How long can species co-exist in chaotic flows?

Theodore Jessop

***Abstract*—We have studied a modified version of the discrete time voter model with neutral selection wherein the interacting constituents are moved according to a chaotic flow, corresponding to a dynamic network. We have developed computer simulations of these systems in one and two dimensions, demonstrating that the limits of maximal and minimal flow correspond to the analytically soluble Moran process and voter process on a square lattice. We demonstrate that the consensus time of the system is maximised for intermediate rates of flow. We have studied evolutionary game theory with a view to extending this model to more complex games in the second semester.**

I. INTRODUCTION

A. Motivation

Evolutionary dynamics is applied in too many fields to enumerate. The voter model, explained in detail below in section II.A., is used to study systems of competition where connections between the competing actors is of paramount importance. Examples include ecological competition where species compete for resources, but only do so directly with their nearest neighbours. An example from the social sciences is the study of the diffusion of an idea or opinion through the population where individuals only communicate with their friends, family and colleagues.

Both these systems might be well represented by the classic voter model. However a clear limitation of this model is that it does not take into account the fact that in practice these networks are not likely to be static on all timescales relevant to evolution. Animals will move around to pursue resources and develop new neighbours as a result. People's social connections will change over time as they acquire new friends and colleagues and lose contact with old ones. If this changing of the network happens on a timescale much larger than that over which relevant dynamics in the evolution of the voters takes place then a static graph may be appropriate. For example, if one species of animal eliminates all of its competitors long before the individuals have had time to move around enough to change their mutual graph then a static voter model would capture all the needed detail. On the other hand this movement of individuals may be much faster than the evolutionary competition such that each individual is exposed to and competes with every other individual. In this case the geometry of the graph is irrelevant and the dynamics reduces to the Moran process wherein every node on the graph is connected to every other node (detailed below in section II.). These are what we refer to as the 'static' and 'well mixed' limits.

However there is no reason to suppose that most systems behave in this way where competitive evolution and flow of the network happen on drastically different timescales. Thus neither the static graph voter model nor Moran process are sufficient to capture the relevant dynamics of such systems, and what is needed is a model which interpolates between the two such that the flow of the network and evolution of the voters may be of a similar timescale. Here we develop such a model, in which the connections of the voter model are defined by Euclidean distance between voters distributed spatially and in which the positions of the voters move over time according to a flow rule. Chaotic flow is studied as it will have the property that at the limit of large flow, the position of a voter after an evolutionary time step will be unpredictable.

Supervisor: Dr. Tobias Galla, Project partner: Cory Emmerson

B. Structure of this report

This report will be presented in three main sections: firstly an overview of the theory behind the system we have studied, secondly details of the simulations we have produced this semester along with their results, and finally a discussion of how we propose to extend this work into the latter semester.

The first section will explain the principles behind our model, including the analytically soluble limits of the Moran process and fixed voter model. The second section details our simulations and their results, which include numerical verification of these analytic limits and the principal result that the consensus time of the system is maximised when the flow and evolution timescales balance. The final section will be a brief overview of the elements of evolutionary game theory that we propose to incorporate into the model to provide asymmetric selection dynamics.

II. THEORY

A. Discrete time voter model

The voter model is a stochastic system which consists of a population of constituents, variously called voters, particles, species, etc, which exist on a the vertices of a graph defines whether each voter is connected. These voters can be of one of two types, representing some differing factor such as political opinion, spin, or biological species. They can be generically labelled 0 and 1. At each step of time, a single voter is selected at random from among all voters. Then a voter connected to the initially chosen voter is also selected. The two then compete according to some predetermined rule, and the loser adopts the type of the winner. Thus the two objects which determine the dynamics of this system are the graph which determines which voters are connected and the selection rule that determines how they compete.

In the most simple dynamics, this selection rule is simply that the second chosen voter always loses and adopts the types of the first. This constitutes neutral selection where the probability of a voter winning an exchange is constant and does not depend on its type or situation. More complicated rules can constitute asymmetric selection and will be discussed in section III.

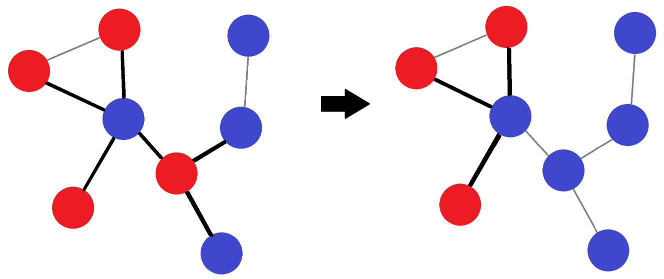

Fig. 1. Illustration of a single evolution time step in the voter model. Red circles are type 1 voters, blue are type 0 and the lines between voters represent connections, defining the 'network' or 'graph'. Light grey lines are 'passive' links between two voters of the same type, heavy black lines are 'active' links between voters of different types. If the voters across an active link are chosen for competition then a reaction can occur where one changes type, whereas should a passive be chosen nothing happens. Here two voters of different types are selected for competition. The type 0 wins the exchange and so the losing type 1 becomes a type 0. Note how the reaction reduces the total number of active links. These reactions continue until every voter is of one type or the other. This may be thought of as the reactions continuing to decrease the proportion of active links until it reaches zero in the consensus state.

Suppose there are N voters, i of which are of type 1 and $(N - i)$ of which are of type 0. During a time step a reaction between two voters takes place. There are three possible outcomes:

1) Two voters of the same type compete, in which case nothing happens as they both remain of the same type.
2) A type 1 meets a type 0 and the type 1 wins the exchange, replacing the type 0 with a 1, increasing i by 1.
3) A type 1 meets a type 0 and the type 0 wins the exchange, replacing the type 1 with a 0, decreasing i by 1.

An illustration of such an exchange is given in figure 1

Each time step will either change i by 1 in either direction or do nothing. Of paticular importance are the states with $i = 0$ and $i = N$, which correspond to every voter being of type 0 and 1 respectively. If the system is in one of these states then nothing more can happen as only the null reaction between two voters of the same type can occur. In this case one species has gone extinct and the other has reached what we call 'fixation' or 'consensus'.

This system is stochastic as which pair of voters compete at each time step is determined at random. Thus each realisation of the system will take a different number of time steps to reach consensus. Of interest is the average number of time steps needed to reach consensus, called the 'mean consensus time' and denoted by t_c. It is useful to define an alternate unit of time called 'sweeps' wherein one sweep is the average number of time steps taken for a single voter to be chosen to compete once, given by $\tau = Nt$. Accordingly we define the mean consensus time in sweeps as $\tau_c = Nt_c$. The mean consensus time will obviously depend on the initial state of the system; i.e. it will vary depending on the nature of the graph on which the voter model is defined and on the type of the voter at each vertex of that graph.

The graph defining the network of voters may be determined in a number of ways. For the case of a static graph we use a square lattice, where voters are equally spaced on the interval $[0, 1]$ in one or two dimensions with connections drawn between nearest neighbours. An illustrative example of such a lattice network in 2D is shown in figure 2.

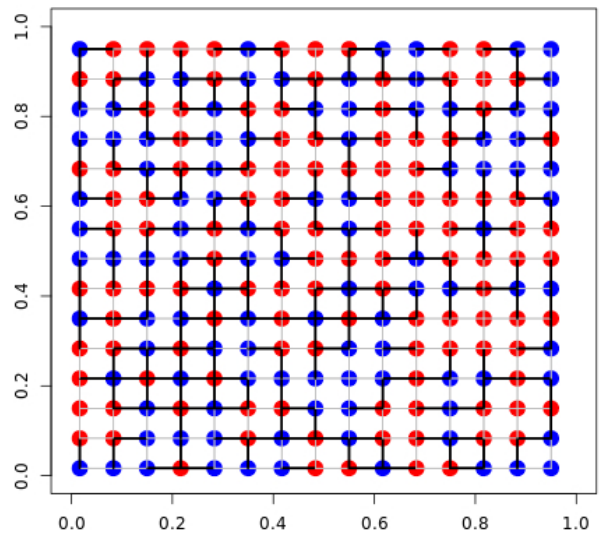

Fig. 2. Example of voters distributed on a square lattice on $[0, 1]$ in two dimensions, $N = 225$, with each voter adopting a random type. Type 0 voters are red, type 1 are blue. The lines represent links in the graph and are here present between 'nearest neighbours'. Grey lines represent 'passive links' where the pair of voters are of the same type and thus only the null reaction can occur when this pair is selected for competition. Heavy black lines represent the active links where, should that pair of voters be chosen to compete, a reaction will take place and one voter will adopt the type of the other.

We also consider flowing graphs where the positions of the voters change. For this we define an interaction distance R whereby voters whose mutual Euclidean is less than R may compete with one another. I.e., a connection on the graph exists if and only if the Euclidean distance between two vertices is less than the interaction distance. An example of this for randomly distributed voters on $[0, 1]$ in two dimensions is shown in figure 3. It is useful to define this interaction distance in units of the average spacing between voters. I.e., if there are $N = n^d$ voters distributed on $[0, 1]$ in d dimensions then we set $R = \frac{\delta}{n}$ and use δ as the variable to control the interaction radius. We do not consider $d > 2$ here.

For voters distributed on a square lattice on $[0, 1]$, setting $\delta \approx 1$ produces the same graph as in figure 2. This is important as it represents the static or no-flow limit of the flowing network model expanded upon below. Crucially this 2D fixed lattice has the property that the consensus time in sweeps scales as:

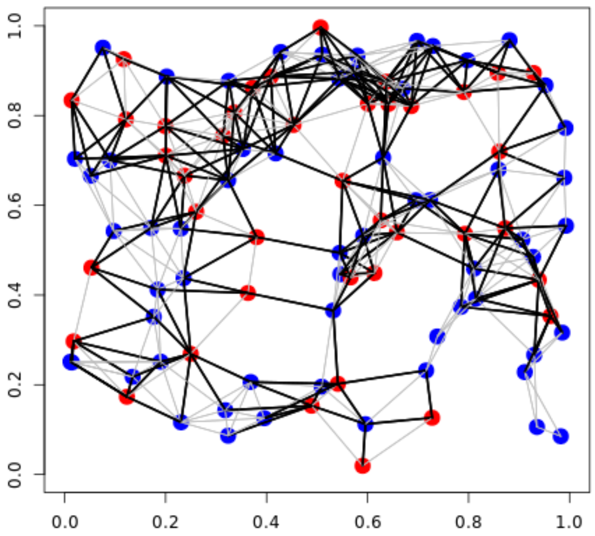

Fig. 3. Example of $N = 100$ voters arrayed randomly by a uniform distribution on $[0, 1]$, with $\delta = 1.8$. The point and line colours have the same meaning as in figure 2

$$\tau_c \sim N \log N, \quad d = 2, \tag{1}$$

demonstrated in reference [2]. The analogue in 1D is that of N voters evenly distributed on a line on $[0, 1]$ where each interacts only with its two neighbours. This has the property that the consensus time in sweeps scales as:

$$\tau_c \sim N^2, \quad d = 1. \tag{2}$$

Finally the number or density of active links as a function of the number of time steps in one dimension is given by:

$$\rho(t) \sim t^{-\frac{1}{2}}, \tag{3}$$

where ρ is the proportion of links which are active, connecting voters of differing types, and t is the number of time steps that have elapsed since the system was initialised. This is demonstrated in [4].

These analytic predictions are strictly only valid in the limit of large N, and are demonstrated by direct simulation in our results section III.

Setting $\delta \gg n$ means that every voter is connected to every other voter, a 'complete graph', an example of which is shown in figure 4. This means that the geometry of how the voters are arranged does not affect the dynamics of evolution and is equivalent to the Moran process, detailed in section II-B.

The voter model may be summarised by the following steps:

1) Select a voter at random from the population.
2) Select a second voter from those that are connected to the first.
3) Have the two voters compete and determine which one is victorious.
4) Have the loser adopt the type of the winner.
5) Repeat ad infinitum, unless one type has reached consensus, at which point stop.

The principal addition to this that we make is that the network changes after a certain number of steps. The details of this are given below in section II-C.

B. The Moran process

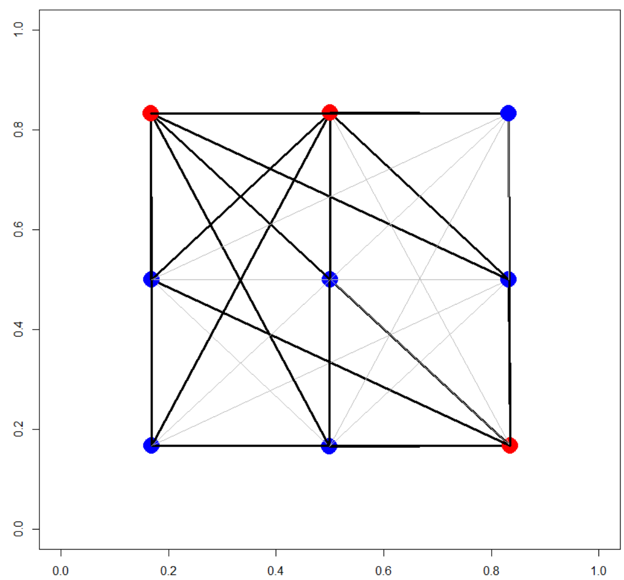

Fig. 4. This is a 'complete graph', where every voter is connected to every other voter, no matter how the voters are arranged spatially. This produces the same dynamics as the Moran process detailed in section II-B

The Moran process is equivalent to the voter model on a complete graph where each voter is connected to every other voter. Thus the spatial arrangement of the voters is not relevant to their evolution and we can describe the state of the system entirely by the number of type 1's, which we call i. This process is important because it corresponds to the limit of maximal flow of voters and is analytically soluble for both the probability that a given type reaches consensus and the mean consensus time for given initial conditions. The main quantity that we study in our simulations is the mean consensus time, so we present the analytic solution for this, following [1]:

Suppose the current state of a system of population size N is i, i.e. there are i type 1's and $N-i$ type 0's. For the system to move to state $i+1$ then a type 1 must be chosen first, a type 0 second, and the type 1 must win the exchange. These are independent events so the probability of all happening is just their product:

$$T_i^+ = \frac{i}{N}\frac{N-i}{N-1}\frac{f_1}{f} \tag{4}$$

where f_1 and f_0 are parameters that determine the relative strengths of the two types, called the fitnesses, and f is the average fitness:

$$f = \frac{if_1+(N-i)f_0}{N}. \tag{5}$$

By the same token, the probability for the system to transition from state i to state $i-1$ is

$$T_i^+ = \frac{N-i}{N}\frac{i}{N-1}\frac{f_0}{f} \tag{6}$$

and $T_i^0 = 1-(T_i^+ + T_i^-)$ is the probability that no reaction happens and the system stays in state i. For neutral selection where the victor is simply the first voter chosen and so $f_1 = f_0 = 1$.

Of interest is $\tau_c(i) = Nt_i$, the expected number of sweeps needed to reach consensus from the state i. The following equation for the mean number of time steps to consensus holds:

$$t_i = 1 + T_i^- t_{i-1} + T_i^+ t_{i+1} + T_i^0 t_i. \tag{7}$$

This may be justified term by term as being one time step (1) plus the expected number of steps remaining for the state adopted after one time step ($T_i^- t_{i-1} + T_i^+ t_{i+1} + T_i^0 t_i$). Since the consensus states $i = 0$ and $i = N$ are absorbing, we impose the boundary conditions $t_0 = t_N = 0$. This may be rearranged:

$$t_{i+1} - t_i = \frac{T_i^-}{T_i^+}(t_i - t_{i-1}) - \frac{1}{T_i^+}. \tag{8}$$

Defining $\gamma_i = \frac{T_i^-}{T_i^+}$ and $y_i = t_i - t_{i-1}$:

$$y_{i+1} = \gamma_i y_i - \frac{1}{T_i^+}, \tag{9}$$

where $y_1 = t_1 - t_0 = t_1$. This may be solved to give:

$$y_i = t_1 \prod_{k=1}^{i-1}\gamma_k - \sum_{j=1}^{i-1}\frac{1}{T_j^+}\prod_{k=j+1}^{i-1}\gamma_k. \tag{10}$$

To determine t_1 we find:

$$\sum_{m=2}^{N} y_m = t_2 - t_1 + t_3 - t_2 ... t_N - t_{N-1} = -t_1, \tag{11}$$

and so:

$$t_1 = -t_1 \sum_{m=2}^{N}\prod_{k=1}^{m-1}\gamma_k + \sum_{m=2}^{N}\sum_{j=1}^{m-1}\frac{1}{T_j^+}\prod_{k=j+1}^{m-1}\gamma_k, \tag{12}$$

which is rearranged to give t_1, the mean consensus time for a single type 1 invader in a population of $N-1$ type 0's as:

$$t_1 = \frac{1}{1+\sum_{k=1}^{N-1}\prod_{l=1}^{k}\gamma_k}\sum_{k=1}^{N-1}\sum_{j=1}^{k}\frac{1}{T_j^-}\prod_{m=j+1}^{k}\gamma_m. \tag{13}$$

In much the same way as before, we find that $t_i = -\sum_{j=i+1}^{N} y_j$, and so general t_i is given by:

$$t_i = -t_1 \sum_{j=i}^{N-1}\prod_{l=1}^{j}\gamma_l + \sum_{j=i}^{N-1}\sum_{l=1}^{j}\frac{1}{T_l^+}\prod_{m=l+1}^{j}\gamma_m. \tag{14}$$

For neutral selection this simplifies greatly as $f_0 = f_1 = 1$, so $T_i^+ = T_i^- = \frac{i(N-i)}{N(N-1)}$, and $\gamma_i = 1$. For large N This may be approximated as:

$$t_i = -N^2\left(\frac{i}{N}\log\frac{i}{N} + (1-\frac{i}{N})\log(1-\frac{i}{N})\right). \tag{15}$$

A special case that we investigate is equal initial proportions of type 1's and type 0's, i.e. with $i = \frac{N}{2}$. For this we find that the consensus time is:

$$t_c = N^2 \log 2, \tag{16}$$

i.e., the consensus time in sweeps is a linear function of N:

$$\tau_c = N \log 2. \tag{17}$$

This final result is particularly important as it represents the expected limit of maximal flow in our model. This is demonstrated numerically in the results section III.

C. Flow of the network

Our principal object of study has been the effect of flow upon a voter model defined by spatially distributed voters with an interaction radius. Here we study a discrete time model in which 'flow events' occur at particular intervals after given numbers of voter model evolution time steps have occurred. A flow event has the effect of moving the positions of the voters according to some rule, at which point the network is redrawn according to the interaction radius method.

The relative timescales for evolution and flow are then controlled by this number of evolution time steps in between flow events. We define the Damköhler number, Da, as the number of sweeps of evolution between flow events. Fractional values of Da correspond to multiple flow events per sweep of evolution.

The flow rule we use here is the well studied logistic map. If we write the coordinates of a voter before a flow event as x and y, with $(x, y) \in [0, 1]$, then the new coordinates are:

$$x' = rx(1 - x), \quad y' = ry(1 - y), \tag{18}$$

where r is some parameter between 0 and 4. For all simulations we have used $r = 3.8$. This maps points within the range $[0, 1]$ to other points within $[0, 1]$. At each flow event, the map is applied to the x and y coordinates of every voter at the same time. After the flow event, the network is redrawn by the interaction radius rule and the usual voter model evolution re-ensues.

The logistic map is one of the most comprehensively investigated objects in all of non-linear dynamics, so to expound on its features here goes beyond the scope of this project. However, we note that for $r > 3.6$ the logistic map has the important property of being chaotic, whereby two points that are initially very close to one another are likely to diverge wildly after only a few iterations of the map. This means that after many such iterations of the map, here realised as

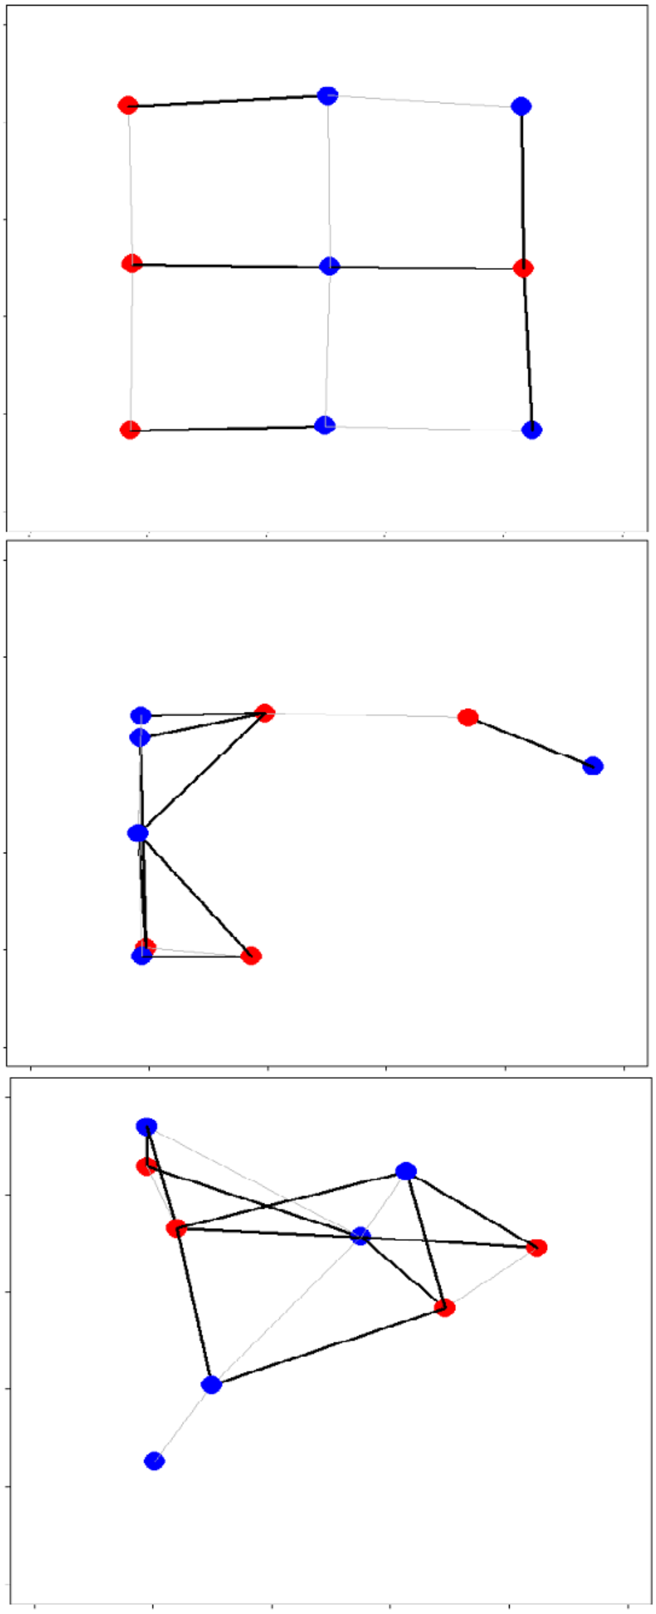

Fig. 5. An N = 9 square lattice with small Gaussian noise after 0, 10 and 30 iterations of the logistic map with $r = 3.8$ and $\delta = 1.3$. Note how the near-symmetry of the original structure is eroded after a few iterations.

flow events, the spatial correlation between two voters disappears and they have adopted new pseudo-random positions. Note however that these new pseudo-random positions are not drawn from a uniform distribution, but rather from a distribution characterised by the logistic map itself, and by the choice of parameter r. This distribution for our choice of $r = 3.8$ is shown in figure 13.

We also investigate the effect of two other similarly chaotic maps. Firstly the Hénon Map, defined by:

$$x' = 1 - ax^2 + y, \quad y' = bx, \tag{19}$$

where a and b are parameters to be chosen and where voters are disributed around $[-1, 1]$ in 2D. Secondly the Chirikov map, defined by:

$$x' = (x + y) \mod 2\pi, y' = y + K \sin x, \tag{20}$$

where K is a parameter to be chosen and voters are distributed on $[0, 2\pi]$.

These maps both have many fascinating properties which go beyond the scope of this project. The key fact is that they are all chaotic.

The key effect is that a single flow event will change the landscape of whatever network we have, breaking links and forming new ones, but spatial correlations will remain intact. This means that a link that existed before the flow event will be more likely than not still to exist after the flow event. I.e., few flow events perceptibly change the system, but do not scramble it entirely such that any voter has an equal chance of being connected to any other. However, many flow events will have this property as the effects of chaos dominate and the voters adopt pseudo-random positions. An illustration of this for the logistic map is shown in figure 5. The limit of $Da \to 0$ produces this effect, as it corresponds to many flow events per step of evolution, or the flow timescale being much larger than the evolution timescale. Intuitively, this limit should be equivalent to the Moran process as it means that any voter is equally likely to compete with any other voter during any step of evolution such that the effective graph is complete. We call this the 'well mixed' or 'maximal flow' limit.

The other limit is $Da \to \infty$, where the flow timescale is vastly slower than the evolution timescale. This corresponds to vanishingly few flow events per step of evolution. For large enough Da the system will have reached consensus long before the first flow event, and so the system is functionally static. We call this the minimal flow or static limit.

III. Simulation and Results

We developed computer simulations of the systems described above. Programming these simulations has taken up the majority of time spent on this project throughout this semester. Particular attention had to be paid to optimising the code as most algorithms we used became much slower to compute results with larger population sizes. Performing voter model simulations to consensus for population sizes greater than $N \approx 1500$ proved prohibitively computationally expensive. A single realisation could take up to an hour to compute, whilst measurements of all interesting quantities required taking averages over hundreds or even thousands of realisations. It seemed an injudicious use of time to determine exactly how these computation times scale with N, but, based on experience of how sensitive we found it to be, we would conjecture that it is a fairly high-degree polynomial. This seems reasonable as there are processes that the algorithms must perform that become more complex with increasing N, such as choosing voters for competition, which requires measuring relative positions to all other voters. So each time step becomes more computationally expensive with increasing N, and this, combined with the fact that the overall number of steps needed scales polynomially with N, means that the overall computation time should increase very fast with N.

Much of what we wanted to measure took the form of scaling laws with N which are based on approximations valid for large N. Ideally we would want to verify these by measuring them over a wide number of values of N which extends to very large values. However, since we found N greater than around 1500 totally infeasible to work with, we opted to measure lower values more accurately by increasing the number of realisations averaged over. We would then have very reliable values for, for example, the mean consensus time at these values and could then check the scaling laws. This is feasible because by construction the overall computation time scales linearly with this number of repeats, and so where we found that more precision was needed it was easy to increase the accuracy without wildly escalating the amount of time such measurements would take.

Initially we developed in C++, but found the environment needlessly cumbersome when performing analysis of the data obtained and when optimising the algorithms. We then decided to convert the program into R, and found that this improved work-flow drastically for a number of reasons. The many graphing functions in R meant that it was relatively easy to create visualisations of the processes being simulated which were impractical with C++, making it possible to get a grasp of exactly how the system evolves. We also found it easier in R to have simulations run in parallel on multiple cores, effectively octupling the processing power available to us. We hope to modify this process next semester for use on the University's Condor system.

A. Static results

We first developed simulations of the simple static system without any flow events at all, using voters arranged on a line for one dimension and a square lattice for two. The evolution of this type of system is fairly intuitive. Each step of evolution reduces the number of active links, or interfaces, by one and so regions with a high density of interfaces rapidly condense into blocs of one species type or another, called domains.

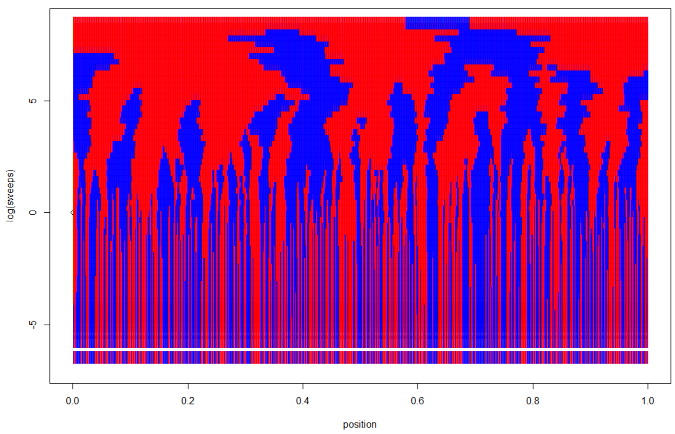

Fig. 6. Visualisation of the evolution of a static nearest neighbour voter model in 1D with $N = 625$. Time is plotted logarithmically on the vertical axis, with each vertical line representing the history of a single voter. Red and blue correspond to type 1 and 0 voters respectively, and the horizontal position of a line corresponds to its position on the interval $[0, 1]$. The voters are initially assigned random types with equal proportions. The initially heterogeneous configuration with a high density of active links rapidly evolves into one with large domains of a single type. These domains then slowly merge into each other until one type has reached consensus. This final process takes a comparatively long time as the low density of active links means that the vast majority of evolution exchanges take place between two voters of the same type where no change of the configuration occurs. This is in agreement with the long tail of the active link density as a function of time, shown in equation 3 and figure 7.

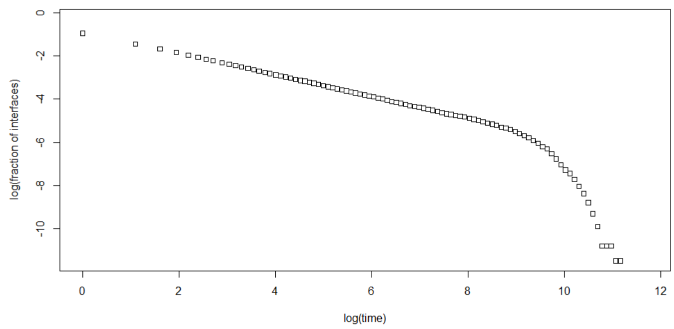

Fig. 7. The number of active links, or interfaces, in an $N = 300$, 1D nearest neighbour voter model as a function of the number of time steps since initialisation. Produced by averaging over the results of 1000 realisations of such systems. The straight line of gradient $-\frac{1}{2}$ part of this log-log plot confirms the $t^{-\frac{1}{2}}$ scaling predicted from equation 3. At large times this scaling breaks down as the number of voters of the type not destined for consensus becomes small enough that the assumptions used in the derivation of 3 cease to be valid, and the log-log plot ceases to be a straight line. The shape of this function makes a certain amount of intuitive sense. Initially ever voter has a random type, and so there are many pairs of neighbours with differing types and so a high interface density. However, because this density of active links is so high, many evolution steps will choose these active links for competition. These exchanges can only ever move interfaces or destroy them, so their density goes down rapidly. Once the density is so reduced, more and more evolution steps will choose passive links for competition where no change to the system occurs. This reduces the rate of decrease of the interface density, and corresponds to the long tail of a $t^{-\frac{1}{2}}$ distribution.

Here we use neutral selection and initial conditions with equal proportions of each type so there is no preference for which type these domains adopt. The static system is also highly symmetric, so there is no preference for where these domains appear. The evolution of the system progresses by these domains slowly merging into one another until one covers the entire network and the system has reached consensus, in this case with type 1 dominating. The dynamics here are similar to the well known Ising model of ferromagnetism at low temperature, with voter types standing in for particle spins. In one dimension, this process is readily visualised, an example of which is shown in figure 6.

The fraction of active links, which in 1D may be thought of as the number of pairs of neighbours on a line which have differing types, is predicted by [4] to scale with time as $t^{-\frac{1}{2}}$. We have verified this prediction, with the results shown in figure 7.

Equation 2 predicts that the mean consensus time for this 1D system should scale as N^2. This was measured, with results shown in figure 8.

Similarly, for a 2D lattice, equation 1 predicts that the mean consensus time should scale as $N \log N$. Similar measurements confirming this are shown in figure 9.

The consensus time itself is a random variable. Whilst the mean consensus time is the main object of interest, we have measured the distribution for which this variable is drawn, with the results shown in figures 10 and 11 for 1D and 2D respectively.

Finally, we simulated the Moran process directly. Results confirming the scaling law of equation 17 and the dependence

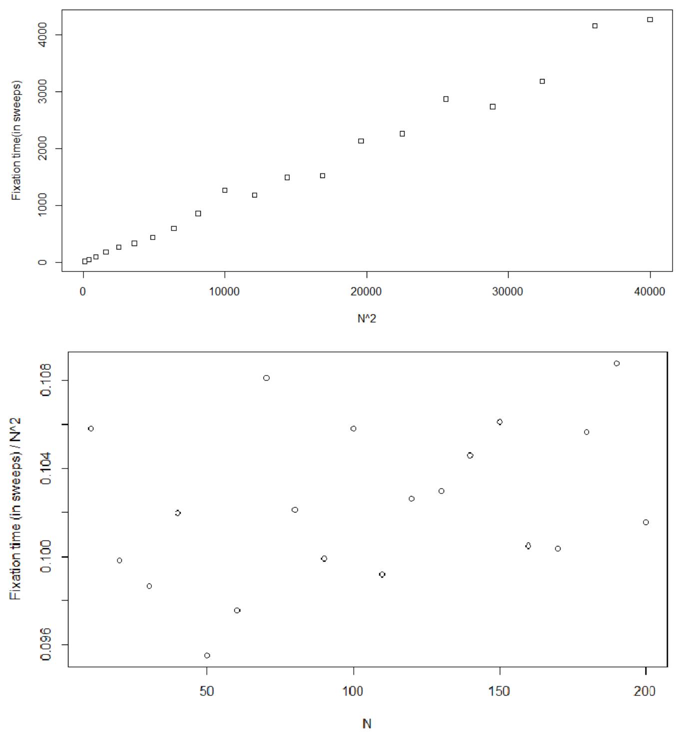

Fig. 8. The scaling of the mean consensus time, τ_c, with N for a 1D nearest neighbour interaction voter model. Produced by averaging over the results of 1000 realisations. Above is shown τ_c as a function of N^2, demonstrating the prediction of equation 2. For confirmation, below is shown $\frac{\tau_c}{N^2}$ as a function of N. As expected, there is insignificant variation in this ratio. These measurements suggest that the proportionality constant of equation 2 is around 0.1.

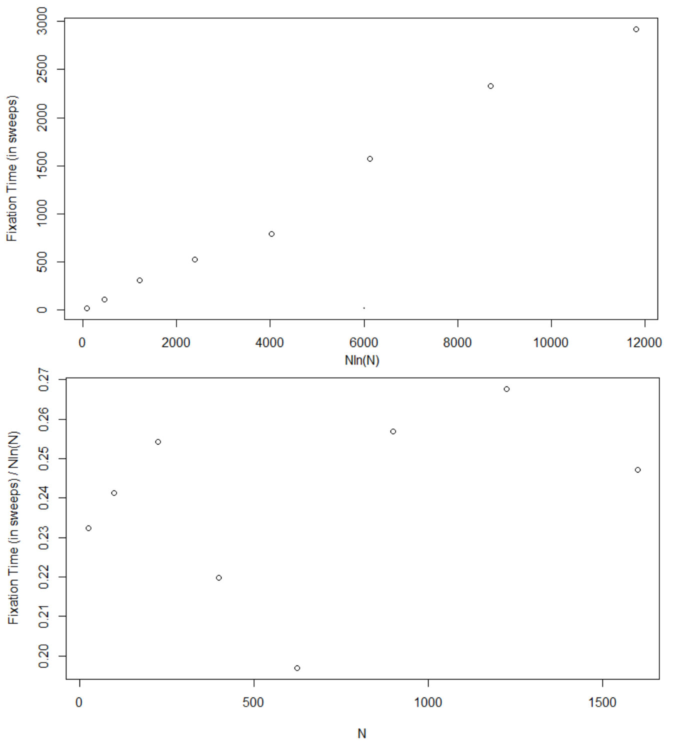

Fig. 9. The scaling of mean consensus time, τ_c, with N for a 2D square lattice voter model. As in figure 8, above is shown τ_c as a function of $N \log N$ and below is the ratio $\frac{\tau_c}{N \log N}$ as a function of N. Both reproduce the predictions of equation 1. Here the proportionality constant is around 0.2.

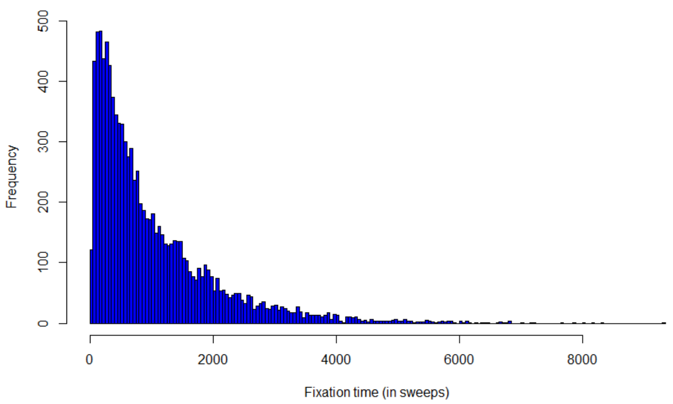

Fig. 10. A histogram of the consensus time of 10,000 realisations of a 1D nearest neighbour voter model with $N = 100$.

of the consensus time on the initial population from equation 15 are shown in figure 12.

B. Behaviour of the logistic map

The logistic map is the main flow rule we use in the simulation to move voters. We have made simulations to determine two relevant properties of this logistic flow: the probability that a flow event breaks a connection between voters, and the clustering behaviour of the logistic flow.

We have chosen $r = 3.8$ in the logistic map for our simulations. This is well into the map's region of chaotic

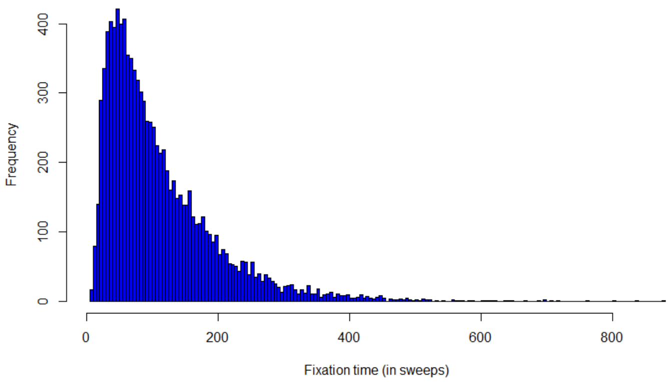

Fig. 11. A histogram of the consensus time of 10,000 realisations of a 2D square lattice voter model with $N = 100$.

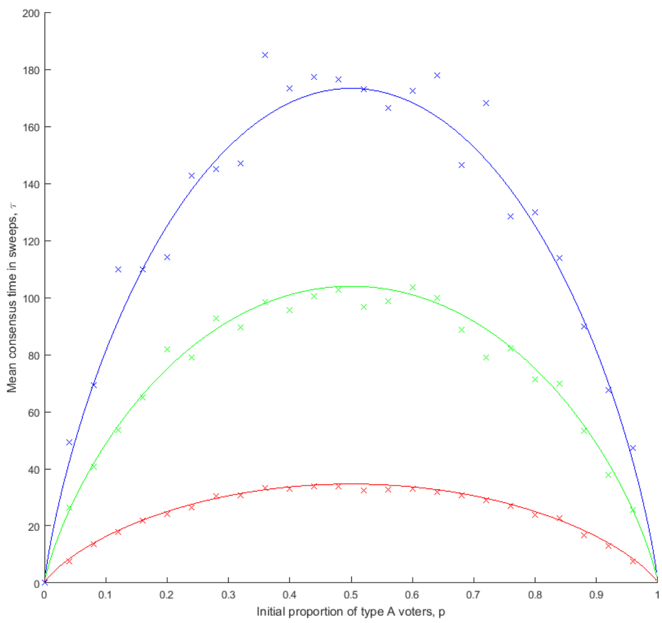

Fig. 12. Results of simulations of the Moran process. Mean consensus time is on the vertical axis, whilst the horizontal is the initial proportion of type 1 voters in the simulation. The solid lines show the exact solutions of equation 15, whilst the crosses show the results of the simulations, averaged over 100 realisations and in good agreement with the model. Red is for $N = 50$, green for $N = 150$ and blue for $N = 250$. The symmetry of these results reflects the fact that this uses neutral selection, so neither type of voter is preferred.

behaviour. Of importance is how voters will cluster under these logistic mappings; these locations will constitute localised groups of nodes which are all connected to one another. This is equivalent to determining the pseudo-random distribution referred to in section II-C. This can be determined by considering how much of its time a single voter spends in different locations over many iterations of the logistic map. I.e., the distribution will be equivalent to a normalised histogram of all the values of a long sequence of iterates of the map. The result of an implementation of this is shown in figure 13.

This distribution is very far from a uniform distribution, so it is important not to confuse a network generated by the maximal flow limit with a uniform distribution on $[0, 1]$ as shown in figure 3. Under maximal logistic flow, voters will

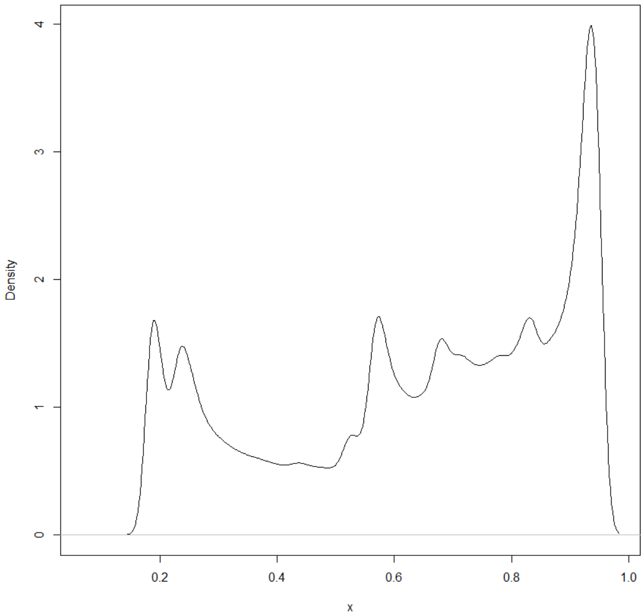

Fig. 13. The probability density function from which is drawn the pseudo-random number of the result of many iterations of the logistic map with $r = 3.8$. Generated by a histogram of the result of 3162278 iterates of the logistic map, starting at 0.1. This may be thought of as illustrating where voters will cluster after many flow events. Note that this is not the same thing as the logistic distribution, which is an entirely different object.

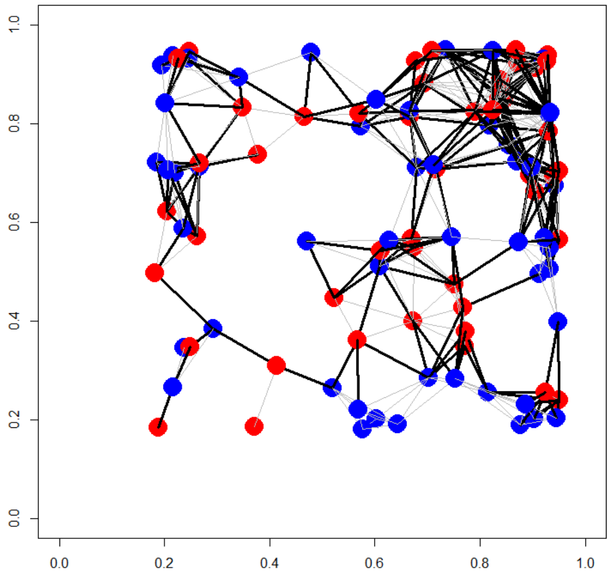

Fig. 14. A typical arrangement of voters after many (100) logistic flow events. Here $N = 100$ and $\delta = 1.6$. Note how the density of voters is much higher in regions who's x and y coordinates correspond to peaks in the density function in figure 13, such as the region around the point $(0.9, 0.9)$.

cluster at points with x and y coordinates around the peaks of this distribution. This clustering is important as connections are defined by an interaction distance and so voters tightly clustered together will likely all be connected to each other. A typical example of such a network with this logistic clustering is shown in figure 14.

We have measured the rate at which the logistic flow breaks connections between voters by considering the following. Generate coordinates of two connected voters in the range $[0, 1]$ by drawing the first, x_1 from a uniform distribution on $[0, 1]$ and the second, x_2, from a uniform distribution on $[x_1 - R, x_1 + R]$, redrawing x_2 if it falls outside $[0, 1]$ and where $R = \frac{\delta}{n}$ is the interaction radius. Then perform j iterations of the logistic map upon both coordinates and check if they are still connected, i.e. if $|x_1 - x_2| < R$. If one performs this process many times, then the proportion of such realisations which satisfy the inequality is a measurement of p_j, the probability that j iterations of the logistic map does not break a connection. The results of such measurements are shown in figure 15.

The fact that $p1$, the probability of a single flow event maintaining a given link between voters, is about 0.6 means that after a flow event only around 40% of connections will be broken, and so the network as a whole will remain similar. Thus the mean consensus time should not change very much if only a single flow event occurs between the system's initiation and its reaching consensus as the network changes only slightly, so a species close to reaching consensus before the flow event will likely still be close to reaching consensus after the flow event. However, for larger values of j, p_j decreases rapidly and reaches its value for random distribution with that value of δ at around $j = 10$. Thus for $Da \approx \frac{1}{10N}$, i.e. when there are approximately 10 flow events per step of evolution, all spatial information between voters has been eroded by the flow. This means that voters will have approximately equal chances of connecting to any other voter during a given evolution exchange. This should have the same dynamics as the Moran process, and so we conclude that the system will reach the behaviour of its well mixed limit at least by $Da = \frac{1}{10N}$, if not sooner.

C. *Effect of flow*

We investigated the effect of chaotic discrete maps upon the dynamics the voter model. We have mostly investigated using the logistic map described above, but have also obtained similar key results using other chaotic maps. This involves applying the map to the positions of all voters at regular intervals of evolution time steps as determined by the Damköhler number. Visualisations of how this affects the 1D nearest neighbour voter model is shown in figure 16.

The main effect we investigate is the effect of varying the Damköhler number, Da, on the mean consensus time. I.e., what effect does the relative size of the evolution and flow timescales have upon how long heterogeneity in the voter model can survive? Our simulations of these systems reproduce the results of the known limits of minimal flow, equations 1 and 2, and maximal flow, equation 17. Of interest

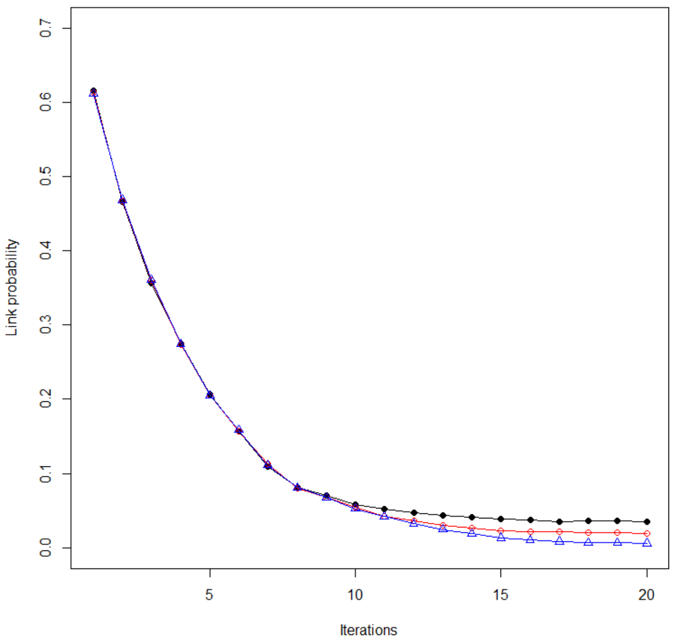

Fig. 15. A measurement of the discrete distribution p_j, the probability that j iterations of the logistic map does not break a connection between two voters in one dimension. Each data point is generated from the average of 100,000 tests of whether j iterations of the logistic map breaks a connection. The black line is for $R = 0.01$, red for $R = 0.005$ and blue for $R = 0.001$. The value of p_1 does not detectably depend on R, with all three taking the value of roughly $p_1 \approx 0.6$. All three decrease rapidly, approaching the connection probability of two voters drawn from the distribution in figure 13 for that value of the interaction radius.

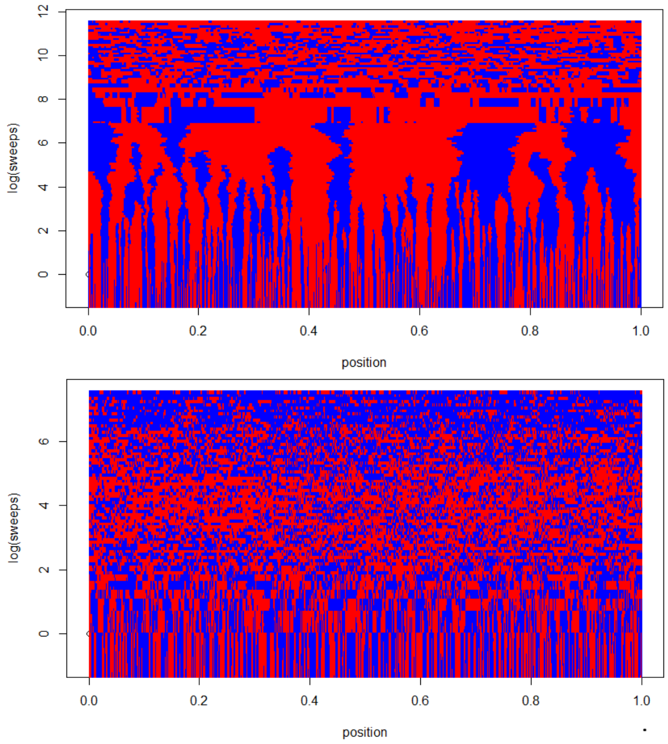

Fig. 16. Visualisations of the configuration over time of the 1D nearest neighbour voter model with logistic flow, $n = 900$. Below is shown $Da = 1$, i.e. one flow event per sweep of evolution. This is close to the well mixed limit, reflected by how the rapid domain condensation present in the static model (figure 6) is not present here. Instead voters are constantly mixed up, scrambling any domains before they have the chance to condense fully. Above is shown $Da = 1000$. This is in the region of Da which produces the maximum mean consensus time, seen in figure 17. Note how the domains seen in the static case do condense out, but that there is enough mixing from logistic flow that they do not similarly reach consensus.

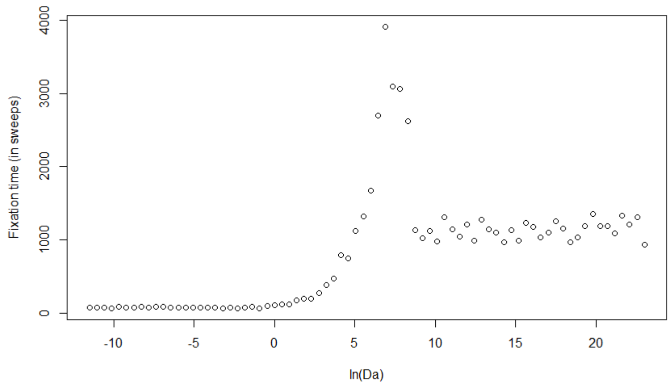

Fig. 17. Mean consensus time, τ_c, as a function of $\log Da$ for a one dimensional nearest neighbour voter model with logistic flow. $N = 100$ and $\delta = 1$. Each data point is constructed as the average consensus time from 100 realisations of the model. The limit of small Da reproduces the Moran process such that $\tau_c = N \log 2 \approx 70$ (see equation 17). The limit of large Da reproduces the static case such that $\tau_c \approx 0.1N^2 \approx 1000$ (see equation 2 and figure 8) . The maximum consensus time is reached for $Da \approx e^7 \approx 1000$. This corresponds to one flow event per thousand sweeps of evolution. Note that the value of Da which maximises $\tau_c(Da)$ is of the order of the static mean consensus time.

is the interpolation between these two limits, particularly how the mean consensus time varies as a function of Da.

Results of simulations to determine this dependency using the logistic map are shown in figures 17 and 18 for 1D and 2D respectively. Results of similar simulations using the Hénon and Chirikov maps are shown in figures 19 and 20. The Hénon map uses parameters $a = 1.4$ and $b = 0.3$. The Chirikov map uses parameter $K = 50$. Both use $N = 100$ and $\delta = 1$.

D. Interpretation

One might expect the interpolation of the mean consensus time between the well-mixed and static limits to be monotonic. Counter-intuitively we find that the mean consensus time is maximised not at either of the limits, but for an intermediate value of Da, corresponding to a critical balance of the two timescales. This is similar to the results of [3], a similar model in continuous time with advective flow. This holds for all of the chaotic maps we have used to provide the flow. Intriguingly, the value of Da which produces the maximum consensus time is consistently of an order with the mean consensus time for the static limit, with both being $\approx 10^3$ for the 1D logistic flow and $\approx 10^2$ for the 2D logistic, Hénon and Chirikov flows. This is to say that:

$$Da^* \approx \tau_c(Da \to \infty), \tag{21}$$

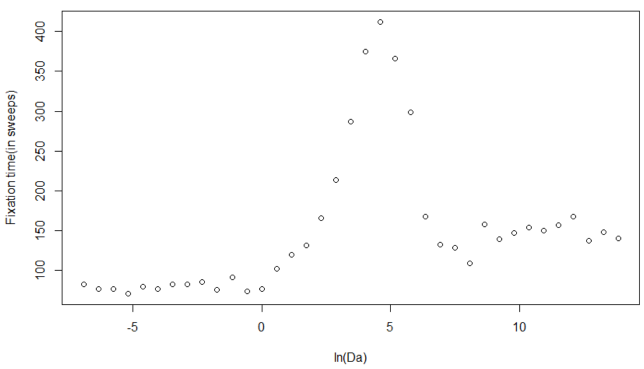

Fig. 18. Mean consensus time, τ_c, as a function of $\log Da$ for a 2D voter model with logistic flow, initially distributed on a square lattice. $N = 100$ and $\delta = 1$. Again, the limits on Da reproduce expected results. $Da \to 0$ again returns the Moran process with $\tau_c = N \log 2 \approx 70$ and $Da \to \infty$ again returns the static result of $N \approx 0.2N \log N \approx 100$ (see equation 1 and figure 9). Here the peak consensus time occurs for $Da \approx e^5 \approx 100$. As in the 1D case, this is of the order of the value of τ_c in the no flow limit.

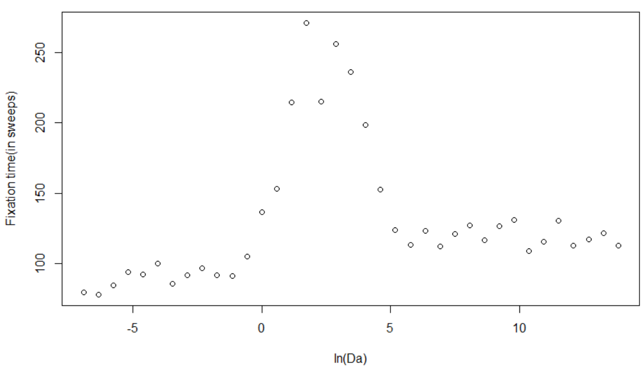

Fig. 19. Mean consensus time, τ_c, as a function of $\log Da$ for a 2D voter model with Hénon flow and with $N = 100$. This map produces a similar effect to the logistic map, with a peak occurring at roughly the value of τ_c in the static limit.

where Da^* is the critical value of Da maximising τ_c. This means that the lifetime of the mixed state where two species of voters coexist is maximised when the interval between flow events is comparable to the expected lifetime of the system without flow.

This makes a certain amount of physical sense. Having Da of this order means that over a timescale where the system would otherwise reach consensus it will expect to experience a small but non-zero number of flow events. These will have the effect of scrambling any condensed domains into one another and increasing the proportion of active links, effectively partially resetting the system back to an earlier more heterogeneous state. This will extend the overall requisite time to achieve consensus as the system must weather the scrambling effects of the chaotic flow, but must still evolve on the comparatively slow dynamics of the fixed lattice graph (as compared to the Moran process).

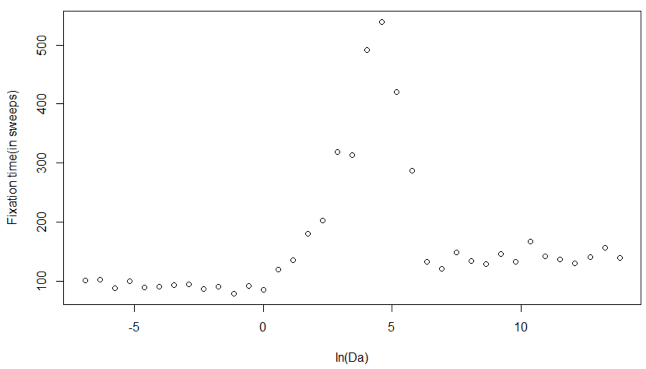

Fig. 20. Mean consensus time, τ_c, as a function of $\log Da$ for a 2D voter model with Chirikov flow and with $N = 100$. Much like the other maps, a peak occurs for an intermediate Da of the order of the static consensus time, here for $Da \approx e^4 \approx 10^2$.

If Da is significantly greater than this critical value however, then the system might expect to experience only one or even zero events before reaching consensus. Thus it would have a good chance of reaching consensus in the usual way for a static graph before a flow event occurs. This will mitigate the effects of the flow on average and cause the system to rapidly approach the dynamics of the fixed lattice once Da is large enough that the majority of realisations reach consensus before any flow events have occurred.

If, on the other hand, Da is significantly lower than the critical value then the system would expect to experience enough flow events in the usual time interval for consensus that the effects of chaos dominate and spatial correlations between voters become irrelevant. This would allow the system to rapidly reach consensus by the dynamics of the Moran process without having to wait for the comparatively longer times required by a static lattice graph.

Considering how these results might reflect real life processes, the expected lifetime of an evolving system represents a natural choice of timescale for that system. If this is adopted as the unit of time then the consensus time maximisation condition of equation 21 simply becomes $Da^* \approx 1$. I.e., that the lifetime of a mixed system under chaotic flow is maximally extended when the timescales of evolution and flow balance.

IV. Proposals for the following semester

In the model we have developed thus far, the outcome of each evolutionary exchange between voters is decided on an entirely neutral basis. For a pair of voters selected for competition there is an equal chance of either one winning the exchange and replacing the loser with a voter of its own type. A more sophisticated model could instead use asymmetric selection rules, where the outcome of an exchange may favour one or the other of the two voters.

Such rules are best phrased through game theory, where the fitnesses of each type are determined through a payoff matrix:

$$\underline{\underline{M}} = \begin{bmatrix} a & b \\ b & c \end{bmatrix}, \tag{22}$$

where M_{ij} determines the payoff received by players of a strategy i against a strategy j. a is the payoff of strategy 0 when competing against another of strategy 0, and b its payoff against one of strategy 1. Similarly, c is the payoff of strategy 1 against 0, and d is the payoff for 1 against 1. Strategies of one the two different types are here interpreted as voters of that type which always adopt the corresponding strategy in exchanges. Payoff for a given exchange is interpreted as fitness. That is to say that a voter with a higher payoff has a higher probability of winning an exchange. Suppose a type 0 and type 1 voter compete. The probability of the type 0 winning is then $\frac{b}{b+c}$ and the probability of the type 1 winning is $\frac{c}{b+c}$.

Many interesting dynamics are possible when asymmetric selection is adopted in this way. A simple example is where one strategy dominates, such that voters of its type are always at an advantage in exchanges. For example $b > c$. Since type 0 voters are then likely to win a majority of exchanges, one expects type 0 to be preferred over type 1 in terms of which one reaches consensus, and for it to do so much faster than neutral selection. This can be used to model the invasion of a small number of a more well adapted species into a population of less well adapted species.

Still more complex dynamics may be achieved if the M_{ij} are allowed to vary depending on the current state of the system. One could imagine dynamics where the fitness of a voter varies depending on the other voters near to it in space. For example, a voter's fitness could be reduced in the presence of many other voters of its type to reflect local competition for resources. Or the fitness could be increased in the same situation to reflect cooperation between members of a herd. Such complex selection rules could lead to interesting results in the evolution of the voter model, potentially leading to stable configurations where neither type of voter reaches consensus and the two coexist stably.

These types of selection dynamics are very well studied for the case of static graphs and infinite populations. We propose to investigate the effect of chaotic flows on these systems. We would aim to determine if similar results are found as for the neutral selection dynamics investigated here.

V. Conclusion

We have studied the effects of chaotic flow implemented by discrete maps on the dynamics of a spatially defined voter model through computerised simulations. The limits of maximal flow and minimal flow are demonstrated to reproduce the results of simulations of the Moran process and static lattice voter model. These results are in agreement with analytic solutions for these limits. Interpolation of the relative flow and evolution timescales between these limits shows that the lifetime of the mixed state is maximised when the two timescales balance. A physical interpretation of this result is offered: balanced timescales maximally extend the length of time the system must spend evolving by quasi-static dynamics by introducing heterogeneity at a critical rate through the chaotic flow. Faster than this critical rate and the effects of chaos dominate such that the system may evolve through the comparatively faster dynamics of the Moran process.

References

[1] A. Traulsen, C. Hauert. Stochastic evolutionary game dynamics. *Reviews of Nonlinear Dynamics and Complexity* Vol. 2, Wiley-VCH, 2009, edited by H.-G. Schuster.

[2] J. T. Cox. Coalescing Random Walks and Voter Model Consensus Times on the Torus in Zd. *The Annals of Probability*, Vol. 17, No. 4 (Oct., 1989), pp. 1333-1366, Institute of Mathematical Statistics.

[3] T. Galla, V. Prez-Muuzuri. Time scales and species coexistence in chaotic flows. 2016. arXiv:1608.05166.

[4] P. L. Krapivsky, S. Redner, E. Ben-Naim. A Kinetic View of Statistical Physics. Cambridge University Press. 2010.

[5] M. A. Nowak. Evolutionary Dynamics: Exploring The Equations Of Life. The Belknap Press of Harvard University Press. 2006.

[6] T. Galla. PHYS40571(M) Advanced Statistical Physics Notes. 2016. University of Manchester.

[7] F. Mandl. Statistical Physics, 2nd Edition, Manchester Physics Series. Wiley-Blackwell. 1988.

[8] S. H. Strogatz, Nonlinear Dynamics and Chaos, Perseus Books, 2000.