NAT-7001Y

2016-17 Natural Sciences Project

A Continuous Mathematical Model of the Extinction Vortex

Author:
100040699

Supervisors:
Dr Mark Blyth (MTH)
Prof Cock van Oosterhout (ENV)

Abstract

The extinction vortex, a feedback loop between low population and low genetic variation, is one of the primary ways in which species go extinct. A mathematical model is presented, derived from first principles based on the fundamental biological processes involved, to investigate this phenomenon and provide useful predictions and insights into its workings.

A thorough mathematical analysis of the model gives predictions for recovery, extinction, and colonisation of populations of sexually reproducing organisms. The results are confirmed by numerical approximations, as well as by an agent-based model with the same set of rules. The definitions and assumptions used in the model are clearly stated and justified.

The mathematical results derived are applied to real world biological situations, offering meaningful applications to conservation of endangered species, and to eradication of populations of invasive species. Possible extensions to the model are considered, based both on mathematical predictions and on biological phenomena not currently covered by the model.

Contents

1 Introduction **2**
1.1 Motivation . . . 2
1.1.1 The Value of Biodiversity . . . 2
1.1.2 The Value of Mathematical Modelling . . . 2
1.2 Definitions . . . 3

2 Equations **5**
2.1 Derivation . . . 5
2.2 Nondimensionalisation . . . 6

3 Methods **8**
3.1 Numerical Modelling . . . 8
3.2 Agent-Based Modelling . . . 9

4 Results **9**
4.1 Spatially Homogeneous Systems . . . 9
4.2 Travelling Waves . . . 12
4.3 Comparison to Agent-Based Model . . . 17

5 Discussion **17**
5.1 Assumptions . . . 17
5.2 Applicability to Real Populations . . . 21
5.2.1 Evidence . . . 21
5.2.2 Calculability . . . 21
5.2.3 Consequences for Endangered Species . . . 22
5.2.4 Consequences for Invasive Species . . . 23
5.3 Future Research . . . 24
5.3.1 Generation Time . . . 24
5.3.2 Heterogeneous Environments . . . 25
5.3.3 Natural Selection . . . 25
5.3.4 Predation . . . 25
5.3.5 Disease . . . 26
5.3.6 Migration . . . 26
5.3.7 Turing Instability . . . 27

6 Conclusion **28**

7 Acknowledgements **28**

References **28**

1 Introduction

1.1 Motivation

In a time dominated by increasing political unrest, social upheaval, inequality, war, and famine, it can be tempting to dismiss extinction - especially of species of no obvious benefit to humanity - as a bourgeois concern, available only to those with the luxury of not having to worry about more material problems.

1.1.1 The Value of Biodiversity

However, biodiversity does have real, meaningful benefits to all of us. Though somewhat trite by now, it bears repeating the usual list of arguments given in favour of maintaining biodiversity:

- Species can be directly useful to us. Many new drugs, for instance, are developed from compounds extracted from the natural world. And we cannot know whether a species will be useful or not until it's been thoroughly studied, so we should conserve all undocumented species just in case. (As a corollary, now that gene editing is becoming more commonplace, it can also be argued that species are valuable as a source of novel genes).
- We rely upon some species for natural resources; particularly, we depend on edible plants and animals for agriculture, on trees for forestry, and on aquatic animals for fisheries. This argument is generally meaningless, firstly because no one would seriously suggest letting species we directly rely upon go extinct, and secondly because these species are generally not threatened in the first place.
- Healthy ecosystems actively support the natural resources, above, that we depend upon. This includes, for instance, pollinators improving crop yields, and prey fish helping fishery yields. However, these secondary species are usually not irreplaceable, and this only covers a minority of species anyway.
- Some species actively improve the wider environment for humans. For instance, trees in urban environments can reduce air pollution, and rainforests and algal blooms take in carbon dioxide from the air and can help mitigate climate change. Again, however, this only applies to a minority of species, and often those which aren't threatened anyway.
- Species have a cultural value. This argument is generally perceived as by far the weakest; most people will never interact with lions or polar bears in any way other than zoos, children's books, and nature documentaries. Who cares if they're extinct in the wild? From most people's perspective, nothing would change.

However, research from the last few decades in the areas of chaos theory and Gaia theory have revealed a rather more urgent and chilling reason for maintaining biodiversity: namely, the risk of wholescale ecosystem collapse.

There is substantial historical evidence of mass extinctions, and smaller scale extinctions are known and documented in modern times. These are well confirmed by theoretical and computational models of chaotic systems (Gribbin, 2005). The most worrying predictions are the fact that responses to stresses on the environment are often highly nonlinear, with the ecosystem remaining stable for a long interval, before suddenly collapsing.

In the event of large scale ecosystem collapse, the consequences for humanity would be dire, and the natural resources we rely upon would be seriously compromised. Environmental damage has been identified as one of the major factors influencing civilisation collapse (Diamond, 2005). At the risk of catastrophising, ecosystem collapse is likely one of the few genuine existential threats we as a species face, along with climate change and nuclear war.

1.1.2 The Value of Mathematical Modelling

Even if we accept that conserving species is worthwhile, it is not obvious why an abstract, theoretical model is of any more help in preventing extinction than actively going out into the habitat and intervening directly.

There is no doubt that such direct intervention is necessary - species cannot be saved by differential equations alone - but in practice, good intentions and honest hard work all too often fail to lead to the desired results. First and foremost, we must maintain a sense of humility in the face of nature and remember that we still know very little about how ecosystems really operate, and that intervention can cause harm as well as good.

Given the harshly limited time and resources available, there is an argument to be made that we should simply make do with what knowledge we have right now, do our best, and hope it is enough. But on the other hand, the more indirect, theoretical approaches are substantially more resource efficient, and their results are often highly generalisable, making them potentially worthwhile even if we cannot afford more traditional, in depth biological investigations.

Some within the biological scientific community have argued for a change in approaches taken (Lazebnik, 2004); namely a shift from the descriptive, reductionist approach traditional in biology in favour of a more abstract and generalised approach. This can be seen, for instance, in the rise of bioinformatics, however there is room for change in all areas of biology.

Even if a mathematical model of extinction proves to be of no use whatsoever, this does not imply that it shouldn't have been done. For one thing, it is impossible to know how useful such a model will be until it has been created. The utility of mathematics is extremely difficult to judge in general, and major advances often come about in attempts to model unrelated, specific physical phenomena (Raymond, 2005).

Though far less significant in comparison to matters of civilisational collapse or major mathematical breakthroughs, on a more personal level, there is the argument to be made that it is better to be a first rate mathematician than a second rate ecologist (Hardy, 1941).

1.2 Definitions

A *population*, for the purposes of this document, is any set of living organisms occupying a well defined, continuous region of space, without complete barriers to migration or breeding. For example, a set of giant pandas living in one valley would constitute a well defined population, but the set of all giant pandas would not, as the giant pandas in zoo enclosures cannot freely migrate or interbreed with wild giant pandas.

The *environment* of a population is the space the population physically occupies, and the properties thereof. These properties include the availability of food and shelter, the local climate, and other species that interact with the population studied, such as predators.

The *population density* is the number of individual organisms per unit area of environment. In the case of populations distributed in three dimensions, i.e. aquatic organisms, it can equivalently be defined as the number of individual organisms per unit volume of environment.

The *census population* is the total number of individual organisms in a population. It can be obtained by integrating the population density u over the environment studied E with respect to area dA.

$$p = \int_E u dA$$

Census population is typically less useful in studying population dynamics than population density.

The *birth rate* and *death rate* are the average number of births and deaths respectively per unit population density per unit time.

The *replacement rate* is the birth rate minus the death rate. These rates are generally variable, and depend on both the environment and the current population density; high population density will cause more competition for resources and hence fewer births and more deaths per unit time. If the replacement rate is positive then the population density is increasing; if it is negative then the population density is decreasing; and if it is zero then the population density remains constant.

The *intrinsic* birth, death, and replacement rates are the above defined rates in the absence of the effects of competition. Equivalently, they can be considered the rates evaluated at extremely low population density. The intrinsic replacement rate can still depend on the properties of the environment and on genetic effects. In this document, the intrinsic replacement rate is given the symbol r.

The *carrying capacity* of an environment is the maximum population density that it can sustainably support in the long term. The carrying capacity can vary in space (some parts of the environment will be more hospitable than others), in time (summer may be more hospitable than winter), and with respect to other variables (for example, local predator population density), but is usually taken to be constant. In general, if the population is below its carrying capacity at a given point then the replacement rate will be positive and the population will increase; if the population is above its carrying capacity then the replacement rate will be negative and the population will decrease. One would intuitively expect populations to equilibrate over time to the carrying capacity, however this is not necessarily the case (see section 5.3.1). In this document, the carrying capacity is given the symbol κ.

The *diffusivity* of an environment is a measure of how easily animals can migrate through it. It can vary in space (terrain may be easier to traverse in some areas than in others), or in time (bodies of water may freeze and thaw with respect to season, for instance), but is usually taken to be constant. In this document, the diffusivity is given the symbol D, with subscripts where relevant.

The *genome* of an organism is the sum of its genetic information; that is, its DNA sequence.

A *gene* is a section of the genome which has some function or purpose. Usually this entails encoding for a protein. Sexually reproducing organisms generally have two copies of every gene.

A *locus* is a point on the genome at which a gene is situated.

An *allele* is a variation of a gene that has a different effect. Most genes have at least two alleles. For example, you may have a gene that dictates eye colour. This gene may have two alleles, one of which produces blue eyes and the other of which produces brown eyes.

A gene is *homozygous* if both copies of the gene are the same allele, or *heterozygous* if the two copies are different alleles.

The *average heterozygosity index* of an organism is the ratio between heterozygous and homozygous loci across the entire genome.

The *inbreeding coefficient* of an organism is one minus the average heterozygosity index.

Inbreeding is breeding of an organism with another organism which it is very close to genetically. Because the two genomes share a lot of alleles, the offspring generally has higher homozygosity. Inbreeding is more common in smaller populations because there are fewer individuals available to breed with that are not direct relatives. Inbreeding is only applicable to sexual reproduction; asexual reproduction does not constitute inbreeding.

Genetic drift is another mechanism by which small populations tend to lose genetic variation. Unless they are strongly acted upon by natural selection, alleles tend to fluctuate in frequency over time due to the random effects of mating. If the frequency of an allele falls to zero then no individuals in the population have it, so it cannot come back. The smaller the population, the more likely it is that random fluctuations will expunge an allele entirely from the population, reducing the overall genetic variation.

Inbreeding depression is the deleterious effects of low genetic variation. Alleles with negative effects are usually recessive, which means they only have an effect if they appear homozygously. Inbreeding increases the homozygosity, and hence increases the likelihood of these negative alleles being expressed (Frankham et al., 2002).

Mutation is the process by which small abnormalities during reproduction can introduce errors in the offspring's genome. Mutation introduces new genetic information, and hence increases genetic variation, and can act to balance the loss of genetic variation from other processes.

The *extinction vortex* is a phenomenon in ecology, constituting a feedback loop between low population, loss of genetic variation due to inbreeding and genetic drift, and reduced fitness due to inbreeding depression. It is one of the main mechanisms by which species go extinct (Lynch et al., 1998).

Colonisation is where a population spreads into an environment that it was not previously in.

An *invasive species* is a non-native species that causes harm to the local ecosystem, typically by either predating upon, parasitising, or competing for resources against native species. Besides the ecological damage, invasive species can also cause substantial economic damage to industries such as agriculture, forestry, fisheries, and tourism. Once colonisation has occurred, invasive species are typically extremely difficult or impossible to completely eradicate. Not all non-native species are invasive.

2 Equations

2.1 Derivation

The equations are based on the Fisher-Kolmogorov equation, defined, in dimensional form (Fisher, 1937, Kolmogorov et al., 1937),

$$\frac{\partial u}{\partial t} = D_1 \frac{\partial^2 u}{\partial x^2} + ru(\kappa - u) \tag{1}$$

where u is population density, x is space, t is time, D_1 is the diffusivity of the population, and r is the intrinsic replacement rate.

Interpreted physically, this equation states that the change in population density over time is the sum of the effects of diffusion (spreading out due to random migration), and reproduction and death. In the final term, the ru portion, being the intrinsic replacement rate multiplied by the current population density, is the maximum rate at which the population could theoretically be increasing. This is multiplied by $\kappa - u$, which corresponds to the effects of the limits of the environment. The closer the population density is to the carrying capacity, the smaller this term becomes, and hence the lower the overall replacement rate. If the population density is greater than the carrying capacity then the term becomes negative, and the replacement rate becomes negative.

The extinction vortex arises not directly due to any properties of the environment, but due to lack of genetic fitness leading to either an increased death rate, or a decreased birth rate, or both; but, ultimately, a decreased replacement rate. This implies that the replacement rate will also be a function of f, where f is the mean inbreeding coefficient. The simplest form which encapsulates the decreased replacement rate due to genetic effects is

$$r = r_0 - I(f - f_{min})$$

where r_0 is the (constant) intrinsic replacement rate in the absence of genetic effects, I is the increased death rate per unit increase in inbreeding coefficient, and f_{min} is the minimum inbreeding coefficient. Inbreeding is minimised when the population is large and healthy. In theory, we would expect that $f_{min} = 0.5$.

Substituting the above form of r into the Fisher-Kolmogorov equation (1) gives

$$\frac{\partial u}{\partial t} = D_1 \frac{\partial^2 u}{\partial x^2} + u(r_0 - I(f - f_{min}))(\kappa - u)$$

However, the mean inbreeding coefficient also varies in space and time, necessitating a similar equation for the rate of change of f. In the Fisher-Kolmogorov equations, there is a diffusive term to account for the migration of animals. As genetic fitness is a property of the animals themselves, the inbreeding coefficient will spread in space just as the animals do. Therefore,

$$\frac{\partial f}{\partial t} = D_2 \frac{\partial^2 f}{\partial x^2} + F(u, f)$$

where F is an unknown function. F must have the following behaviour:

- f is bounded between two constant limits, f_{min} and f_{max}.
- At intermediate values, f increases at a rate that is higher the lower the population is.

f_{min} is defined above. By definition, f cannot increase above 1 (corresponding to complete homozygosity), and in practice above a certain value death of any individual at that level of inbreeding is inevitable. This can be estimated by considering the number of *lethal equivalents* in the genome; that is, the number of heterozygous genes that would cause death if they were homozygous. For reference, humans have approximately three lethal equivalents.

The simplest form which encapsulates the above behaviour is:

$$\frac{\partial f}{\partial t} = D_2 \frac{\partial^2 f}{\partial x^2} + J(f - f_{min})(f_{max} - f)(\kappa - u)$$

where J is a coefficient representing the maximum rate of increase of inbreeding coefficient (or, equivalently, loss of heterozygosity) due to inbreeding. This maximum occurs when f is halfway between f_{min} and f_{max}, and when $u \ll \kappa$.

As it stands, the rate of change of the inbreeding coefficient is always greater than or equal to zero, unless there is a constant influx of new genetic variation introduced via the diffusive term. This implies that any population that ever goes below the carrying capacity will eventually go extinct, which is unrealistic; real populations can recover given time. In reality, loss of genetic variation due to processes such as inbreeding and genetic drift (as well as other evolutionary forces that reduce genetic variation, such as natural selection) are balanced by a background mutation rate.

This can be modelled by adding an extra term, such that

$$\frac{\partial f}{\partial t} = D_2 \frac{\partial^2 f}{\partial x^2} + J(f - f_{min})(f_{max} - f)(\kappa - u) - \mu u(f - f_{min})$$

where μ is the mutation rate per unit population density. The mutation term is zero when $f = f_{min}$, because at this point it is assumed that mutations are just as likely to cause hetetrozygosity as homozygosity. The rate of change of the mean inbreeding coefficient is taken to be proportional to the population density because, intuitively, more individuals mutating at once will result in more mutations, and hence a greater net increase in genetic variation.

2.2 Nondimensionalisation

The full equations, in dimensional terms, have the form

$$\frac{\partial u}{\partial t} = D_1 \frac{\partial^2 u}{\partial x^2} + u\big(r_0 - I(f - f_{min})\big)(\kappa - u) \tag{2}$$

$$\frac{\partial v}{\partial t} = D_2 \frac{\partial^2 v}{\partial x^2} + J(f - f_{min})(f_{max} - f)(\kappa - u) - \mu u f \tag{3}$$

where u is the population density, f is the mean inbreeding coefficient, D_1 and D_2 are the diffusivities of the population with respect to spread of population and mean inbreeding coefficient respectively, r_0 is the intrinsic replacement rate of the population (the mean rate of increase of population in the absence of inbreeding or limitations in the environment), I is the additional death rate per unit increase in mean inbreeding coefficient, κ is the carrying capacity (the maximum population density that the local environment can support), J is the maximum rate of increase of the mean inbreeding coefficient due to inbreeding and genetic drift, and μ is the maximum rate at which mutation can increase genetic variation.

Let the new, nondimensional variables be defined

$$\begin{aligned} u &= \kappa \tilde{u} \\ t &= \frac{1}{\mu\kappa}\tilde{t} \\ x &= \left(\frac{D_1}{\mu\kappa}\right)^{1/2}\tilde{x} \end{aligned}$$

The coefficients have dimensions of

$$\begin{aligned} [\kappa] &\sim \frac{1}{L} \\ [r_0, I, J, \mu] &\sim \frac{1}{LT} \\ [D_1, D_2] &\sim \frac{L^2}{T} \end{aligned}$$

where L is the lengthscale and T is the timescale.

Furthermore, let v be a rescaled form of f such that

$$v = \frac{f - f_{min}}{f_{max} - f_{min}}$$

noting that the range $[f_{min}, f_{max}]$ in f corresponds to the range $[0, 1]$ in v. The inbreeding coefficient is defined as one minus the heterozygosity of the genome, where the heterozygosity is the ratio between heterozygous and homozygous loci. As a ratio, it is necessarily dimensionless, hence f and v are also dimensionless and we do not require new nondimensionalised variables for either.

Substituting the nondimensionalised variables, and the new variable v into the population density equation 2 gives, in full,

$$\kappa^2 \mu \frac{\partial \tilde{u}}{\partial \tilde{t}} = D_1 \kappa \frac{\mu \kappa}{D_1} \frac{\partial^2 \tilde{u}}{\partial \tilde{x}^2} + \kappa \tilde{u}(r_0 - I(f_{max} - f_{min})v)(\kappa - \kappa u)$$

which simplifies to

$$\frac{\partial \tilde{u}}{\partial \tilde{t}} = \frac{\partial^2 \tilde{u}}{\partial \tilde{x}^2} + \frac{r_o}{\mu} \tilde{u} \left(1 - \frac{I(f_{max} - f_{min})}{r_0} v \right) (1 - \tilde{u})$$

Similarly, substituting into the inbreeding equation 3 gives, in full

$$(f_{max} - f_{min}) \mu \kappa \frac{\partial v}{\partial \tilde{t}} = D_2 (f_{max} - f_{min}) \frac{\mu \kappa}{D_1} \frac{\partial^2 v}{\partial \tilde{x}^2} + J(f_{max} - f_{min}) v \big(f_{max} - f_{min} - (f_{max} - f_{min}) v \big) (\kappa - \kappa \tilde{u}) - \mu \kappa \tilde{u} v$$

which simplifies to

$$\frac{\partial v}{\partial \tilde{t}} = \frac{D_2}{D_1} \frac{\partial^2 v}{\partial \tilde{x}^2} + \frac{J(f_{max} - f_{min})}{\mu} v(1 - v)(1 - \tilde{u}) - \tilde{u} v$$

Dropping tildes and assigning symbols to the dimensionless groups, the nondimensionalised system can be expressed as

$$\frac{\partial u}{\partial t} = \frac{\partial^2 u}{\partial x^2} + \alpha u (1 - \beta v)(1 - u) \tag{4}$$

$$\frac{\partial v}{\partial t} = D \frac{\partial^2 v}{\partial x^2} + \gamma v (1 - v)(1 - u) - uv \tag{5}$$

where

$$\begin{aligned} \alpha &= \frac{r_0}{\mu} \\ \beta &= \frac{I(f_{max} - f_{min})}{r_0} \\ D &= \frac{D_1}{D_2} \\ \gamma &= \frac{J(f_{max} - f_{min})}{\mu} \end{aligned}$$

For realistic systems, we expect that

- The intrinsic replacement rate is greater than the mutation rate.
- The maximum additional death rate due to inbreeding is greater than the intrinsic replacement rate, to ensure that loss of population is possible.
- The genetic diffusivity is less than or equal to the population diffusivity, as animals cannot have an effect on the local mean genetic variation before they have physically migrated there.
- The maximum rate of loss of genetic variation due to inbreeding is greater than the maximum mutation rate.

Mathematically, this implies that α, β, and γ are greater than 1, and that D is greater than or equal to 1.

3 Methods

3.1 Numerical Modelling

The equations were solved numerically to confirm the results of the mathematical analysis, and to investigate features which couldn't be studied using analysis alone.

The equations were approximated using an explicit method. From elementary calculus we know that

$$\frac{df}{dx} = \lim_{h \to 0} \frac{f(x+h) - f(x)}{h}$$

This can be approximated as

$$\frac{df}{dx} \approx \frac{f(x+\delta x) - f(x)}{\delta x}$$

where δx is small. Applying this approximation to all derivative terms in the full nondimensionalised equations 4 and 5, noting that the scale of variation in space δx and in time δt may not be the same, gives

$$\begin{aligned}
\frac{u(x,t+\delta t) - u(x,t)}{\delta t} &= \frac{u(x+\delta x,t) - 2u(x,t) + u(x-\delta x,t)}{\delta x^2} + \alpha u(x,t)\big(1-\beta v(x,t)\big)\big(1-u(x,t)\big) \\
\frac{v(x,t+\delta t) - v(x,t)}{\delta t} &= D\frac{v(x+\delta x,t) - 2v(x,t) + v(x-\delta x,t)}{\delta x^2} + \gamma v(x,t)\big(1-v(x,t)\big)\big(1-u(x,t)\big) - u(x,t)v(x,t)
\end{aligned}$$

In a numerical simulation, u and v are represented by arrays, with each entry being the value at a particular point in space or time. The points form a grid with a resolution of δx in the spatial dimension and δt in the temporal dimension.

For simplicity, let the ith entry in the spatial dimension and jth entry in the temporal dimension of u and v be denoted $u_{i,j}$ and $v_{i,j}$ respectively. Then the system can be written

$$\begin{aligned}
\frac{u_{i,j+1} - u_{i,j}}{\delta t} &= \frac{u_{i+1,j} - 2u_{i,j} + u_{i-1,j}}{\delta x^2} + \alpha u_{i,j}\big(1-\beta v_{i,j}\big)\big(1-u_{i,j}\big) \\
\frac{v_{i,j+1} - v_{i,j}}{\delta t} &= D\frac{v_{i+1,j} - 2v_{i,j} + v_{i-1,j}}{\delta x^2} + \gamma v_{i,j}\big(1-v_{i,j}\big)\big(1-u_{i,j}\big) - u_{i,j}v_{i,j}
\end{aligned}$$

In general, the initial distributions of u and v will be known, and the aim is to calculate the distributions at subsequent timesteps. If the system can be rearranged to have quantities evaluated at timestep $j+1$ only on the left hand side, and at timestep j only on the right hand side, then it is possible to obtain a way of explicitly calculating the distributions at the next timestep using only information from the current timestep. Rearranging thus gives

$$\begin{aligned}
u_{i,j+1} &= u_{i,j} + \frac{\delta t}{\delta x^2}\big(u_{i+1,j} - 2u_{i,j} + u_{i-1,j}\big) + \delta t \alpha u_{i,j}\big(1-\beta v_{i,j}\big)\big(1-u_{i,j}\big) \\
v_{i,j+1} &= v_{i,j} + \frac{D\delta t}{\delta x^2}\big(v_{i+1,j} - 2v_{i,j} + v_{i-1,j}\big) + \delta t \gamma v_{i,j}\big(1-v_{i,j}\big)\big(1-u_{i,j}\big) - \delta t u_{i,j}v_{i,j}
\end{aligned}$$

This is an example of an explicit method. Explicit methods are possibly the simplest technique for numerically approximating partial differential equations, however they have serious limitations. In order to converge onto the correct value, it is required that (see, for example, (Cheney and Kincaid, 2012))

$$\delta t \leq \frac{\delta x^2}{2}$$

For reasonably small δx, this necessitates an extremely small value of δt. The explicit method is thus sufficient for small, simple applications, but can quickly become prohibitively computationally expensive, especially for long timespans.

3.2 Agent-Based Modelling

An agent-based model was composed using the Python programming language in order to confirm whether or not the results of the continuous system are also applicable to a large but finite simulated population. The agents were based on diploid, sexually reproducing organisms, which is what most species vulnerable to the extinction vortex are.

The agent-based model was composed of discrete subpopulations, each containing a number of agents representing individual organisms. The agents each had their own “genome” of a given length, split into two “chromosomes”, composed of binary data (1s or 0s). The inbreeding coefficient was calculated as one minus the heterozygosity, which was itself calculated as the ratio between loci at which both chromosomes had opposite values and loci where they had the same value.

At initialisation, the model generated a population of agents distributed in the environment, each with a genome populated by random data, with an externally specified mean inbreeding coefficient.

At each subsequent timestep, the agents, with externally specified probabilities, performed any of the following actions:

- **Migration:** the agent moved, at random, either up or down one place in the environment, unless at the upper or lower limit.
- **Reproduction:** the agent chose another agent in the same subpopulation, if not alone, and produced a new agent in the same subpopulation with a genome composed of a random length section of each parent’s genome, with a small probability of mutation swapping the value at a random locus.
- **Death:** the agent was removed from the subpopulation. This could occur in one of three ways:
  - At random; there was a background death rate.
  - Due to exceeding the carrying capacity; for any new agent entering a subpopulation at carrying capacity (either due to breeding or migration), an agent in the subpopulation chosen at random was killed.
  - Due to inbreeding; there was an additional death rate that increased proportionally to the inbreeding coefficient.

There were no other rules provided. Notably, exactly the same rules were applied to the agent-based model as were in the continuous model: random migration, random reproduction, and random death, with additional deaths due to competition and inbreeding.

4 Results

4.1 Spatially Homogeneous Systems

If the population is assumed to be spatially homogeneous, then the spatial derivatives (i.e. the diffusive terms) can be neglected, which simplifies the equations to a system of ordinary differential equations:

$$\frac{du}{dt} = \alpha u(1-u)(1-\beta v) \tag{6}$$

$$\frac{dv}{dt} = \gamma v(1-v)(1-u) - uv \tag{7}$$

where u and v are now functions of time only.

To identify equilibrium points, we impose that

$$\frac{du}{dt} = 0, \quad \frac{dv}{dt} = 0$$

This gives four equilibrium points:

$$(u,v) = (0,0), (1,0), (0,1), \left(\frac{\gamma - \frac{\gamma}{\beta}}{1 + \gamma - \frac{\gamma}{\beta}}, \frac{1}{\beta} \right)$$

These correspond to, respectively: a population that is extinct but genetically healthy, a population that is at carrying capacity and genetically healthy, a population that is extinct and genetically as unhealthy as possible, and a borderline case with intermediate population density and inbreeding coefficient.

In order to find the stability and classification of the equilibrium points, we use perturbation analysis of a generalised equilibrium point given by (u_0, v_0). For simplicity, we denote

$$\begin{aligned} F(u,v) &= \alpha u(1-u)(1-\beta v) \\ G(u,v) &= \gamma v(1-v)(1-u) - uv \end{aligned}$$

Now we introduce a small perturbation about the equilibrium point:

$$\begin{aligned} u &= u_0 + \tilde{u} \\ v &= v_0 + \tilde{v} \end{aligned}$$

where $\tilde{u}, \tilde{v} \ll 1$. Substituting this into the full equations, we obtain

$$\begin{aligned} \frac{d\tilde{u}}{dt} &= F(u_0 + \tilde{u}, v_0 + \tilde{v}) \\ \frac{d\tilde{v}}{dt} &= G(u_0 + \tilde{u}, v_0 + \tilde{v}) \end{aligned}$$

Expanding as a Taylor series, and neglecting $O(\tilde{u}^2)$,

$$\begin{aligned} \frac{d\tilde{u}}{dt} &= F(u_0, v_0) + \tilde{u}F_u(u_0, v_0) + \tilde{v}F_v(u_0, v_0) + ... \\ \frac{d\tilde{v}}{dt} &= G(u_0, v_0) + \tilde{u}G_u(u_0, v_0) + \tilde{v}G_v(u_0, v_0) + ... \end{aligned}$$

where F_u is the derivative of F with respect to u. Given that $F(u_0, v_0) = G(u_0, v_0) = 0$ by definition, this simplifies to

$$\frac{d}{dt} \begin{pmatrix} \tilde{u} \\ \tilde{v} \end{pmatrix} = J \begin{pmatrix} \tilde{u} \\ \tilde{v} \end{pmatrix}$$

where J is the Jacobian matrix, defined

$$J = \begin{pmatrix} F_u & F_v \\ G_u & G_v \end{pmatrix}$$

with all derivatives evaluated at the equilibrium point (u_0, v_0).

We now assume that the small perturbations have the form

$$\tilde{u} = \tilde{A}e^{\lambda t}, \quad \tilde{v} = \tilde{B}e^{\lambda t}$$

Substituting this into the above system and rearranging gives

$$(J - \lambda I) \begin{pmatrix} \tilde{A} \\ \tilde{B} \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$$

where I is the identity matrix. For nontrivial solutions, we take the determinant of the matrix $J - \lambda I$ to be zero.

For this model, the general form of the entries of the Jacobian are

$$\begin{aligned}
F_u &= \alpha(1-\beta v_0)(1-2u_0) \\
F_V &= -\alpha\beta u_0(1-u0) \\
G_u &= -\gamma v_0(1-v_0) - v_0 \\
G_v &= \gamma(1-2v_0)(1-u_0) - u_0
\end{aligned}$$

We now consider the four equilibrium points in turn. At $(u_0, v_0) = (1, 0)$,

$$\begin{pmatrix} -\alpha-\lambda & 0 \\ 0 & -1-\lambda \end{pmatrix} \begin{pmatrix} \tilde{A} \\ \tilde{B} \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$$

with eigenvalues of

$$\lambda = -\alpha, \quad -1.$$

Given that $\alpha > 1$, both eigenvalues are real and negative, which implies that the point is a stable node. This makes intuitive sense, as we would expect a large population with minimal inbreeding to be able to recover from small perturbations.

At $(u_0, v_0) = (0, 1)$,

$$\begin{pmatrix} -\alpha(1-\beta)-\lambda & 0 \\ 0 & -\gamma-\lambda \end{pmatrix} \begin{pmatrix} \tilde{A} \\ \tilde{B} \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$$

with eigenvalues of

$$\lambda = -\alpha(\beta-1), \quad -\gamma.$$

Given that α, β, γ are all greater than 1, both eigenvalues are again real and negative, which implies that the point is also a stable node. Again this makes intuitive sense, because we would expect a small perturbation (here equivalent to a small, highly inbred population) to vanish with time.

At $(u_0, v_0) = (0, 0)$,

$$\begin{pmatrix} \alpha-\lambda & 0 \\ 0 & \gamma-\lambda \end{pmatrix} \begin{pmatrix} \tilde{A} \\ \tilde{B} \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$$

with eigenvalues of

$$\lambda = \alpha, \quad \gamma$$

Since α and γ are both greater than 1, both eigenvalues are real and positive, which implies that the point is an unstable node. This makes sense because a population that is small but genetically optimal can be expected to recover to full health.

At the final point, $(u_0, v_0) = \left(\frac{\gamma-\frac{\gamma}{\beta}}{1+\gamma-\frac{\gamma}{\beta}}, \frac{1}{\beta}\right)$, the algebra becomes too cumbersome to transcribe. Nevertheless we find that one eigenvalue is positive and the other negative, implying that it is a saddle point. This finding is confirmed by the numerical results presented below.

It is now clear why it was assumed in section 2.2 that the nondimensional groups are greater than one. Besides the arguments put forward there, letting them be less than or equal to one can lead to some of the equilibrium points changing classifications, leading to them no longer being biologically meaningful. For example, letting β be less than one would mean that the point $(0, 1)$ would become unstable, meaning that the population would be unable to go extinct.

Figure 1 shows the full phase plane for the spatially homogeneous problem, generated numerically. The stable, unstable and saddle points are clearly identifiable, confirming the analytical results derived above.

This shows that, in the spatially homogeneous case, with perturbation, there are only two long term behaviours possible: complete recovery (the system goes to $(1, 0)$), or extinction (the system goes to $(0, 1)$). Figure 2 shows the long term result of numerically integrating (6) using Euler's method from any point in a grid within the phase plane. Though there is some inaccuracy owing to the use of Euler's method, it broadly agrees with what would be expected from the phase plane.

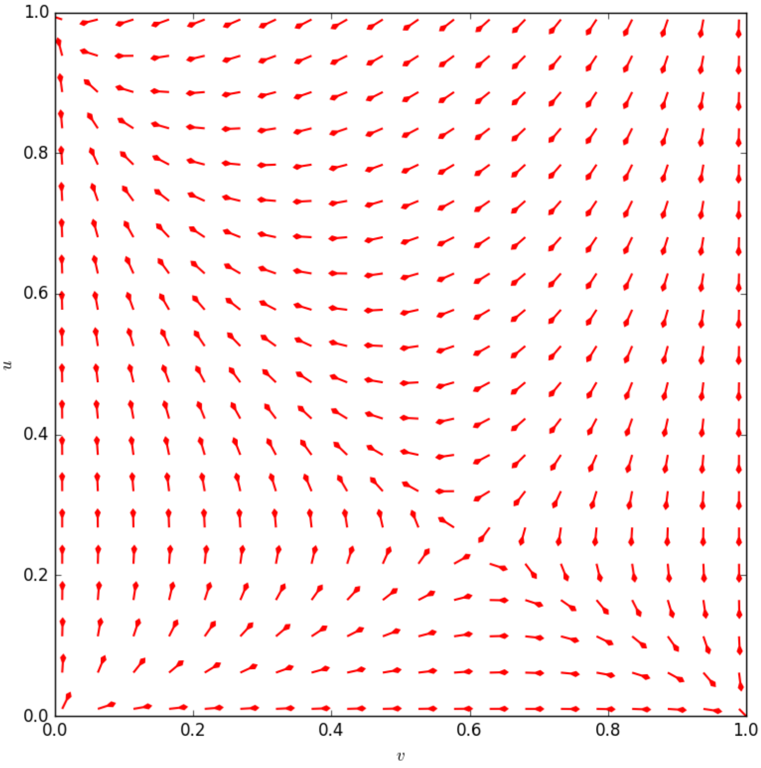

Figure 1: The phase plane for the spatially homogeneous case. The horizontal axis, v, is the rescaled mean inbreeding coefficient, while the vertical axis, u, is the population density. Generated using parameter values $\alpha = 1$, $\beta = 4$, $\gamma = 2$.

4.2 Travelling Waves

Waves travelling at constant speed are equivalent to distributions in space and time that are static in a frame of reference moving at the same speed as the wave. This moving frame of reference can be imposed by introducing a new coordinate

$$\zeta = x - ct$$

where c is the (constant) wavespeed. In order to convert the system into the new coordinate system, we first consider the derivatives in the equations, for a general function w. Using the chain rule,

$$\begin{aligned}
\frac{\partial w}{\partial t} &= \frac{dw}{d\zeta}\frac{\partial \zeta}{\partial t} \\
&= -c\frac{dw}{d\zeta} \\
\frac{\partial w}{\partial x} &= \frac{dw}{d\zeta}\frac{\partial \zeta}{\partial x} \\
&= \frac{dw}{d\zeta} \\
\frac{\partial^2 w}{\partial x^2} &= \frac{\partial}{\partial x}\left(\frac{\partial w}{\partial x}\right) \\
&= \frac{\partial \zeta}{\partial x}\frac{d}{d\zeta}\left(\frac{dw}{d\zeta}\right) \\
&= \frac{d^2 w}{d\zeta^2}
\end{aligned}$$

Substituting these derivatives into the full, nondimensionalised equations 4 and 5 gives

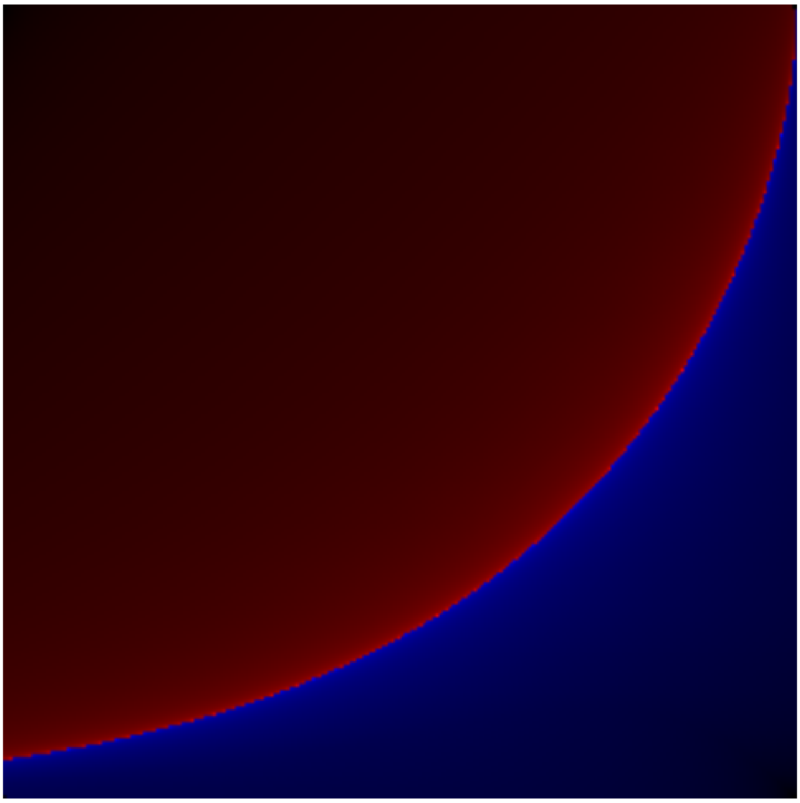

Figure 2: A colour map of the phase plane, as in figure 1. Red indicates initial positions from which the population ultimately recovers; blue indicates initial positions from which the population ultimately goes extinct. Lightness is proportional to the number of timesteps taken to come within 0.005 of either $(1, 0)$ or $(0, 1)$. Generated using the same parameter values: $\alpha = 1$, $\beta = 4$, $\gamma = 2$, with a timestep for Euler's method of $dt = 0.02$.

$$\begin{aligned}-cu' &= u'' + \alpha u(1-u)(1-\beta v)\\ -cv' &= Dv'' + \gamma v(1-v)(1-u) - uv\end{aligned}$$

where primes denote differentiation with respect to ζ. Introducing new variables $p = u'$ and $q = v'$, this can be rearranged into a dynamical system of first order ordinary differential equations:

$$\begin{aligned}u' &= p\\ v' &= q\\ p' &= -cp - \alpha u(1-\beta v)(1-u)\\ q' &= \frac{1}{D}\big(-cq - \gamma v(1-v)(1-u) - uv\big)\end{aligned}$$

Equilibrium points can be found by imposing that $u' = v' = p' = q' = 0$. This gives four equilibrium points:

$$(u_0, v_0, p_0, q_0) = (0,0,0,0), (1,0,0,0), (0,1,0,0), \left(\frac{\gamma - \frac{\gamma}{\beta}}{1+\gamma-\frac{\gamma}{\beta}}, \frac{1}{\beta}, 0, 0\right)$$

These are the same as the equilibrium points observed in the spatially homogeneous case in section 4.1, with the addition of p and q both being zero. Though the equilibrium points may be the same, the stability and classifications of the equilibrium points can be expected to be different, both because the system of equations being solved is different, and because there are two extra dimensions of phase space.

The perturbation analysis outlined in section 4.1 applies equally well to the four dimensional problem. For simplicity, we denote

$$\begin{aligned}F(u,v,p,q) &= p\\ G(u,v,p,q) &= q\\ H(u,v,p,q) &= -cp - \alpha u(1-\beta v)(1-u)\\ K(u,v,p,q) &= \frac{1}{D}\big(-cq - \gamma v(1-v)(1-u) - uv\big)\end{aligned}$$

As before, we now introduce a small perturbation about the equilibrium point:

$$\begin{aligned}u &= u_0 + \tilde{u}\\ v &= v_0 + \tilde{v}\\ p &= p_0 + \tilde{p}\\ q &= q_0 + \tilde{q}\end{aligned}$$

where $\tilde{u}, \tilde{v}, \tilde{p}, \tilde{q} \ll 1$.

Substituting into the equations, expanding as a Taylor series, and neglecting small terms then leaves

$$\frac{d}{d\zeta}\begin{pmatrix}\tilde{u}\\ \tilde{v}\\ \tilde{p}\\ \tilde{q}\end{pmatrix} = J\begin{pmatrix}\tilde{u}\\ \tilde{v}\\ \tilde{p}\\ \tilde{q}\end{pmatrix}$$

where J is the four dimensional Jacobian matrix, with the form

$$J = \begin{pmatrix}F_u & F_v & F_p & F_q\\ G_u & G_v & G_p & G_q\\ H_u & H_v & H_p & H_q\\ K_u & K_v & K_p & K_q\end{pmatrix}$$

with all derivatives evaluated at the equilibrium point (u_0, v_0, p_0, q_0).

Now assuming that the small perturbations have the form

$$\tilde{u} = \tilde{A}e^{\lambda\zeta}, \quad \tilde{v} = \tilde{B}e^{\lambda\zeta}, \quad \tilde{p} = \tilde{C}e^{\lambda\zeta}, \quad \tilde{q} = \tilde{D}e^{\lambda\zeta}$$

and rearranging gives

$$(J - \lambda I)\begin{pmatrix}\tilde{A}\\ \tilde{B}\\ \tilde{C}\\ \tilde{D}\end{pmatrix} = \begin{pmatrix}0\\0\\0\\0\end{pmatrix}$$

or, in full,

$$\begin{pmatrix} -\lambda & 0 & 1 & 0 \\ 0 & -\lambda & 0 & 1 \\ -\alpha(1-\beta v_0)(1-2u_0) & \alpha\beta u_0(1-_0) & -c-\lambda & 0 \\ \frac{\gamma}{D}v_0(1-v_0)+\frac{1}{D}v_0 & -\frac{\gamma}{D}(1-2v_0)(1-u_0)+\frac{1}{D}u_0 & 0 & -\frac{c}{D}-\lambda \end{pmatrix}\begin{pmatrix}\tilde{A}\\ \tilde{B}\\ \tilde{C}\\ \tilde{D}\end{pmatrix} = \begin{pmatrix}0\\0\\0\\0\end{pmatrix}$$

When calculating the determinant, and finding the roots, the algebra becomes prohibitively complicated to display in full. Using computational methods, it was found that the eigenvalues were, at $(1, 0, 0, 0)$,

$$\lambda = -\frac{c}{2D} \pm \frac{1}{2D}\sqrt{c^2 + 4D}, \quad -\frac{c}{2} \pm \frac{1}{2}\sqrt{c^2 + 4\alpha}$$

Since $\alpha > 0$ and $D > 0$, there are two positive and two negative eigenvalues, and hence this point is unstable. Similarly, at $(0, 1, 0, 0)$,

$$\lambda = -\frac{c}{2D} \pm \frac{1}{2D}\sqrt{c^2 + 4D\gamma}, \quad -\frac{c}{2} \pm \frac{1}{2}\sqrt{c^2 + 4\alpha(\beta - 1)}$$

which, since $\gamma > 0$ and $\beta > 1$ is also unstable.

At $(0, 0, 0, 0)$,

$$\lambda = -\frac{c}{2D} \pm \frac{1}{2D}\sqrt{c^2 - 4D\gamma}, \quad -\frac{c}{2} \pm \frac{1}{2}\sqrt{c^2 - 4\alpha}$$

As all eigenvalues have negative real part, this point is stable.

The closed form of the eigenvalues at the remaining stationary point are too cumbersome to write in full, but we find that this point is unstable.

We cannot allow spirals at $u = 0$, as this requires the population density to become negative, which makes no physical sense. On this basis, we require that

$$c > 2\sqrt{D\gamma}$$
$$c > 2\sqrt{\alpha}$$

This implies that, for physically meaningful travelling wave solutions, the wave has a minimum speed which is either $2\sqrt{D\gamma}$ or $2\sqrt{\alpha}$. If the latter, the minimum wavespeed is proportional to $\sqrt{\alpha}$, which is the intrinsic replacement rate relative to the mutation rate. A greater replacement rate will generally result in a greater rate of increase of population at the leading edge, which intuitively will result in a faster wavespeed of u.

Otherwise the wave speed is proportional to $\sqrt{D\gamma}$. A larger value of D will mean that population diffuses faster than the genetic variation, while a larger value of γ will mean that the population will locally lose genetic variation more easily. A greater value of $sqrtD\gamma$ will thus intuitively result in a faster wavespeed of v.

For physically plausible values of α, β, γ, D and c, we find that the point $(0, 0, 0, 0)$ is a stable node, and the remaining three equilibrium points are unstable.

However, in order to explicitly show that travelling wave solutions exist, we need to show that a corresponding trajectory plausibly exists in phase space. This trajectory would have the properties that

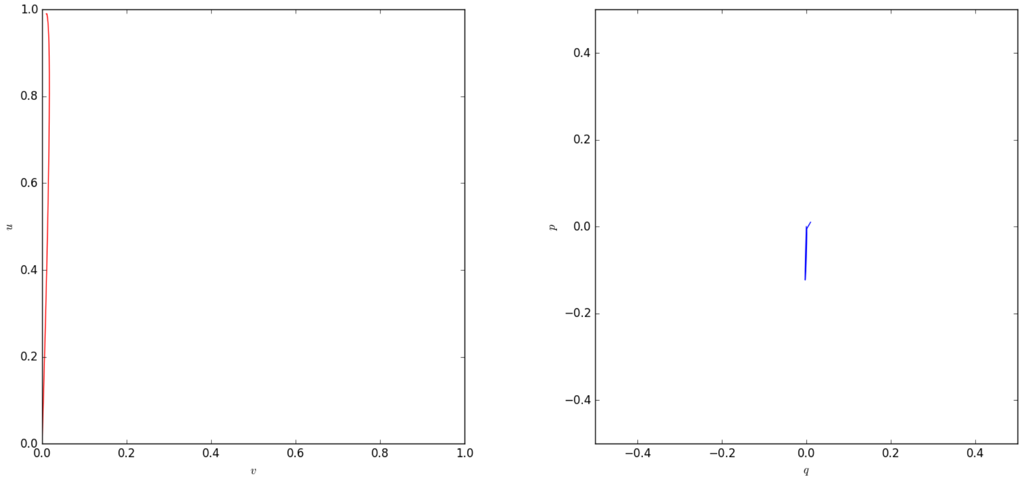

Figure 3: A plausible trajectory corresponding to a population front. Left: the projection of the trajectory onto the u, v plane. Right: the projection of the trajectory onto the p, q plane.

- The trajectory would start from a healthy equilibrium state; as $\zeta \to -\infty$, $u \to 1$ and $v \to 0$.
- The trajectory would end in some state of zero population; as $\zeta \to \infty$, $u \to 0$.
- Between these two points, the trajectory would not pass through any states that are not physically meaningful, for instance where u is negative.

The point $(1, 0, 0, 0)$ plausibly fulfils the first condition; it is unstable, so trajectories starting there will not stay there. Similarly, the point $(0, 0, 0, 0)$ plausibly fulfils the second condition, as it is stable. Not only that, but it is the only stable state of the entire system, so we would expect all trajectories to ultimately end there.

Proving the existence of such a trajectory would be reasonably straightforward for a two-dimensional phase plane, but for the full four-dimensional phase space it is substantially more difficult. Instead, figure 3 shows the result of a numerical integration of the problem, performed using Euler's method.

In the u, v plane, the trajectory begins from a small perturbation away from the point $(1, 0)$, and travels downwards to the point $(0, 0)$, where it stays indefinitely. The trajectory does not leave the unit square. In the p, q plane, it begins from a small perturbation away from $(0, 0)$, travels down in p, before returning to $(0, 0)$.

Physically, this would constitute a case where, at large negative x, the population is at the carrying capacity, with minimum inbreeding; at large positive x, the population is at zero, with minimum inbreeding; and between the two, a travelling wave front moves at speed c towards positive x, with a small peak of inbreeding.

Though less plausible biologically, another travelling wave solution is the trajectory from $(0, 1, 0, 0)$ to $(0, 0, 0, 0)$, as shown in figure 4. This is essentially the same case, but with the population density and inbreeding coefficient swapped. That is, a small peak of population travelling with an inbreeding coefficient front towards positive x.

This is not a situation seen in reality, because it would require a population small enough to be endangered to be constantly migrating. Unlike a population near carrying capacity, this population would be unable to recover from any significant losses, and would likely collapse as soon as it reached the edge of its habitat.

The Fisher-Kolmogorov equation is known to have travelling wave solutions, which travel with wavespeed $c = 2$, and which can be found numerically as well (see figure 5).

Similarly, the travelling waves predicted for the full system of equations above can be found numerically for values of constants where inbreeding is well balanced by mutation (figure 6), but for values where inbreeding is significantly greater than mutation, the travelling waves come to a stop (figure 7). As travelling waves with non-constant wavespeed cannot be found by the above

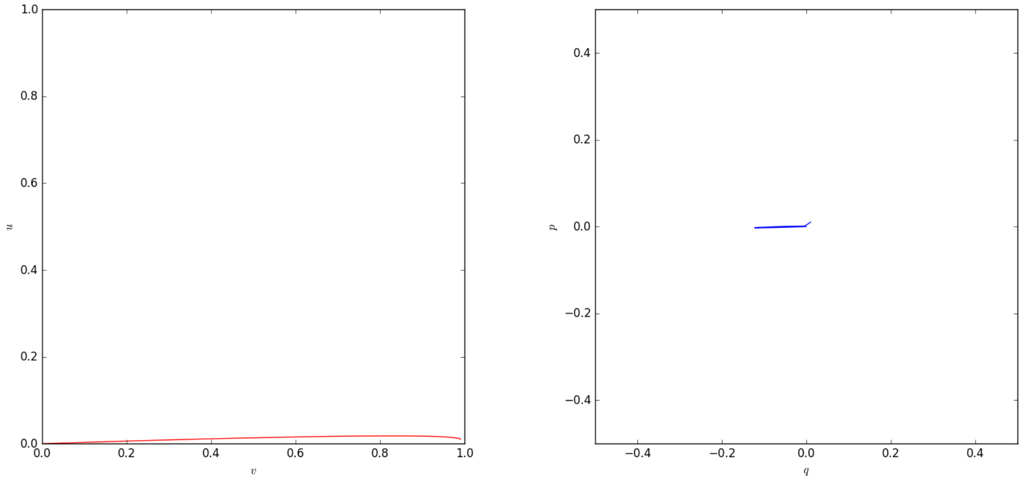

Figure 4: A plausible trajectory corresponding to an inbreeding front. Left: the projection of the trajectory onto the u, v plane. Right: the projection of the trajectory onto the p, q plane.

mathematical analysis, further investigation is needed to confirm whether this is real behaviour of the equations or simply an artefact of the numerical approximation.

4.3 Comparison to Agent-Based Model

The agent based model showed, at least qualitatively, all of the same behaviour as the continuous model. The (non-rescaled) mean inbreeding coefficient generally stayed close to $f = 0.5$, apart from in areas with a population far below the carrying capacity, where it increased.

In spatially homogeneous cases, the population always either reached a stable recovery state or went entirely extinct, although this was influenced by random variation.

Travelling waves could be observed under appropriate conditions (see figure 8), and appeared to travel faster with higher diffusivity or replacement rate, and with lower sensitivity to inbreeding.

5 Discussion

5.1 Assumptions

In order for the model to be applicable to real populations, it is important to make clear the assumptions implicit in the above derivation.

The first and most obvious assumption is that the population density and inbreeding coefficient can be modelled using a continuous, deterministic system. In reality, this is a stochastic process acting on discrete agents.

The Fisher-Kolmogorov equation, on which the model is partially based, provides a precedent in treating relatively large populations of animals as continuous, however the objective of investigating the extinction vortex necessarily means that it must be applicable to small populations as well.

Even if the assumption of continuity does break down when the population is reduced to the scale of individuals, however, it does tend towards the correct expected behaviour as the population decreases, in either space or time, from a larger population.

More specifically, as a population decreases over time, the model correctly predicts that the population tends towards extinction; and as a population decreases in space, it correctly predicts that it approaches zero population density.

The assumption of continuity may be less well founded for the inbreeding coefficient. The problem is that the mean inbreeding coefficient is only well defined where the population is non-zero. At best, the limit of v as x approaches infinity can be used as a substitute for the value it approaches at the edge of the physical population. But continuous distributions, which the model deals with, do not have a well defined edge.

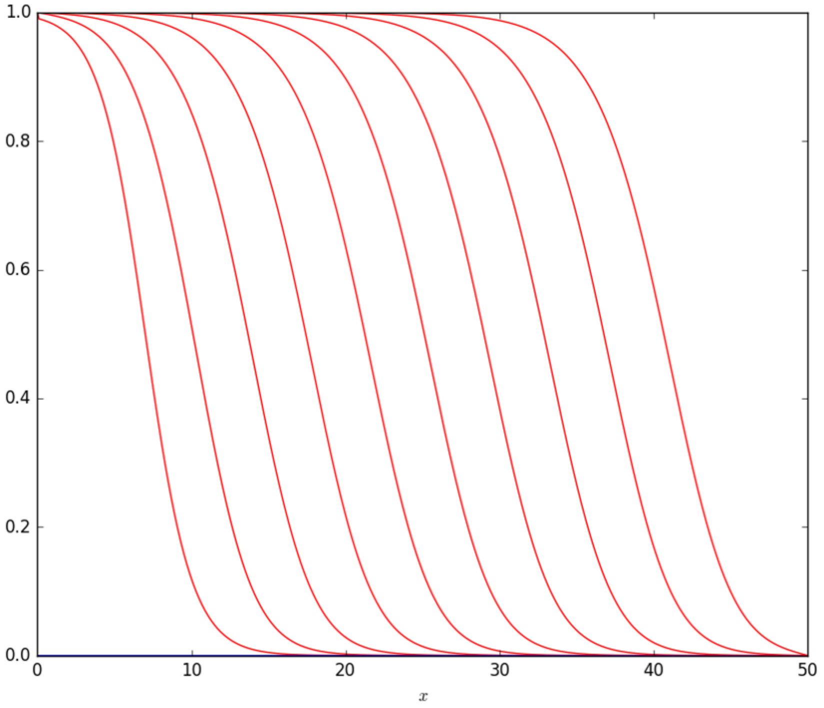

Figure 5: A numerical integration of the Fisher-Kolmogorov equation. Red is population density, u. Lines are taken at equal time intervals. The parameters used were $r = 1$, $\kappa = 1$, $D = 1$. The integration was performed with lengthscale $\delta x = 0.1$ and timescale 0.002.

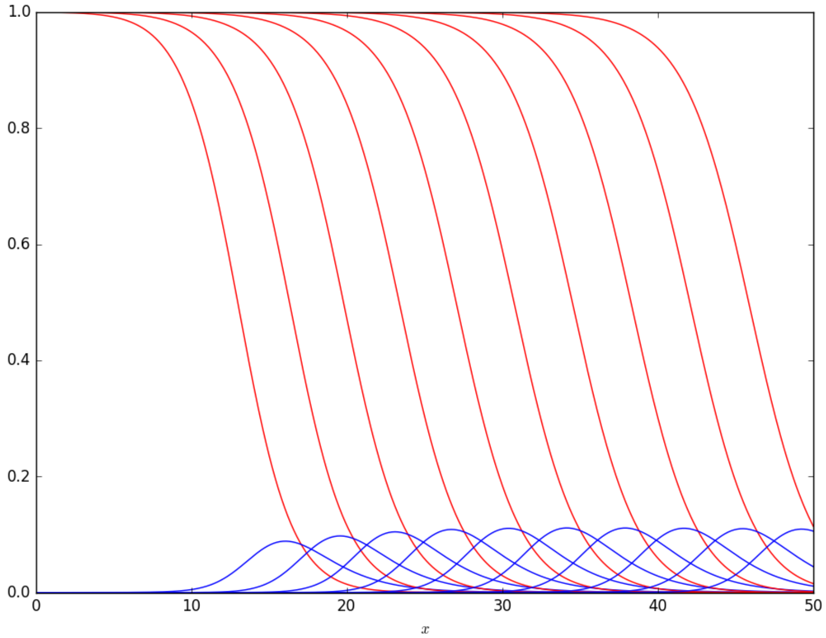

Figure 6: A numerical integration of the full system of equations, with a relatively high mutation rate. Red is population density, u, and blue is rescaled mean inbreeding coefficient, v. Lines are taken at equal time intervals. The parameters used were $\alpha = 2$, $\beta = 1.2$, $\gamma = 1.5$, $D = 1$.

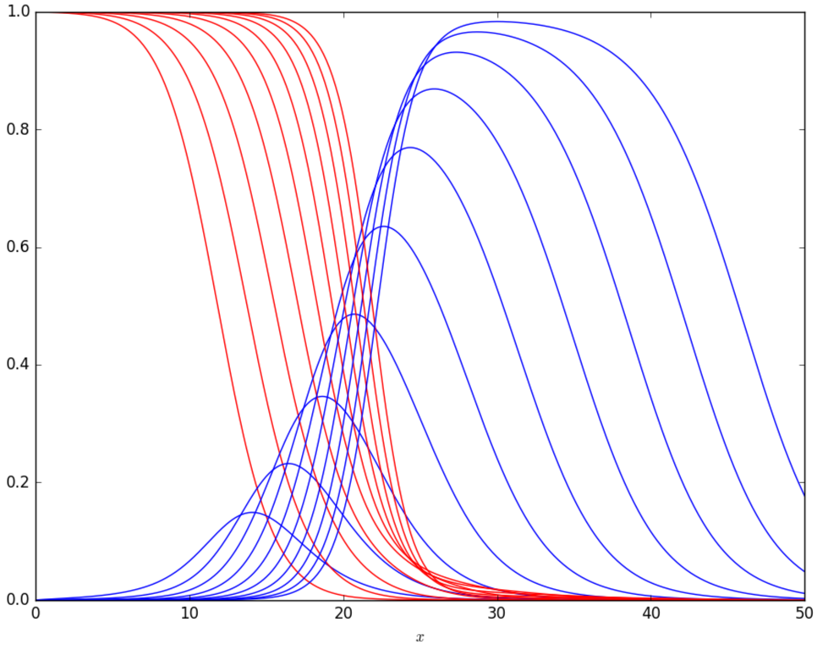

Figure 7: A numerical integration of the full system of equations, with a relatively low mutation rate. Red is population density, u, and blue is rescaled mean inbreeding coefficient, v. Lines are taken at equal time intervals. The parameters used were $\alpha = 4$, $\beta = 2$, $\gamma = 3$, $D = 1$.

All analysis of the model outlined in this document assumes that the carrying capacity is constant. This is not as unreasonable an assumption as it may seem at first glance. One may point out that the properties of the environment, such as the amount of food available, can be expected to fluctuate substantially over time, and will vary relative to the current or recent population density (high populations will outstrip the environment and leave less food for the next generation). However, it is important to remember that the carrying capacity is not simply the maximum population density that the environment can support *at that point in time*, but rather the maximum population that the environment can support *for an indefinite period of time*.

The other, non-constant properties of the environment, are assumed to be in equilibrium at the carrying capacity, and hence generally not relevant for small fluctuations, but for populations far from equilibrium it may be necessary to model these explicitly (see section 5.3 for examples).

The carrying capacity may be non-constant, due to long term changes in the environment, but it is sufficient to assume that the carrying capacity varies on a much longer timescale than the population density and inbreeding coefficient to justify approximating it as constant. This technique is known as *two-timing*, and gives results much more accurate than may be expected, even for timescales only a few times greater. (Strogatz, 2015)

The model also assumes, however, that the population will generally equilibrate to the carrying capacity, which is less well founded. Real populations, even those not threatened by significant changes in the environment, will often fluctuate dramatically without obvious cause. Though this is not observed in the model, nor in the Fisher-Kolmogorov equation, it is apparent in the logistic map, defined

$$u_{n+1} = ru_n(\kappa - u_n)$$

where u_n is the population density at the discrete timestep (or generation) n, and r and κ are the intrinsic replacement rate and carrying capacity. For such a simple model, it can display surprisingly complex behaviour. For low values of r, the population will equilibrate to the carrying capacity, as expected. For higher values, the population instead equilibrates to a steady, stable cycle of fluctuating above and below the carrying capacity. For higher values still, the population does not equilibrate at all, and instead fluctuates around the carrying capacity unpredictably and chaotically (May et al., 1976).

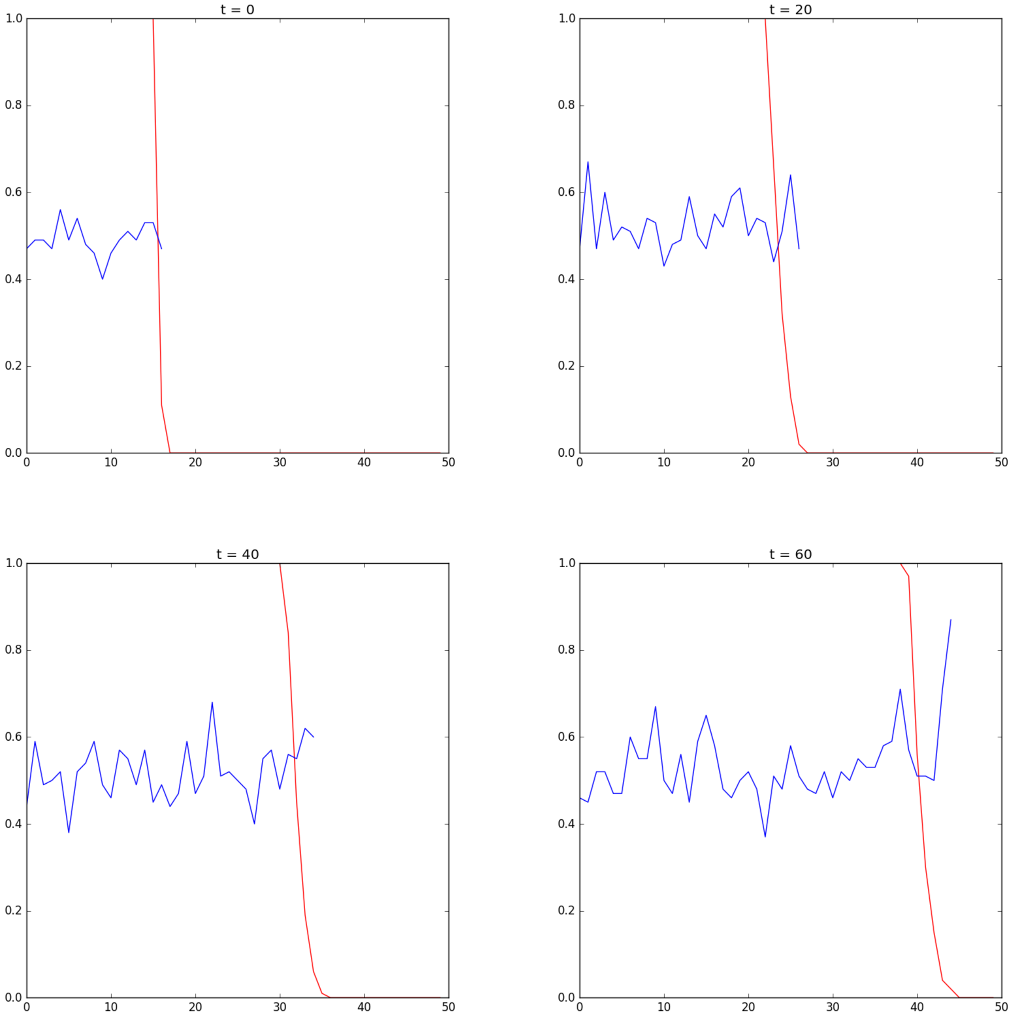

Figure 8: A travelling wave in the agent-based model. The horizontal axis is space. The red line is the population density, u, and the blue line is the mean inbreeding coefficient, f. The model had a migration rate of 0.7, breeding rate of 0.7, background death rate of 0.2, additional death rate per unit increase in f of 0.8, and a mutation rate of 0.02 per locus. All rates were probabilities per agent per timestep. The chromosomes had 100 loci each, and the carrying capacity was 100 for each subpopulation.

5.2 Applicability to Real Populations

5.2.1 Evidence

The verifiability of the model is greatly limited by the scarcity of high precision data of population density and genetic variation within endangered or otherwise limited populations. What data sets of natural populations do exist often have timesteps between measurements of as much as a year, and typically do not include detailed genetic information or accurate distributions in space.

In the absence of more reliable data, it is argued that the general agreement of the continuous model with both the results of the agent-based model and the existing theory and literature are sufficient for the model to be as well confirmed as possible at the present time.

Experimental data may be able to confirm or deny the results of the model. In particular, the common model organism *Drosophila melanogaster* satisfies most of the assumptions outlined in section 5.1, and lends itself well to laboratory experiments.

To investigate the critical value for recovery in the spatially homogeneous case, one could maintain small populations in enclosed environments, varying the initial population density and mean inbreeding coefficient, and measuring the long term dynamics of the population over the course of several generations. The carrying capacity could be held constant by providing a constant amount of food regardless of the local population.

To investigate travelling waves would require more sophisticated apparatus, but is still possible in theory. One possibility would be a long, thin enclosed tank of length substantially greater than the mean distance an individual will migrate over the course of its lifetime. Food would need to be dispensed throughout the length of the tank in order to maintain a constant carrying capacity. An initial population would then be released at one end of the tank, and it would be seen if, over the course of multiple generations, a travelling wave forms, and, if so, how fast it would travel.

5.2.2 Calculability

All variables and parameters in the model are at least theoretically measurable or calculable. The population density can be calculated by separating the environment into small sub-environments, measuring the census population in each sub-environment, and dividing by the area of the sub-environment. The census population can be estimated by, for instance, the mark-and-recapture method.

The mean inbreeding coefficient can be calculated as one minus the heterozygosity, and the heterozygosity can be measured by sequencing at least part of the genomes of several individuals in each sub-environment (as above). Thanks to next generation sequencing methods, even sequencing several hundred genomes is no longer prohibitively expensive.

Both above methods are vulnerable to the *modifiable areal unit problem*, meaning that the calculated quantity can vary dramatically depending on the size, shape and position of the sub-environments. To minimise the effects of this, it is advisable to use constant size, constant shape sub-environments, for instance a grid.

The carrying capacity can be observed directly if there are reasonably long term records of the population density of a healthy population of the species, in which case it is simply what value the population density equilibrates to over time. If such records are not available, it can be estimated based on the amount of food and space available in the environment and the amount used per individual. This is less reliable, and requires knowledge of what the minimum limiting factor is to population growth.

Diffusivity can be estimated based on how far individuals tend to migrate over the course of their lifetimes. This can be measured using, for instance, tracking collars.

The intrinsic replacement rate, r_0, can be estimated by monitoring the birth and death rates of genetically healthy individuals in isolation. The actual replacement rate, r, can be estimated by monitoring the same things in the actual environment.

If all of the above are known, along with the rate of change of the population density, then, in a region without a strong gradient of population density in space,

$$\begin{aligned}\frac{\partial u}{\partial t} &= ru(\kappa - u)\\ &= u(r_0 - Iv)(\kappa - u)\end{aligned}$$

so the coefficient I can be deduced. This would be very difficult to measure otherwise, but may be estimable from the number of lethal equivalents in the genome.

5.2.3 Consequences for Endangered Species

For conserving a population of a given endangered species, the aims are

- To identify whether or not the population will recover without intervention.
- If the population is outside the extinction vortex, to maximise its resilience to future stresses.
- If the population is within the extinction vortex, to intervene so as to allow it to recover.
- To do the above in the most time and resource efficient way possible.

The first thing that becomes apparent is that the environment, and the distribution of the population within the environment, are extremely important. Endangered species often have highly fragmented habitats, with subpopulations effectively isolated from one another.

Fully, or almost fully isolated subpopulations are more prone to extinction relative to well connected subpopulations because losses in population density cannot be offset by migration. Hence, removing or minimising barriers in the environment (e.g. providing bridges or fords across rivers) could go a long way towards improving the resilience of the population as a whole. Similarly, increasing the diffusivity of the environment (e.g. clearing small paths through the environment) will enable faster migration to counteract localised loss of population density.

The (dimensional) constants μ, I and J are intrinsic properties of the species and cannot be changed. By contrast, the intrinsic replacement rate, r, and the carrying capacity κ are mutable properties of the environment.

This ties in with the *r/K selection theory* (Pianka, 1970), wherein organisms can be classified as *r-type* or *K-type*. r-type organisms are generally opportunistic, being quick to recover and to exploit new environments, whereas K-type organisms are generally more conservative, relying on having long-term dominance of the environment through larger populations. These correspond to maximising the intrinsic replacement rate r, or the carrying capacity κ, respectively.

The theory has fallen out of favour in recent times, as it has become apparent that most species do not fit neatly into either category, or even necessarily on a spectrum between the two extremes. However, the concepts can still be of use: for small, endangered populations, r-type behaviour is optimal for quickly recovering, whereas for large, stable populations, K-type behaviour is optimal for maximising resilience.

The intrinsic replacement rate can be maximised by maximising the breeding rate and minimising the death rate. Physically, this could include

- Breeding in captivity and releasing offspring back into the environment.
- Reducing predation, by excluding predators from the environment or, in extreme cases, culling.
- Making reproduction easier, for instance by providing nest boxes.

By contrast, maximising the carrying capacity would constitute maximising whatever is the limiting factor in population growth. Usually, this will mean promoting the populations of food plants or prey, as appropriate, but some species may be limited by the availability of territory, or the climate.

Notably, the r-type strategies are generally resource intensive and temporary measures, whereas the K-type strategies are generally easier and more permanent. This highlights the fact that early intervention and a preventative approach are far more efficient and effective than waiting for species to become critically endangered before intervening.

It is important also to remember that r-type strategies will do nothing to help a population at or close to carrying capacity, as the limiting factor is a lack of resources, and similarly K-type strategies will do nothing to help an endangered population, as there are plenty of resources to go around.

The model does highlight some non-obvious strategies for conserving endangered species. For example, if the population is highly spread out in many subpopulations, each with low population density, it may be best to actively concentrate the population into a smaller area. At first glance, this may seem absurd - if the population is already low, then forcing them together so that they have to compete for resources as well is only going to reduce the population further - but the genetic benefits of having a higher population density are potentially great enough to overcome this.

Another strategy is to introduce a few breeding individuals with a highly different genome into an inbred population to provide a boost to the local mean genetic variation. This is known as *genetic rescue*, and has been successful in enabling long term recovery of, for instance, a Scandinavian wolf population (Åkesson et al., 2016).

One important thing to note from the spatially homogeneous results, as shown in figure 1, is that the rate of change of population density is not necessarily a good indication of the long term survival of the population. In the top-right of the phase plane, the population density is decreasing precipitously, but still ultimately recovers; whereas, towards the bottom left, the population density may be increasing but the population is still ultimately headed for extinction.

This result is especially important because the IUCN red list (IUCN, 2001), which is usually used to classify how endangered a given species is, has in its criteria a strong focus on the rate of change of population, but has no mention of inbreeding or genetic effects whatsoever. Some criteria are vague enough that an argument could be put forward of “inferred or suspected population size reduction” based on genetic effects, however.

Given how strict many of the other criteria are, this presents a serious risk of misclassification of endangered species, which could carry severe consequences.

5.2.4 Consequences for Invasive Species

For eradicating a population of a given invasive species, the aims are

- To prevent the population from spreading into new environments.
- To drive the current population extinct.
- To do the above in the most time and resource efficient way possible.

This can be seen as the reverse of the endangered species case, and hence in many ways the appropriate strategies are the reverse of the above.

Preventing spread can be accomplished by reducing the replacement rate, reducing the diffusivity or introducing barriers to migration, and, in particular, by having a highly inbred population near the leading edge. Care must be taken with measures that affect the environment itself, as these may also negatively affect native species, which may themselves be endangered due to competition with the invasive species. The introduction of highly inbred subpopulations near the leading edge is therefore an especially useful strategy, as it does not directly affect other species nearby.

Eradicating well established populations is likely to be difficult, but the extinction vortex offers a good target and a means of preventing recovery. The goal is then to minimise population density and to maximise the mean inbreeding coefficient.

Unfortunately, increasing inbreeding simply by releasing inbred individuals, as a reverse of genetic rescue, is unlikely to be effective. Unless the inbred individuals form a majority, they will simply be outcompeted by genetically healthy individuals within a few generations. Rather, inbreeding will have to be increased simply by maintaining a low population density.

For a healthy population close to carrying capacity, a three step approach is likely to be most effective:

- Firstly, minimise the diffusivity and carrying capacity to increase the vulnerability of the population. Where possible, this should be done with the minimum collateral harm to native species.
- Secondly, the population density should be reduced, either by culling, or by waiting for a natural population drop due to environmental factors.
- Finally, when the population is at its most vulnerable, the replacement rate should be minimised and the population density kept low until the extinction vortex takes hold.

The first step is optional, but potentially useful. The last step is especially difficult, because even a single subpopulation recovering can result in the recolonisation of the entire environment. If it is possible to fully divide the environment into discrete pieces, with no migration between them, then focusing on one piece at a time is likely to be easiest.

The naive approach to eradicating invasive species (namely, just killing as many as possible by any means) is typically ineffective, because cullings and intentionally introduced diseases will generally not kill enough, or maintain a population loss for long enough, for the population to fall into the extinction vortex.

Given how difficult and expensive the above strategies are likely to be, by far the most efficient way to counteract invasive species is simply to prevent them from invading in the first place. Hence, preventative measures should be a high priority.

5.3 Future Research

The model works fairly well as a closed system, and appears to contain all important evolutionary forces for modelling the extinction vortex. However, as outlined in section 5.1, there are many situations in which it is inapplicable. Far from being limitations, these offer interesting possible avenues for future research.

5.3.1 Generation Time

The *generation time* of a species is the average time between an individual being born and it reaching sexual maturity. The model implicitly assumes that the generation time is zero; that is, an increase in population density due to births contributes instantaneously to more births.

Incorporating this would require explicitly separating the replacement rate into separate birth and death rates, and having the increase in population due to births be proportional not to the total population density, but to the population density of sexually mature individuals. There are several approaches that could be used to model this.

Perhaps the most obvious would be to have separate mature and immature subpopulations. The limitation of population increase due to approaching the carrying capacity, and the diffusion, would then be proportional to the sum of mature and immature subpopulations, but the rate of increase of population would only be proportional to the mature subpopulation.

A simple alternative would be to have the population increase due to births be proportional not to the current population density but to the population density some amount of time ago. This assumes that the proportion of the population that is sexually mature is approximately constant. Though less accurate, this captures the fundamental behaviour without requiring the addition of another variable, however it would make the analysis substantially more challenging.

An even more radical simplification would be to change the model from being a set of differential equations to a set of discrete mappings; that is, to discretise the model in time, not entirely unlike the approximation taken in section 3.1. This would correspond to having separate, but potentially overlapping, generations; a common feature of real animals. The equations would then have a form along the lines of

$$\begin{aligned}
u_{n+1} &= \frac{\partial^2 u_n}{\partial x^2} + (r_b u_{n-G} - (r_d + Iv)u_n)(1 - u_n) \\
v_{n+1} &= \frac{\partial^2 v_n}{\partial x^2} + \gamma v_{n-G}(1 - v_{n-G})(1 - u_n) - u_n v_{n-G}
\end{aligned}$$

where u_n is the population density at the nth timestep, G is the generation time in timesteps, r_b is the intrinsic birth rate, and r_d is the intrinsic death rate. Note that all terms involving breeding are proportional to the population density or inbreeding coefficient at the timestep at which the currently breeding individuals were born, rather than the current timestep.

Any of the above approximations has the potential to dramatically alter the dynamics of the system. The first approximation introduces a third dimension in phase space; the second makes the equations nonautonomous (explicitly depending on time), and the third would make phase space discontinuous. These introduce the possibility of chaotic solutions, which are precluded by the equations in their current form, due to the fact that in two dimensions the trajectories in phase space cannot loop back on themselves and must instead approach some orderly behaviour.

This would serve to make the behaviour of the system much less predictable, but in many ways more realistic. For example, there would be the possibility of serious population booms (jumping above the carrying capacity) and crashes (potentially landing it into the extinction vortex with little warning).

Given the relative mathematical ease of discrete mappings, and the potential gains in terms of interesting behaviour of the system, this represents an especially promising modification to the system.

5.3.2 Heterogeneous Environments

The analysis put forward in this document has been limited to an environment with constant, spatially homogeneous parameters; the carrying capacity, diffusivity, and intrinsic replacement rate are all the same across the entire environment. In reality, it is likely that some parts of the environment will be more hospitable than others, and all species have limits to what habitats they can tolerate, so all will have certain environments which they cannot colonise.

Shigesada et al. (1986) analysed the Fisher-Kolmogorov equation in a periodic, heterogeneous environment, with the diffusivity and replacement rate varying in space. They found that travelling waves could pass even through regions with negative replacement rate, provided that it is above a given critical value.

A similar approach taken towards this model would be valuable in investigating the limits in which colonisation is still possible, and how genetic effects are involved.

5.3.3 Natural Selection

A major evolutionary force that is notably absent from the model is natural selection. Though this generally operates on a longer timespan than is considered for the extinction vortex, it could become relevant for modelling long-term recovery, or in the case of extremely strong selective pressures.

Natural selection acts to decrease the genetic variation of the population, and hence increase the inbreeding coefficient, by purging the population of deleterious alleles - yet the purged population would be healthier and hence would not have the decreased replacement rate characteristic of the extinction vortex.

The inclusion of natural selection in the model would almost certainly necessitate the addition of a new variable to measure fitness. Qualitatively, the expected dynamics of the system would be that

- The population density would generally increase up to the carrying capacity, as in the current model. The replacement rate would be decreased by low fitness, rather than directly by a high inbreeding coefficient.
- The inbreeding coefficient would increase at low populations (due to inbreeding) and at high rate of change of fitness (due to deleterious alleles being purged) and decrease with mutation.
- The fitness would generally increase up to a constant maximum fitness. It would be decreased by low populations, due to the effects of inbreeding and genetic drift.

Intuitively, one would expect this system to have broadly similar dynamics to the basic model, with stable states corresponding to a healthy population (maximum population density, minimum inbreeding coefficient, maximum fitness) and to extinction (zero population density, maximum inbreeding coefficient, minimum fitness). Without an explicit form of the equations it is impossible to give any more detailed analysis.

5.3.4 Predation

Predation is already included in the model in a very crude form as part of the replacement rate. This assumes the only effect of predation is an additional constant death rate. This may be a reasonable approximation if the predator population is generally constant (for example, a generalist predator that has many other sources of food), but in the case of a predator that feeds only on the species studied, the predator population - and hence the decrease in replacement rate - will vary with the prey population.

The above is equally applicable to populations of parasites. If the effects of the other population are inverted (i.e. the other population increases the replacement rate, rather than decreasing it) then it can be used to model the effects of a mutualistic relationship with another species.

A synthesis of the Fisher-Kolmogorov equation and Lotka-Volterra equations was studied by Cruywagen et al. (1996) in the context of colonisation by genetically modified organisms, including the effects of spatially heterogeneous environments. However, this analysis assumed a constant replacement rate, and considered the two populations to be in direct competition, rather than predating on one another.

Modifications of the equations of this sort may be especially valuable in studying highly interdependent species that are both endangered.

5.3.5 Disease

Besides simply competing for resources, another way in which high population density can lower the replacement rate and hence limit population growth is through the spread of disease. In the model, population growth is linear with respect to competition for resources; the $\kappa - u$ term can be loosely interpreted as the amount of resources available to the population. The rate of change of population density is linearly proportional to this.

By contrast, limits to population growth imposed by disease will be extremely nonlinear, and not straightforward to model. One possibility would be to use a variation of the SIR equations. These separate the population into three subpopulations: susceptible, infected, and recovered. In their original formulation, these equations assume a constant census population, and neglect any spatial heterogeneity. The assumption of constant population also implies that the disease is never fatal.

The system of equations would likely be similar, with the additional condition that

$$u = S + I + R$$

where S, I and R are the population densities of the susceptible, infected and recovered subpopulations respectively. The u equation would also include an additional death rate proportional to the infected population density I. There would be three equations required to describe the changes in S, I and R with respect to time, noting that offspring will be born susceptible. The exact form of these hypothetical equations is beyond the scope of this document, and the resultant system would be unlikely to produce many precise analytic results without the use of substantial simplifications.

Even as a purely computational model, the results could be interesting, and the dynamics of the equations provide some interesting possibilities. One would expect a general cycle of epidemic and recovery, with the possibility of extinction if an especially severe epidemic reduces the population to below the brink of the extinction vortex.

Furthermore, the known existence of travelling waves in population, inbreeding, and infection raises interesting prospects of how the different waves would interact, which has serious consequences for the proposed use of deliberately introduced diseases to control invasive species populations. If the goal is to prevent spread, then unless the infection wavespeed is greater than the population wavespeed, it will fail to prevent colonisation. If the goal is complete eradication, then unless the population density is zero in the wake of the infection wave, it will likely be followed by a resistant population recolonisation wave.

5.3.6 Migration

The model assumes that migration is entirely random, with no preferred direction. In reality, many species migrate with some preferred direction. If the entire population is moving in one direction at a constant speed, without any other migration, then this is equivalent to being static in a moving frame of reference and simplifies to the spatially homogeneous system considered in section 4.1.

For aquatic organisms, migration is often dominated by ocean currents. If the velocity field of the body of water is known, then this can be approximated by replacing the time derivatives in the equations with material derivatives, defined for some general function F,

$$\frac{DF}{Dt} = \frac{\partial F}{\partial t} + \underline{u}.\nabla F$$

where $\underline{u}$ is the (vector) velocity field, not to be confused with u, which is the (scalar) population density. In comparison to the effects of the ocean currents, the diffusive terms can probably be neglected, giving a system of equations of the form

$$\frac{Du}{Dt} = \alpha u(1-u)(1-\beta v)$$
$$\frac{Dv}{Dt} = \gamma v(1-v)(1-u) - uv$$

The effects of this change depend on the properties of the velocity field. In areas of divergence (where the flow is spreading out, or accelerating), if present, there will be a net loss of population density, making the local population, and potentially also the population downstream, more prone to extinction and less able to recover from population losses. The opposite would be true in areas of convergence (where the flow is coming together or decelerating).

Provided that the velocity field is non-zero, there would probably still be travelling waves, but their wavespeed would likely depend on the local velocity field, and they would only be able to spread in the same direction as the velocity field itself. Furthermore, there would be the possibility of waves of extinction, as well as of colonisation.

5.3.7 Turing Instability

The model constitutes a reaction-diffusion system. Reaction-diffusion systems have the general form

$$\frac{\partial u}{\partial t} = \nabla^2 u + F(u,v)$$
$$\frac{\partial v}{\partial t} = D\nabla^2 v + G(u,v)$$

where u and v are concentrations or densities of some sort, D is the differential diffusivity (the ratio between the diffusivities of u and of v), and F and G are known functions.

For $D \neq 1$, it is possible for the effects of the diffusion and reaction terms to balance in such a way that heterogeneous distributions in space of u and v can be stable in time. This is known as *Turing instability.*

Turing instability is especially notable because systems featuring diffusion typically act to smooth out variation and ultimately result in spatially homogeneous distributions. For example, the long-term behaviour of the Fisher-Kolmogorov equation, for non-zero initial conditions, is always $u = 1$, corresponding to the population uniformly staying at carrying capacity.

For a general reaction-diffusion system to exhibit Turing instability (Murray, 2001):

1. There must exist an equilibrium point, (u_0, v_0) such that $F(u_0, v_0) = G(u_0, v_0) = 0$.
2. The equilibrium point must be stable if u and v are spatially homogeneous.
3. The equilibrium point must be unstable if u and v are spatially heterogeneous.
4. The differential diffusivity, D, must not be equal to one.

It was shown in section 4.1 that the spatially homogeneous case has two stable equilibrium points, $(1, 0)$ and $(0, 1)$, and in section 4.2 that these points are both unstable in the spatially heterogeneous case with respect to travelling waves. Further analysis is needed to prove definitively whether or not the model can display Turing instability, but preliminarily it is certainly possible.

If the model were to show Turing instability, this would imply the possibility of spatial heterogeneity arising spontaneously, including unhealthy or uninhabited patches within the environment that may be steady over time. If this is true, then it could go some way towards explaining the way in which real populations tend to coalesce into discrete subpopulations separated by relatively uninhabited regions. What is especially notable is that this does not even require any heterogeneity in the environment.

6 Conclusion

The model presented in this document has a good backing in biological theory. In terms of physical processes, it includes all of, and only, birth, death, migration, inbreeding, and mutation. The terms of the equations intuitively correspond to real processes. The coefficients can be written in terms of meaningful biological constants.

In the absence of spatial effects, the model predicts that the population will tend towards either recovery or extinction. There is a clearly identifiable boundary to the extinction vortex. Given accurate measurements of the physical quantities involved, the model can make predictions of whether the population will recover or not without intervention.

With spatial effects, the model predicts travelling waves corresponding to colonisation events to occur under some circumstances. These are characterised by a smooth front in population density, with a small peak of greater than equilibrium inbreeding coefficient at the leading edge. The wavespeed can be predicted based on physical parameters.

The results are generally confirmed by an agent-based model, operating under exactly the same set of physical processes.

Both the spatially homogeneous and spatially heterogeneous cases offer useful, meaningful, and detailed predictions about conservation of endangered species and eradication of invasive species, including some which run counter to the prevailing techniques and classifications currently used.

The applicability to real populations is extremely difficult to judge. Qualitatively, the behaviour appears to be generally accurate, but the quantitative accuracy cannot be judged due to insufficient data. If confirmed, many promising modifications to the model are available that could improve its accuracy, or extend it to cover situations specific to certain species.

In conclusion, the model presented here is successful in its goal of modelling the extinction vortex, as well as many related phenomena in population dynamics. Though usable as it is, it warrants further research to achieve its true potential.

7 Acknowledgements

Sincere and heartfelt thanks are offered to Dr Mark Blyth of the UEA school of mathematics and professor Cock van Oosterhout of the UEA school of environmental sciences for their guidance over the course of this project, without which it would not have been possible.

All numerical work was performed in the Python programming language, using the NumPy, PyPlot and SymPy libraries. Source code is available on request.

References

Mikael Åkesson, Olof Liberg, Håkan Sand, Petter Wabakken, Staffan Bensch, and Øystein Flagstad. Genetic rescue in a severely inbred wolf population. *Molecular Ecology*, 25(19):4745–4756, 2016.

Ward Cheney and David Kincaid. *Numerical Mathematics and Computing*. Nelson Education, 2012.

Gerhard C Cruywagen, P Kareiva, MA Lewis, and JD Murray. Comptetition in a spatially heterogeneous environment: Modelling the risk of spread of a genetically engineered population. *theoretical population biology*, 49(1):1–38, 1996.

Jared Diamond. *Collapse: How societies choose to fail or succeed*. Penguin, 2005.

Ronald Aylmer Fisher. The wave of advance of advantageous genes. *Annals of eugenics*, 7(4): 355–369, 1937.

Richard Frankham, David A Briscoe, and Jonathan D Ballou. *Introduction to conservation genetics*. Cambridge university press, 2002.

John Gribbin. *Deep simplicity: Chaos, Complexity and the Emergence of Life*. Penguin UK, 2005.

Godfrey Harold Hardy. A mathematician's apology, 1941.

IUCN. Iucn categories and criteria (version 3.1), 2001. URL http://www.iucnredlist.org/static/categories_criteria_3_1. Accessed 22/04/17.

Andrei Nikolaevitch Kolmogorov, Ivan G Petrovskii, and Nikolaï S Piskunov. A study of the equation of diffusion with increase in the quantity of matter, and its application to a biological problem. *Bjul. Moskovskogo Gos. Univ*, 1(7):1–26, 1937.

Yuri Lazebnik. Can a biologist fix a radio? — or, what i learned while studying apoptosis. *Biochemistry (Moscow)*, 69(12):1403–1406, 2004.

Michael Lynch, Bruce Walsh, et al. *Genetics and analysis of quantitative traits*, volume 1. Sinauer Sunderland, MA, 1998.

Robert M May et al. Simple mathematical models with very complicated dynamics. *Nature*, 261 (5560):459–467, 1976.

James D Murray. *Mathematical Biology*. Springer-Verlag New York Incorporated, 2001.

Eric R Pianka. On r-and k-selection. *The American Naturalist*, 104(940):592–597, 1970.

Eric S Raymond. The utility of mathematics, 2005. URL `http://www.catb.org/~esr/writings/utility-of-math/`. Accessed 18/04/17.

Nanako Shigesada, Kohkichi Kawasaki, and Ei Teramoto. Traveling [sic] periodic waves in heterogeneous environments. *Theoretical Population Biology*, 30(1):143–160, 1986.

Steven Strogatz. *Nonlinear dynamics and chaos: with applications to physics, biology, chemistry, and engineering*. Westview Press, 2nd edition, 2015.