing budgets to become popular hits when many *a priori* indistinguishable efforts fail to rise above the noise. A much more formal analysis of the importance of social networks in determining success or failure in the film industry is [3].

In many social and economic contexts, individuals are faced with a choice between two alternative actions, and their decision depends, at least in part, on the actions of other individuals. Ref [4] describes this class of problem as one of ‘binary decisions with externalities’. An important feature of such systems is that they are ‘robust yet fragile’ [5,6]. In other words, behaviour may remain stable for long periods of time and then suddenly exhibit a cascade in which behaviour changes on a large scale across the individuals within the system.

Two recent American studies [7,8] using the Framingham Heart Study data base [9] have demonstrated the importance of social networks in determining the behaviour of individuals on matters of public health, specifically obesity and smoking. The Framingham data base contains detailed information on over 12,000 individuals, monitored over more than three decades since 1971.

The social networks of individuals on this data base have been important determinants of both the spread of obesity and the reduction in smoking over this period. In terms of obesity, for example, the chance of any individual being obese increased by 57 per cent if he or she had a friend who became obese. When a spouse stopped smoking, the other was 67 per cent less likely to smoke.

The aim of this paper is to examine the extent to which the sudden emergence of the binge drinking problem in the UK can be explained as a social network phenomenon.

We use the methodology developed in [10]. We consider a small amount of straightforward and readily accessible information. An agent based model is set up in which agents face the binary decision on whether or not to binge drink. Transmission of binge drinking behaviour across agents connected on a social network is determined according to a threshold rule.

The theoretical model is calibrated against empirical evidence. We deduce from this, using an agent-based model, the type of social network across which information flows and agents influence each other's behaviour in this context. Specifically, we show that information appears to flow across a small world network.

The approach described here can be used more generally in areas where policy makers are interested in regulating and altering agent behaviour. An important aspect of the methodology is that it is feasible both to construct empirically grounded agent based models and to draw useful implications from them, whilst at the same time requiring only small amounts of data.

Section 2 describes the basic data, Section 3 the initial evidence for the existence of imitative behaviour on social networks, Section 4 the theoretical model and results. Section 5 examines the robustness of the results, and Section 6 gives a brief discussion.

2. The data

In this particular context, no longitudinal survey such as the Framingham Heart Study exists. So data was gathered using standard survey techniques. The market research company FDS interviewed 504 18-24 year-olds in the UK using an online survey based on MyVoice Panel. Of the respondents, 258 (51 per cent) were male and 246 (49 per cent) were female. The sample group was selected to reflect a demographic which is believed to represent a particular problem in terms of alcohol consumption.

Definitions of heavy drinking vary widely [11] and changes to the standard definitions can have a significant impact on the reported incidence of alcohol misuse. For example, the latest Office for National Statistics report on alcohol consumption in the UK [12] introduced a revised methodology for estimating the proportions of heavy drinkers within the population, taking into account increased alcohol strengths and larger drink sizes. This results in increased counts of heavy drinkers in all age and gender categories, even though the underlying data have not changed. For people aged 16-24, for example, the proportion of women identified as heavy drinkers rises from 29 per cent to 40 per cent.

The focus of this study is not on heavy drinking as such, but on drinking behaviour which is likely to lead to anti-social behaviour i.e. binge drinking.

An individual might regularly drink a fairly large quantity of alcohol but (being habituated) might not subjectively experience this as ‘bingeing’, i.e. might not actually feel that they are particularly drunk. Thus, in order to distinguish between regular *binge* drinkers and those who are simply regular *big* drinkers, our definition is based upon a combination of consumption of alcohol (anyone drinking more than 10 drinks in a single session is considered to be drinking enough to get very drunk, regardless of their own perception), and subjective perception – those who at least once a week drink an amount that they had previously specified as being, for them, “enough to get very drunk”.

We have therefore defined ‘binge drinking’ as follows:

For men, getting drunk on 4 or more drinks OR having 10 or more drinks (but not necessarily getting drunk) at least once a week and for women, getting drunk on 3 or more drinks OR having 10 or more drinks (but not necessarily getting drunk) at least once a week.

This definition therefore captures behaviour that is directed at purposefully getting drunk, and also includes those who drink excessively (i.e. ten or more drinks in a single session) even if the excessive drinking does not cause the drinker to feel drunk. It is similar to the definition of very heavy drinking/intoxication used in [13].

Overall, 16.2 per cent of respondents qualified as binge drinkers using the definition described above. Of this group, the vast majority reported anti-social behaviour as a result of binge drinking such as shouting or vomiting in the street, getting into a fight.

Scaling up the survey, the figures indicate there are around 950,000 binge drinkers in the UK 18-24 year old population, participating in 1.5 million binge drinking ‘events’ each week.

3. Initial evidence for interaction on social networks as a factor in binge drinking

We analysed the patterns of social interaction for those classified as binge drinkers and compared them to the patterns of non-binge drinkers. We looked at three types of social group which might have an influence on a person’s drinking behaviour:

- Family
- Work colleagues
- Friends

Everyone in the survey was asked what they thought about the binge drinking behaviour of people in their social groups. Table 1 shows the results for family members.

Table 1: Proportion of family thought to be binge drinkers: for binge drinkers and non-binge drinkers

Proportion of family thought to be binge drinkers	Proportion (%) for binge drinkers	Proportion (%) for non-binge drinkers
All of them	9	3
Almost all of them	9	3
Most of them	11	10
Some of them	33	28
Hardly any of them	29	36
None of them	10	19

So, for example, amongst people who binge drink themselves, 18 per cent think that ‘all’ or ‘almost all’ their family members also binge drink. This compares to non-binge drinkers, 6 per cent of which think ‘all’ or ‘almost all’ their family members binge drink. There may of course be a genetic component to familial resemblance in drinking behaviour. But, fortunately, we do not need to get involved in this difficult issue since there are considerably more marked differences between the social networks of binge and non-binge drinkers when we consider work colleagues and friends.

Table 2: Proportion of colleagues thought to be binge drinkers: for binge drinkers and non-binge drinkers

Proportion of colleagues thought to be binge drinkers	Proportion (%) for binge drinkers	Proportion (%) for non-binge drinkers
All of them	13	2
Almost all of them	21	10
Most of them	31	22
Some of them	24	42
Hardly any of them	7	18
None of them	6	6

Here, for example, no less than 34 per cent of binge drinkers think that ‘all’ or ‘almost all’ of their work colleagues binge drink, compared to only 12 per cent of non-binge drinkers.

But by far the most dramatic difference is seen in the behaviour of friends¹.

¹ This is confirmed in formal analysis by calculating both the Manhattan and Euclidean norms between the two columns

Table 3: *Proportion of friends thought to be binge drinkers: for binge drinkers and non-binge drinkers*

Proportion of friends thought to be binge drinkers	Proportion (%) for binge drinkers	Proportion (%) for non-binge drinkers
All of them	24	6
Almost all of them	30	11
Most of them	31	24
Some of them	12	36
Hardly any of them	1	15
None of them	2	7

Table 3 shows that 54 per cent of binge drinkers think that all or almost all of their friends are binge drinkers, compared to 17 per cent of non-binge drinkers for whom all or almost all friends are binge drinkers. Conversely, only 3 per cent of binge drinkers have no or hardly any friends that binge drink, compared to 22 per cent of non-binge drinkers.

4. The theoretical model and its calibration

Our aim is to establish whether social network effects are a *sufficient* condition to account for the observed binge drinking behaviour in the UK. We know from Tables 1-3 above, especially Table 3 which is now our specific focus, that binge drinkers have different sets of social networks to non-binge drinkers.

We set up a simple agent based model, in which the decision of an agent to become a binge drinker is determined solely by the proportion of friends on his/her network who already are binge drinkers. The paper follows the methodology described in [10], where the issue analysed was whether people on benefits had bank accounts.

In the model, we connect agents on different types of social network, specifically a random, a small world and a scale-free network, using both replacement and additional re-wiring in the latter.

Initially, all agents in the model are in state 0, i.e. they are not binge-drinkers. A small percentage (2 per cent) of the total is selected at random to become binge drinkers (state 1).

Each agent is allocated a threshold above which he or she will convert from state 0 to state 1. This is drawn at random from a uniform distribution on $[0, U_1]$, where U_1 is a variable of the model. The threshold is the proportion of friends who in state 1, above which the agent will switch from state 0 to state 1, otherwise stay in state 0.

We monitor the percolation of state 1 behaviour across the network, and halt the solution when the proportion reaches 16.2 per cent, the estimated number of state 1 agents from the empirical data.

We then examine the networks of the friends of agents in both state 0 and state 1, and to see how closely they correspond to the observed structure set out in Table 3 above. More precisely, we simplify Table 2 slightly, and merge the categories ‘all’ and ‘almost all’ into a single one, and do the same for ‘none’ and ‘hardly any’.

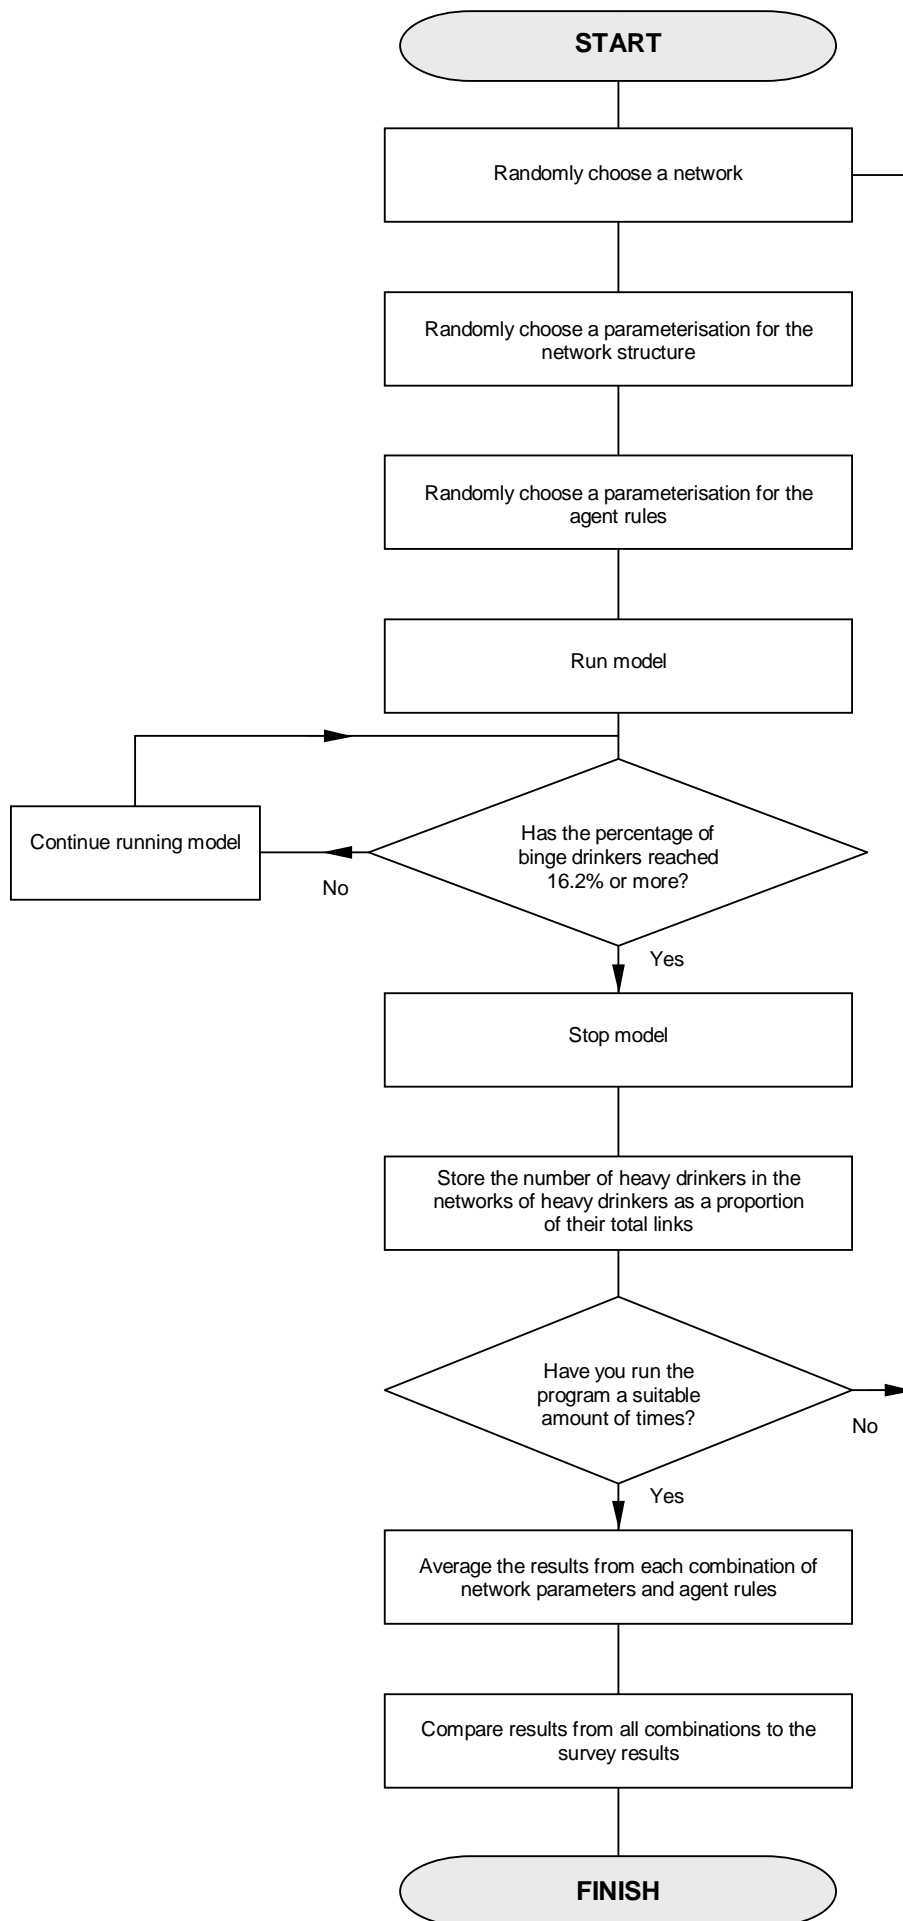
A final assumption needs to be made as to what the categories ‘all/almost all’, ‘most’, ‘some’ and ‘hardly any/none’ mean in percentages. We use the following:

Table 4: Assigned values for the questionnaire responses.

Questionnaire Response	Assigned Corresponding Value	Value For Quartile Denoted As
‘Hardly any’ and ‘None’	≥ 0 and $\leq 25\%$	Q_1
‘Some’	$>25\%$ and $\leq 50\%$	Q_2
‘Most’	$>50\%$ and $\leq 75\%$	Q_3
‘All’ and ‘Almost all’	$>75\%$ and $\leq 100\%$	Q_4

We conducted extensive searches for the best combination(s) of relevant parameters in each of the three types of networks examined.

A graphical representation of the calibration approach adopted is set out in Figure 1.



The options are small world (replacement rewiring), small world (additional rewiring), scale free or random

See table 5

Agents will change from non-binge drinkers to binge drinkers if the proportion of binge drinkers in their network is above their individual threshold.

This threshold is randomly drawn from a uniform threshold between $L1$ and $U1$ and is constant throughout the run of a model.

The model is run for up to 20 time steps. If the run fails to reach 16.2% it is discarded.

In each time step each agent is able to assess their network against their threshold and take up binge drinking.

The program must be run a suitably large amount of times so that each possible combination of network parameters and agent rules has been generated and tested. Furthermore each combination must be run a significantly large number of times so that the average behaviour can be determined as not all runs of the same parameterisation will be equivalent.

In the initial sweep the program was run 40,000 times. Candidate final model parameterisations were run 1000 times each.

Figure 1 *Flow chart of the methodology use to parameterize the model*

The initial sweeping of the combinations of model parameters was performed 40,000 times which equated to averaging each parameterisation over 300-1000 runs. The candidate models taken forward from this sifting were then run an additional 1000 times. The range of parameters examined is set out in Table 5.

Table 5: Parameters used in the generation of the three types of networks

Application	Parameter	Description	Value/Range
General parameters	n	Number of agents in network.	1000
	L1	The lower limit to the distribution for the threshold of agents to switch from not binge drinking to binge drinking based on an evaluation of agents connected to them by their social network.	0
	U1	The upper limit to the distribution for the threshold of agents to switch from not binge drinking to binge drinking based on an evaluation of agents connected to them by their social network.	0.4-0.8
	L2	The lower limit to the distribution for the threshold of agents to switch from binge drinking to not binge drinking based on an evaluation of agents connected to them by their social network.	1.2
	U2	The upper limit to the distribution for the threshold of agents to switch from binge drinking to not binge drinking based on an evaluation of agents connected to them by their social network.	1.2
Small world network	k	Number of adjacent agents each agent is linked to on either side.	2-10
	ϕ	Probability of rewiring a link (either additionally or replacement) when generating network.	0-0.1
Scale free network	q	Average number of links each agent makes when it is added to network.	0.5-2
	α	Number of initially completely connected agents before generating network.	2-8
Random network	p	Probability that any two agents are connected.	0.002-0.025

In order to compare the results for the various models they were scored using the following equation.

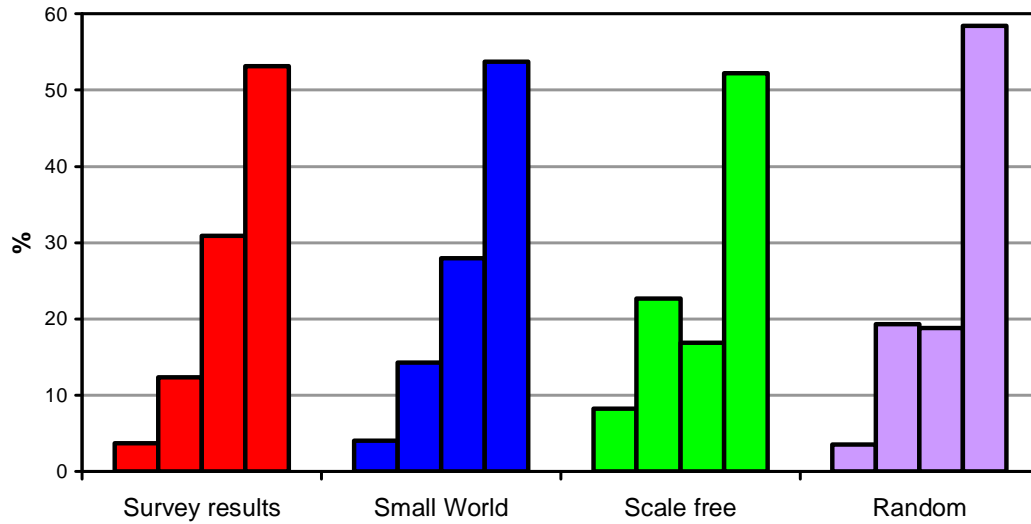
$$S = |q_1 - Q_1| + |q_2 - Q_2| + |q_3 - Q_3| + |q_4 - Q_4|$$

where q_i is the percentage in the relevant quartile of the basic data on friends described in Table 3 above, and Q_i is the model-generated proportion when overall 16.2 per cent of agents are in state 1.

Models with a lower S , or score value, will more closely resemble the survey results.

The models from each type of network with the lowest score are shown in Figure 2². The corresponding parameters are shown in Table 6. More detailed results across the various combinations of parameters are available on request from the authors.

Figure 2: *The final candidate models for each type of network with the lowest score value and the questionnaire results for the proportion of the binge drinker's friends who are thought to be binge drinkers*



It can be seen that the candidate model for the small world most accurately reproduces the survey results and is therefore the chosen model. It significantly out performs the other types of networks including the random network which was used as a control. The candidate network models do not reproduce the profile of the quartiles.

² Note that only one type of small world network is shown, this is for the version with additional wiring. The level of rewiring is low in the small world ($\phi \ll 1$) so results for both types of small world networks are the almost identical.

Table 6: Parameters for the final candidate models

Network	Parameter	Optimised Value
Small world network	k	4
	ϕ	0.005
	U1	0.5
Scale free network	q	1
	α	4
	U1	0.8
Random network	p	0.002
	U1	0.8

The optimized value of k implies that in the context of drinking behaviour, binge drinkers regard 8 people as their friends.

6 Robustness of the results

6.1 Effect of including a 1-0 transition

The model rules explored so far have only considered social influence causing agents to take up binge drinking (the 0-1 transition), the possibility that social pressure could stimulate people to give up binge drinking (1-0 transitions) has not been explored.

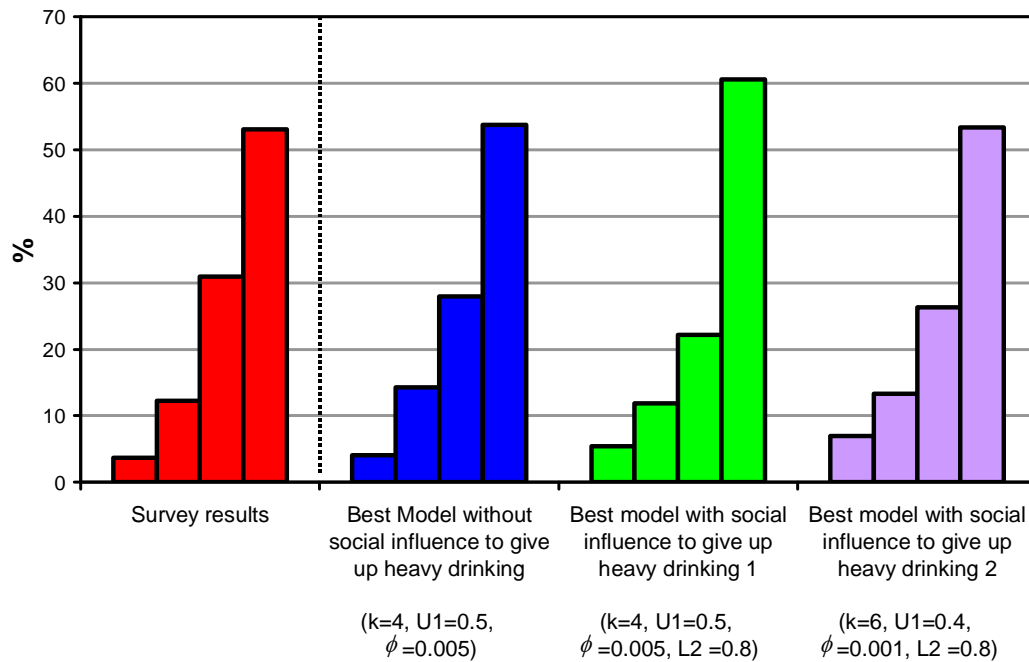
In order to identify candidate models which had similar or better scores to the chosen model and that included 0-1 transitions the small world model space described in Table 6 was swept again with L2 values of 0.8, 0.6 and 0.5 and a value of U2 of 1.2 added to the combinations. This means that a proportion of agents will never be able to stop binge drinking if they take it up, irrespective of their network, but that the remaining fraction will give binge drinking if their network is sufficiently connected to non-binge drinkers.

The candidates for models including 1-0 transitions are shown in Figure 3 alongside the survey results and the chosen small world model. Two candidate models are shown, firstly a model with the same k , ϕ and U1 values as the chosen small world model and secondly the optimum model from the entire parameter space.

Figure 3 shows that the chosen model, without behavioural rules to give up binge drinking, outperforms the candidates of those that do. The fact that in both of the optimised candidate models the scale of social pressure to give up binge drinking is much lower than that to take it up provides more evidence of the robustness of the chosen model

and that social influence to give up binge drinking (in this age group) can be approximated to zero.

Figure 3: Results from introducing behavioural rules which allow agents to stop binge drinking based on their social network. Shown in red is the questionnaire results and blue is the optimised model. Green shows the model with the best score when U2 was introduced into the optimised model while the purple results show the best model from all available parameterisations.



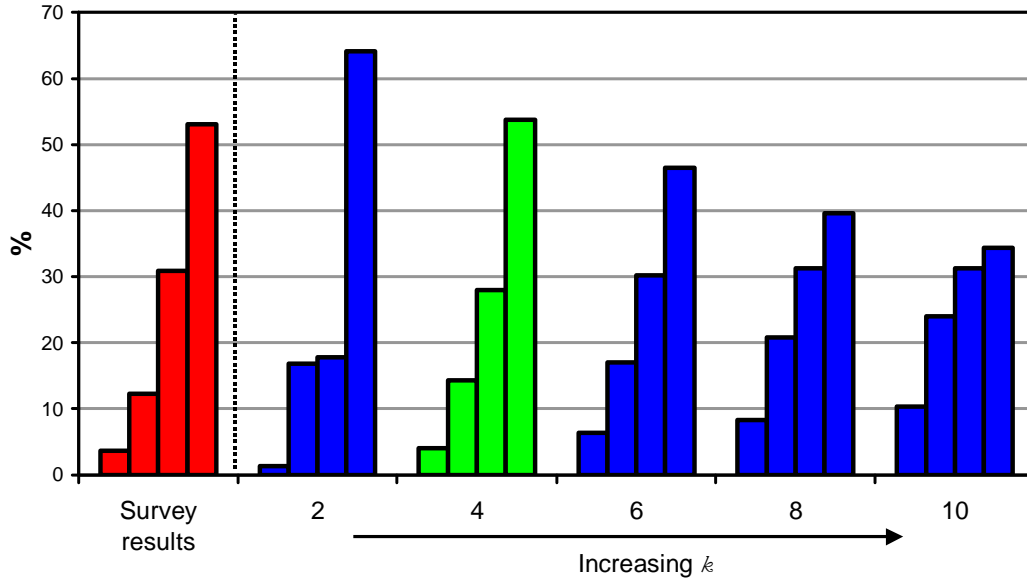
This result is intuitively plausible, given that the focus of the research is young people agents 18-24. The transition from 1 to 0 i.e. giving up binge drinking is likely to be closely connected to age, so as these particular agents get older, they will cease to be binge drinkers for a variety of reasons.

Effect of varying parameters from the optimized values

We demonstrate this with respect to the network parameters, k and ϕ and, for illustration. We also show results for one of the ranges of the general parameters $U1$, and demonstrate that the number of runs used, 1000, is suitably large to reveal the properties of the model in an unbiased way.

Figure 4 shows the effect of keeping all of the model parameters the same except for k . It shows that $k=4$ outperforms perturbations values between $k=2$ and $k=10$. The model exhibits the strongest sensitivity to this parameter. The results for $k=2$ and $k=6$ fall within the top 37th percentile of models with the former exhibiting a different profile shape to the survey results.

Figure 4: Results for varying k for the optimised model, shown in red are the questionnaire results and green is the optimised model



The reason for this increased sensitivity is demonstrated in table 7, a change in k means that the average number of binge drinkers required in the network of a non-binge drinking agent in order for them to switch behaviour varies. When k is reduced the average number of required binge drinkers falls as well. Similarly when k increases the average links required increases. In a model where agents make decisions on only integers (i.e. five of seven connections are binge drinkers) and where the increments of k that can be tested are restricted, changes in k represent strong transitions in the model which explain the changes in profile observed in Figure 4.

Table 7: Analysis of the impact of perturbations of k from the optimised model.

k	\bar{U}_1	Average links to agents who are binge drinkers required to take up binge drinking
2	0.25	0.5
4	0.25	1.0
6	0.25	1.5

Figure 5 shows similar results as figure 4 but ϕ is varied rather than k . The model still performs well for $\phi=0.001$, five times less than the optimised value, as well as $\phi=0.01$. These models scored within the top 3rd percentile of the small world models considered.

Figure 5: Results for varying ϕ for the optimised model, shown in red are the questionnaire results and green is the optimised model.

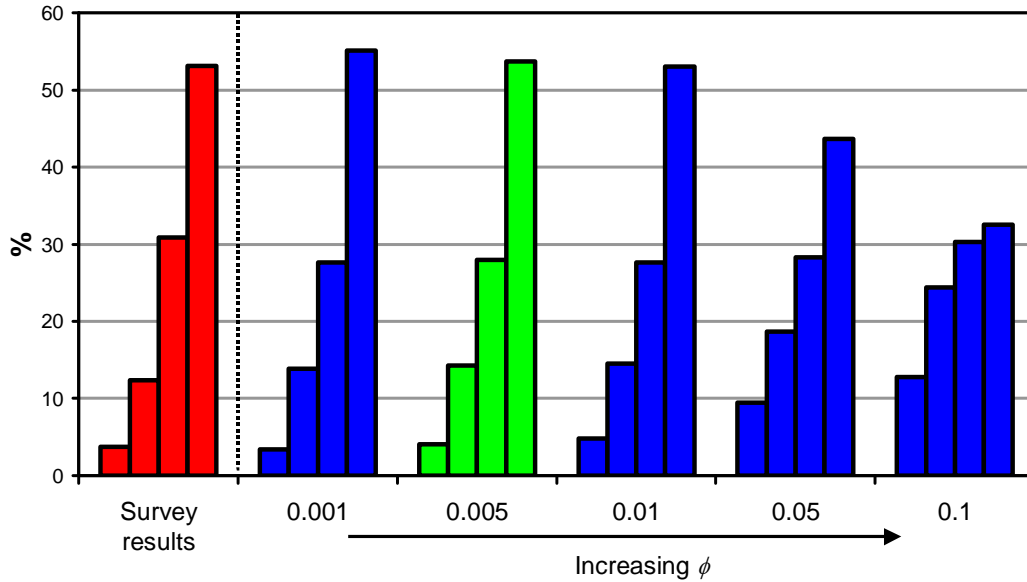
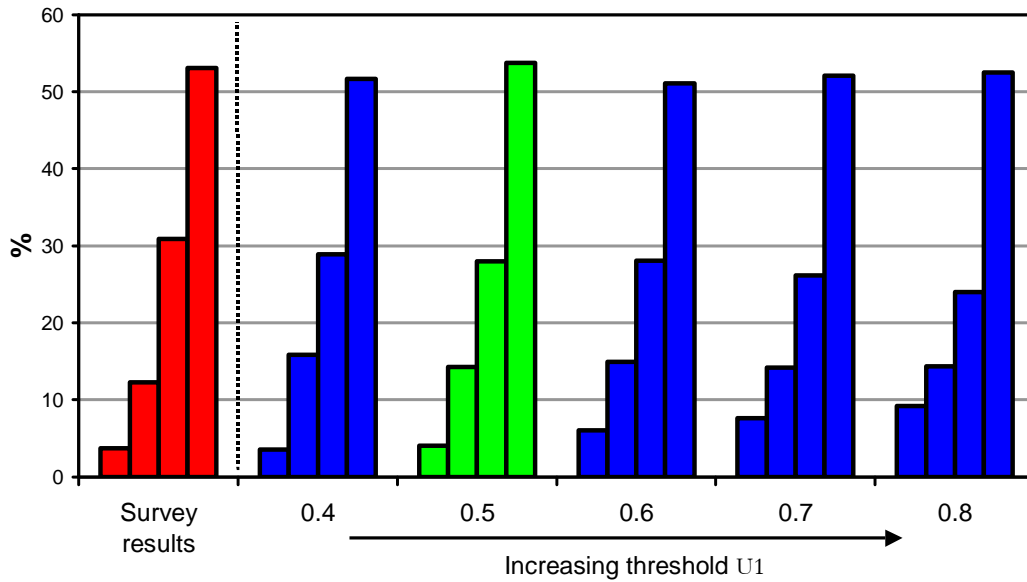


Figure 6 shows the final test for model robustness where $U1$ was varied while k and ϕ were held constant. It shows that $U1=0.5$ out performs perturbations values between $U1=0.4$ and $U1=0.8$. Importantly the model still performs well for both $U1=0.4$ and $U1=0.6$, the broad shape is similar. These two models were both within the top 5th percentile of the small world models considered.

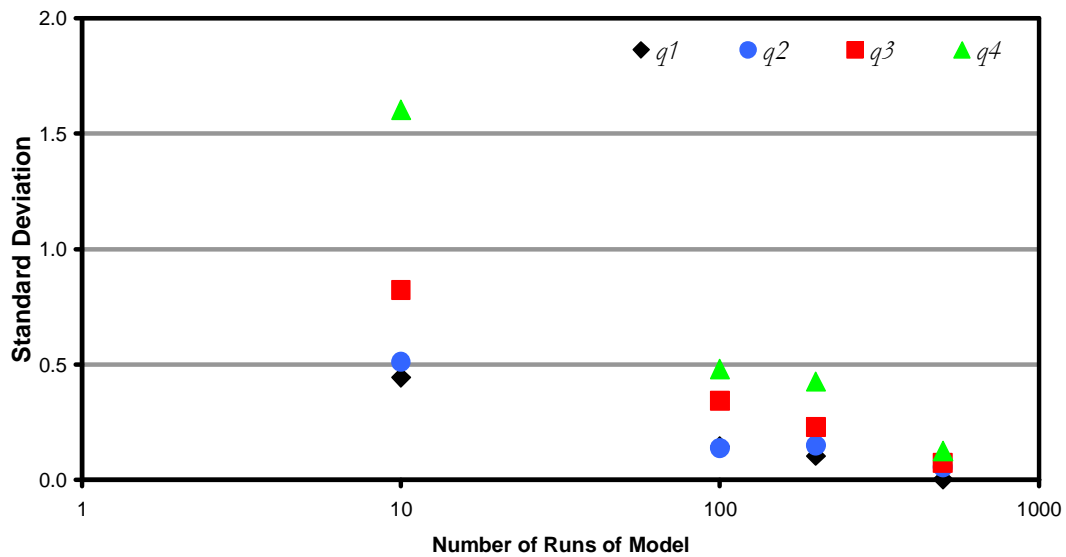
Figure 6: Results for varying U_1 for the optimised model, shown in red are the questionnaire results and green is the optimised model



As shown in Figure 1 the initial sweeping of the combinations of model parameters was performed 40,000 times which equated to averaging each parameterisation over 300-1000 runs. The candidate models taken forward from this sifting were then run an additional 1000 times. It is important to determine how the results of the models depend on this averaging and how much the results, from any given number of runs, can be expected to vary.

Figure 7 shows the standard deviations of the averages of set of runs of the model using different numbers of runs. For example the model was run 10 times and the average result calculated, this was repeated 16 times. The standard deviations of these 16 averages are plotted in Figure 7 for q_1 through to q_4 . It can be seen that the standard deviation drops dramatically as the number of runs in the averaging increases but that it plateaus off for 200 or more runs. Therefore the 1000 runs used in evaluating the candidate models is suitably large that the results can be considered statistically unbiased.

Figure 7: Standard deviations of each quartile as a function of the number of runs of the model



6.3 Distribution of quartile values for the chosen model

It is important that the results for q_1 through to q_4 of the chosen model are suitably distributed across the 1000 runs. If they were uniformly distributed then the results would hold little value. It is important that each result is suitably distributed around its mean. Figures 8-11 show histograms of the values of q_1 through to q_4 from 1000 runs respectively. Each result is suitably distributed, however it is noticeable that q_4 has the largest full width half maximum (FWHM) and is therefore the most variable result from the model

Figure 8: *Histogram of the value of q_1 from 1000 runs of the chosen model*

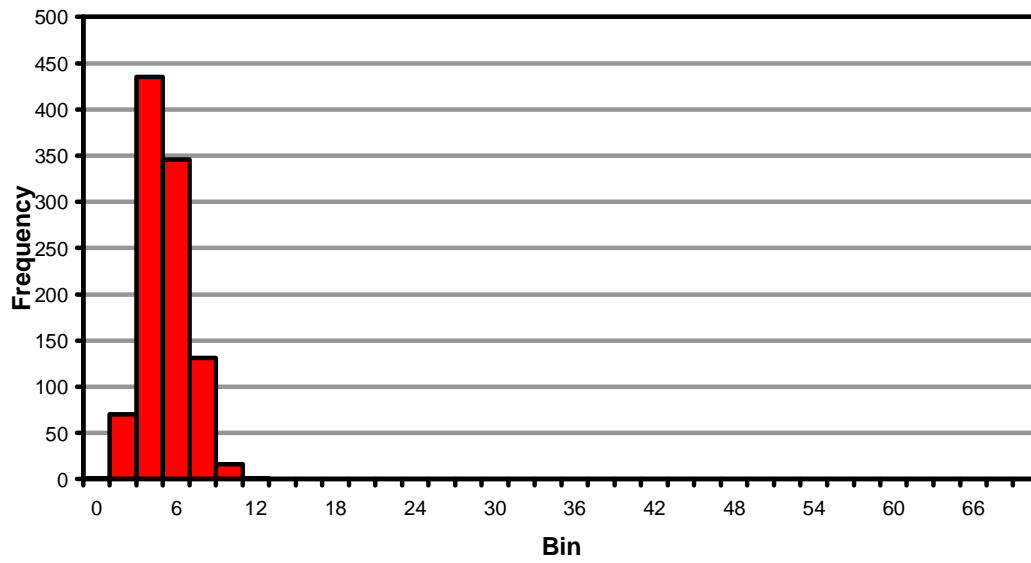


Figure 9: *Histogram of the value of q_2 from 1000 runs of the chosen model*

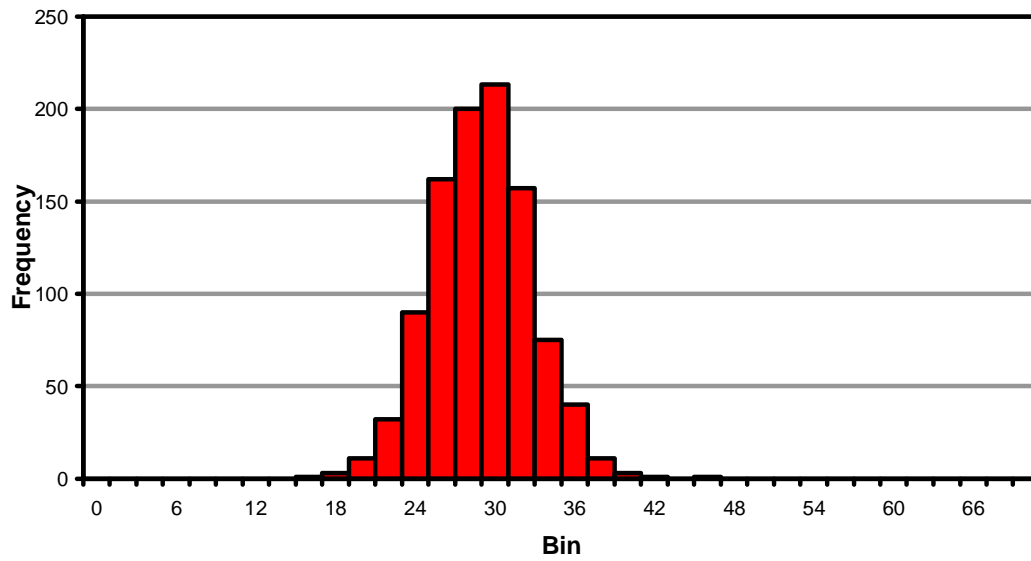


Figure 10: *Histogram of the value of q_3 from 1000 runs of the chosen model*

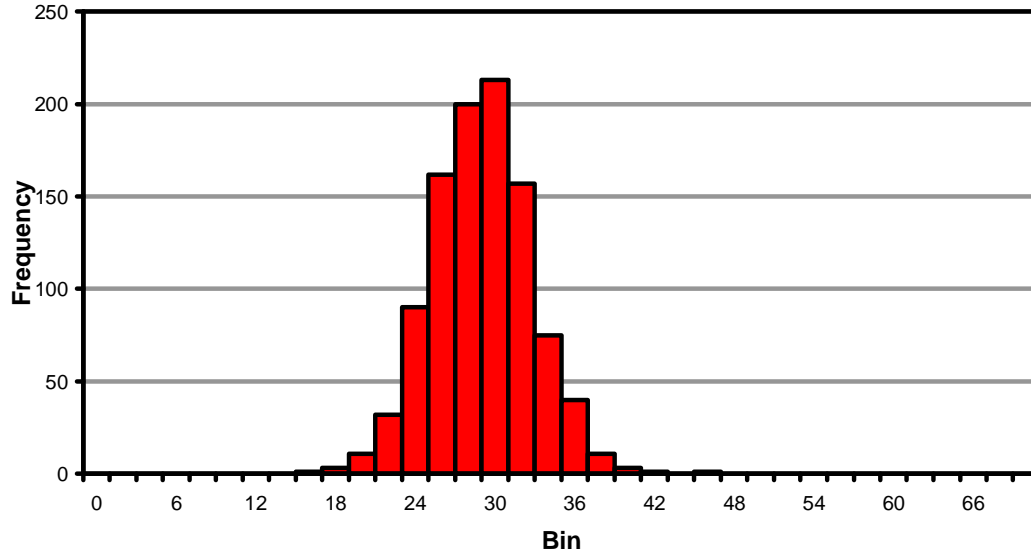
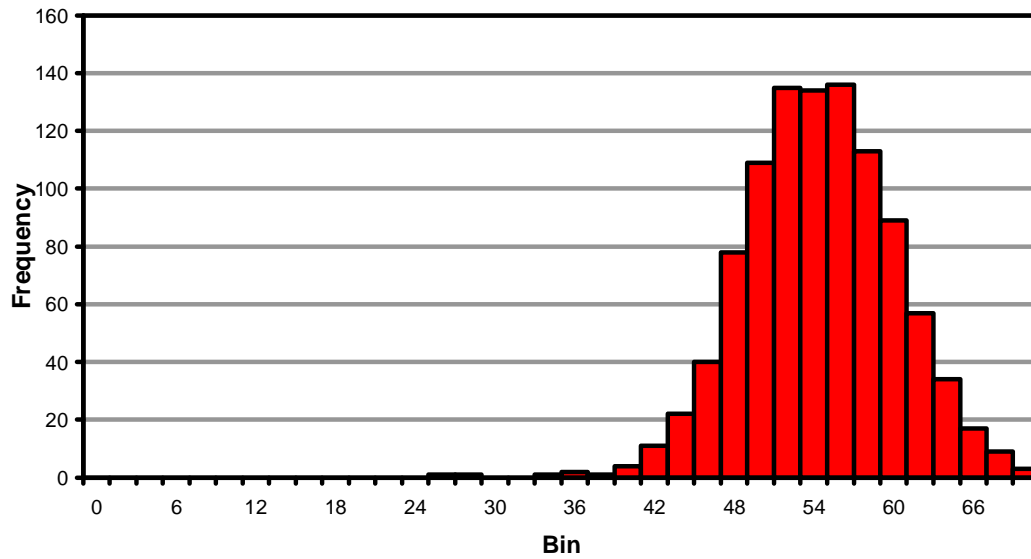


Figure 11: *Histogram of the value of q_4 from 1000 runs of the chosen model*



6. Brief discussion

Binge drinking in the UK has grown rapidly in recent years and has become a matter of serious policy concern. ‘Binge’ means the rapid consumption of large amounts of alcohol, especially by young people, leading to anti-social behaviour in urban centres.

Increasingly, policy makers in the West are concerned about how not just to regulate but to alter social behaviour. Smoking and obesity are obvious examples, and in the UK ‘binge’ drinking has become a focus of acute policy concern.

We develop a simple agent based theoretical model which requires a limited amount of easily acquired information in order to calibrate scientifically. The research is an example of the feasibility of building empirically grounded agent models with a small amount of data, which are able to offer useful insights into social phenomena.

We examine whether the spread of imitative behaviour across friendship networks is a *sufficient* condition to account for the observed patterns of binge drinking behaviour in the UK.

A standard market research survey was carried out in order to discover both the number of binge drinkers in the 18-24 year old population, where the problem is most acute, and their friendship patterns in terms of drinking behaviour. There are decisive differences in the drinking behaviour of friends of binge drinkers compared to the drinking behaviour of non-binger drinkers.

We examined different types of potential networks, random, scale-free and small world with both additional and re-wired links. We conducted extensive searches for the best combination(s) of relevant parameters in each of the three types of networks considered.

A small world network was the optimal choice of network, and was able to generate a close approximation to the observed patterns of behaviour.

There are two reservations which must be stated about the approach we have adopted. First, we lack any longitudinal data, so we are not able to analyse how the behaviour of any given individual evolves over time. This is a feature which is common to many issues which capture the attention of policy makers. Second, we have no direct evidence on the behaviour of the binge drinkers’ friends, but rely on the perceptions which the drinkers themselves have about their friends. There is therefore a risk that the respondents exhibit a certain amount of cognitive dissonance about their own behaviour in order to rationalise it and to protect their self-image. That said, the results do have a basic plausibility, since drinking, after all, is in general a social activity.

Both these potential problems could be overcome in straightforward enough ways, but at the cost of a huge increase in the level of time and resources required to obtain the basic data. The approach shows that it is possible to obtain interesting and useful results with what are essentially minimal data requirements.

It could also be argued that we are simply assuming the existence of a ‘contagion’ effect amongst friends rather than demonstrating its existence, particularly given that the basic data is cross-sectional rather than longitudinal. However, the phenomenon of binge drinking, whilst it has existed at a low level for many years in the UK, has grown very rapidly over a short period of time. This is shown by evidence such as very large increases in admissions to hospital of young people for binge drinking-related problems and police records on arrests of young people who have been binge drinking.

But it would be curious, to say the least, if large numbers of young people had suddenly decided quite independently of each other to binge drink, and then had happened to congregate together in friendship networks. So whilst the existence of a contagion effect amongst friendship networks is not technically proved by our results, it seems a far more likely explanation than the alternative one posited in the previous sentence.

The research does not demonstrate that imitation on social networks is necessarily the only significant reason for the recent rapid and dramatic rise in binge drinking. But it offers strong evidence that this factor is important, indeed it is sufficient to describe current behaviour. So policy makers have to take this into account when they try to devise strategies to combat the problem.

The discovery that the relevant network has a small world structure is also helpful to policy makers. It does not tell them precisely what to do, but it suggests, for example, that strategies based upon the concept that there is a small number of ‘influentials’ who are important in the spread of this anti-social behaviour are not likely to be very successful.

If the network had been a scale-free one, then of course such an approach might well work very well, provided always that the ‘influentials’ can be identified. This finding provides empirical support for the theoretical proposition developed in [14] that that it is rarely the case that highly influential individuals are responsible for bringing about shifts in public opinion and/or behaviour.

References

1. M Martinic and F Measham, *Swimming with Crocodiles: the Culture of Extreme Drinking*, Routledge, London, 2008
2. M. Gladwell, *The Tipping Point: How Little Things Can Make a Big Difference*, Little, Brown, New York, 2000
3. AS de Vany and WD Walls, ‘Bose-Einstein dynamics and adaptive contracting in the motion picture industry’, *Economic Journal*, 106, 1493-1514, 1996
4. TC Schelling, ‘Hockey helmets, concealed weapons and daylight saving: a study of binary choices with externalities’, *Journal of Conflict Resolution*, 17, 381-428, 1973
5. DJ Watts, ‘A simple model of global cascades on random networks’, *Proc. National Academy of Science*, 99, 5776-5771, 2002

6. P Ormerod and R Colbaugh, 'Cascades of failure and extinction in evolving complex systems', *Journal of Artificial Societies and Social Simulation*, 9(4), 2006
7. NA Christakis and JH Fowler, 'The spread of obesity in a large social network over 32 years', *New England Journal of Medicine*, 357, 370-379, 2007
8. NA Christakis and JH Fowler, 'The collective dynamics of smoking in a large social network', *New England Journal of Medicine*, 358, 2249-2258, 2008
9. <http://www.nhlbi.nih.gov/about/framingham/>
10. P Ormerod, 'Extracting deep knowledge from limited information', *Physica A*, 378, 48-52, 2007
11. *Practical Guides for Alcohol Policy and Prevention Approaches*, International Centre for Alcohol Policies, 2005
12. *Drinking: adults' behaviour and knowledge in 2007*, Office of National Statistics, London, 2008
13. G Knupfer, 'The risks of drunkenness (*or, ebrietas eesurrecta*): A comparison of frequent intoxication indices and of population sub-groups as to problem risks', *Addiction*, 79, 185-196, 2006
14. DJ Watts and PS Dodds, 'Influentials, networks and public opinion formation', *Journal of Consumer Research*, 34, 441-458, 2007

Acknowledgements

We are grateful to the UK Advertising Association for sponsoring this research. The authors have full responsibility for the content.

