PhD Macroeconomics 1
2. Phase Diagrams

Karl Whelan

School of Economics, UCD

Autumn 2023

Part I

Introduction to Phase Diagrams

Illustrating Dynamics in a Two Variable Model

- While there are various tricks that work for solving specific types of nonlinear differential equations, there is no general solution technique that works.
- And linearisation around a particular point is not much help if we want to understand the global dynamics of the model: Often the dynamics of a model around one point can be quite different to its dynamics around a different point.
- Phase diagrams are a popular technique that allow us to see the joint dynamics of two variables. Many models in economics can be boiled down to two variables so this method is widely used.
- Let's start with a very general example

$$\begin{aligned} \dot{x} &= g(x, y) \quad \frac{\partial g}{\partial x} < 0 \quad \frac{\partial g}{\partial y} < 0 \\ \dot{y} &= f(x, y) \quad \frac{\partial f}{\partial x} < 0 \quad \frac{\partial f}{\partial y} > 0 \end{aligned}$$

where g and f are continuous differentiable functions.

Deriving the Dynamics of x and y

- Now let's think about the set of points that correspond to $\dot{x} = 0$. Start off at one such point. Since g is a decreasing function of both x and y, then if we increase x, then for $g(x, y)$ to remain equal to zero, we would have to decrease y. This means the set of points corresponding to $\dot{x} = 0$ is characterised as a downward-sloping line in $y - x$ space. We have drawn this here as a straight line but in general this won't be the case.
- Away from the $\dot{x} = 0$ line, what are the dynamics? Well the points above the line feature higher values of y than the points on the line. Since $\frac{\partial g}{\partial y} < 0$, this means that $\dot{x} < 0$, meaning x is declining. Phase diagrams mark this with an arrow pointing left. Points under the line correspond to x increasing so we draw an arrow to the right.
- Similar logic tells you that the $\dot{y} = 0$ line is upward-sloping and that, at points above it, y is increasing while it is decreasing at points below the $\dot{y} = 0$ line.
- This divides up the (x, y) space into four quadrants and we know in which directions that x and y are going in each of these quadrants.
- This, my friends, is a **phase diagram**

The $\frac{dx}{dt} = 0$ Locus

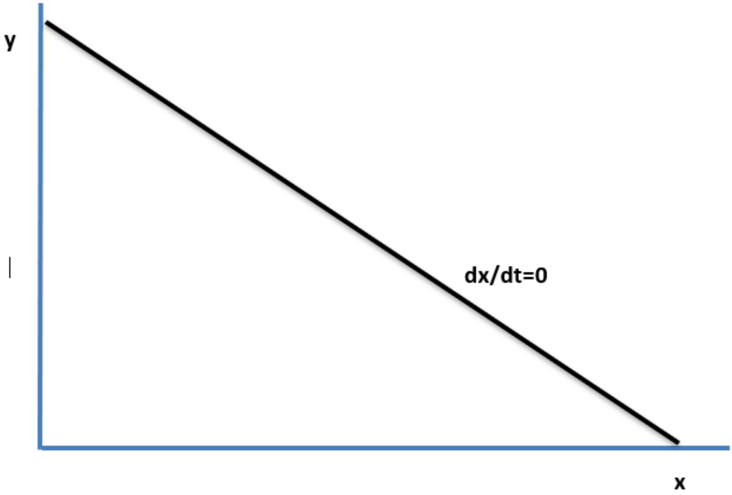

Dynamics of x

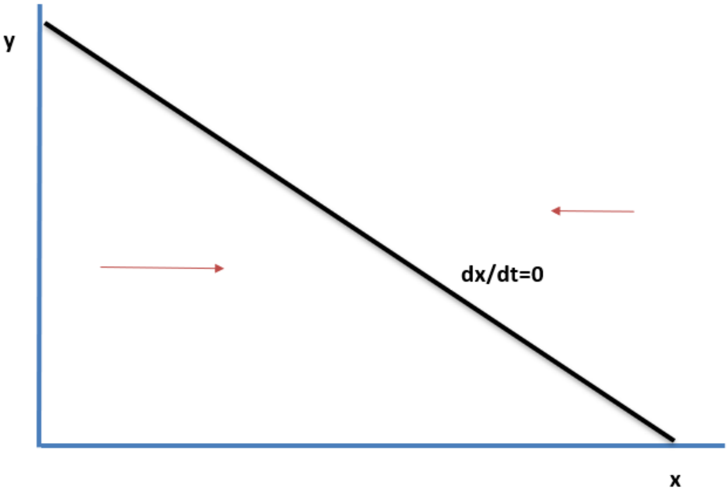

The $\frac{dy}{dt} = 0$ Locus and Dynamics of y

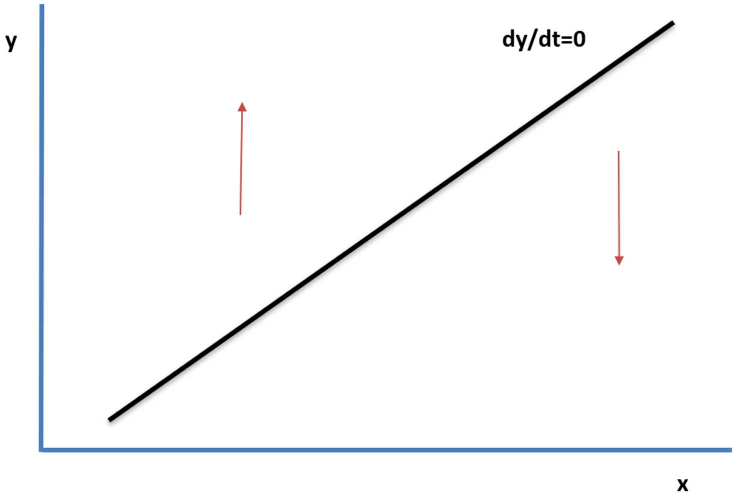

Full System Dynamics

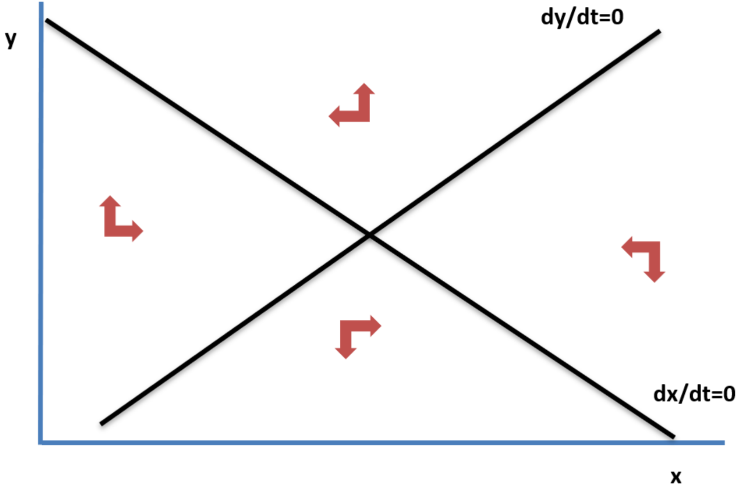

Convergence to Equilibrium?

- The equilibrium of this model is where the $\dot{y} = 0$ line intersects the $\dot{x} = 0$ line.
- Will the model generally tend to head for this equilibrium no matter where the system starts?
- We can see from the previous picture that, in this case, it appears the answer is no. In the quadrants directly above and below the equilibrium, the dynamics take the variables away from equilibrium.
- We would need more detail to figure out if there is any chance of the model converging to equilibrium. However, many macroeconomic models have the feature that the variables must take values along a unique so-called **saddle path** that propels the variables steadily towards the equilibrium value.
- Note that g and f are continuous functions so the dynamics cannot be "jumpy". So, for example, if we go from a part of the phase space where y is increasing and then pass over into a part where it is decreasing, then the rate at which it is increasing must gradually approach zero from a positive direction and then gradually move away from zero in a negative direction.
- To illustrate, I have drawn an example of what a saddle path might look like for this model.

A Convergent Saddle Path

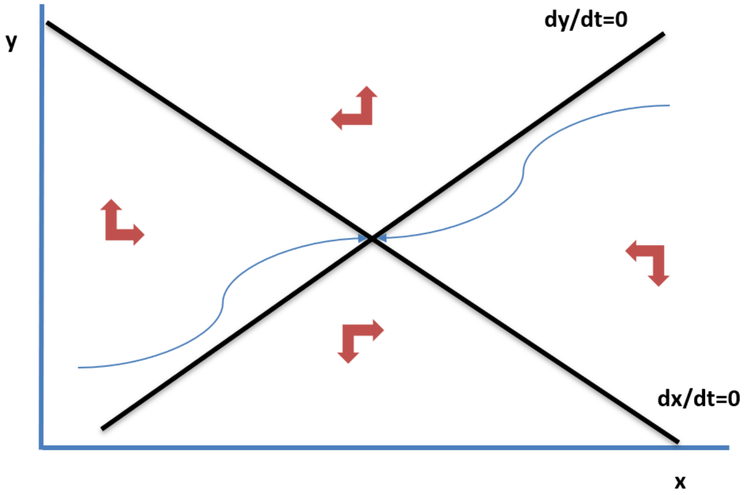

Part II

Example: The SIR Epidemiology Model

Introducing the SIR Model

- This is the classic model of epidemiology which has been the basis for many analyses of the COVID19 pandemic.
- People can be in one of three states:
  - S_t is the number of people that are susceptible to being infected.
  - I_t is the number of people that are currently infected.
  - R_t is the number of people that have recovered.
- The model's dynamics are as follows

$$\begin{aligned}
\dot{S}_t &= -\beta I_t S_t \\
\dot{I}_t &= \beta I_t S_t - \gamma I_t \\
R_t &= 1 - S_t - I_t
\end{aligned}$$

- The discussion of the model is mainly borrowed from lecture notes by Ben Moll of LSE. So are the Matlab programmes generating the graphs.

Assumptions of the Basic SIR Model

1. The number of people getting the infection is proportional to both S_t and I_t. This effectively assumes people are randomly meeting and some fraction of the meetings between susceptible and infected result in new infections. The parameter β determines the rate of new infections: It depends on both the intensity with which people meet others and the infectiousness of the disease.
2. Everyone either has the infection now, has had it in the past and recovered (and is now immune) or is susceptible to getting it now or in the future. There are no people who are immune to the infection at the start of the outbreak.
3. People recover from being infected at rate γ. There is no distinction between being infected and being infectious (e.g. a period during which the person can pass on the disease but does not feel infected or a period in which the person feels sick but is no longer infectious.)
4. We have normalised the total population to one.

We will look at loosening the first and second assumptions after we have worked through the dynamics of the basic model.

A Phase Diagram for the SIR Model

- We are going to make a phase diagram in S, I space. Once we know what these variables are doing, we automatically get R_t.
- To make a phase diagram, we need to know the conditions under which our two variables are increasing or decreasing.
- Let's re-express the dynamics of S_t and I_t in growth rate terms:

$$\begin{aligned} \frac{\dot{S}_t}{S_t} &= -\beta I_t \\ \frac{\dot{I}_t}{I_t} &= \beta S_t - \gamma \end{aligned}$$

- S_t is always decreasing with the pace depending positively on the number of infected.
- I_t is unchanged when $S_t = S^* = \frac{\gamma}{\beta}$.
- Since the growth rate of I_t depends positively on S_t, we know the number of infections is growing if $S_t > S^*$ and is falling if $S_t < S^*$.
- That's enough information to make the arrows for a phase diagram.

Ben Moll's Phase Diagram Arrows

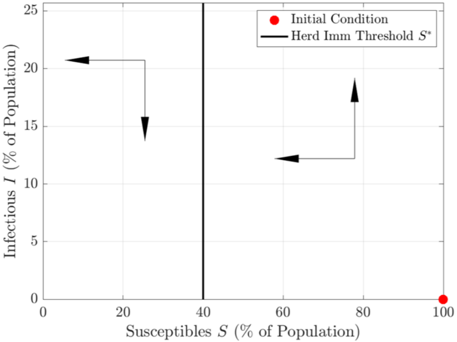

Path Starting from 100% Susceptible Looks Like This

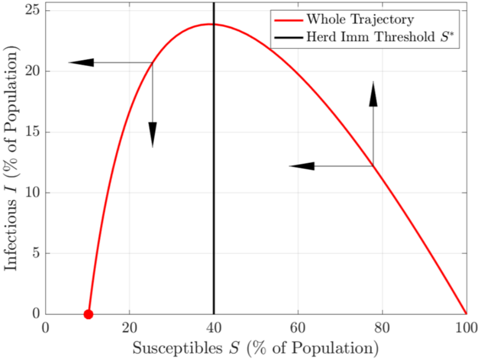

Some (Sort of Confusing) Terminology

- Notice from Ben Moll's graph that he labels $S^* = \frac{\gamma}{\beta}$ (equal to 0.4 in his implementation) as "**the herd immunity threshold**" level of susceptibility.
- A few comments on "herd immunity" in this model.
  - In many models, variables labelled with a star represent the long-run levels that the model ends up at. There is nothing in this model that makes the herd immunity level of susceptibility be the long-run outcome. You can see from Ben Moll's chart that when you start at $S_0 = 1$, the level of susceptibility way over-shoots S^* to end up at a much lower level.
  - S^* is only important because once you reach that point, you know that the number of infections will steadily decrease. How long it takes infections to go to zero depends on how many infections you have by the time you arrive at the herd immunity threshold.
- Epidemiologists often refer to $\frac{\beta}{\gamma}$ (the inverse of S^*) as "**the basic reproductive number**" and sometimes call this R-zero. I'll denote this as $\mathcal{R}_0$ to avoid confusion with our existing R variable. For each infected person, $\mathcal{R}_0$ measures the ratio of people being newly infected to people recovering. If $\mathcal{R}_0 = 2.5$, then you only have a falling infection rate when you reach $S_t = 0.4$.

An Analytical Solution

- One useful result is that the derivative of log of a variable gives its instantaneous growth rate. In other words,

$$\frac{d \log X_t}{dt} = \frac{d \log X_t}{dX} \frac{dX_t}{dt} = \frac{\dot{X}_t}{X_t}$$

- Looking back at the model, we have

$$\begin{aligned} \frac{\dot{S}_t}{S_t} &= -\beta I_t \\ \dot{R}_t &= -\dot{S}_t = \dot{I}_t = \gamma I_t \end{aligned}$$

- This means that

$$\frac{dR_t}{dt} = \dot{R}_t = -\frac{\gamma}{\beta} \frac{\dot{S}_t}{S_t} = -\mathcal{R}_0 \frac{d \log S_t}{dt}$$

- Integrating from 0 to t, we get

$$\int_0^t \frac{dR_k}{dk} dk = -\mathcal{R}_0 \int_0^t \frac{d \log S_k}{dk} dk \Longrightarrow R_t - R_0 = -\mathcal{R}_0 \left(\log \left(\frac{S_t}{S_0} \right) \right)$$

An Analytical Solution

- We have derived
$$R_t - R_0 = -\mathcal{R}_0 \left(\log \left(\frac{S_t}{S_0} \right) \right)$$
- This can be re-written as
$$1 - I_t - S_t - R_0 = -\mathcal{R}_0 \left(\log \left(\frac{S_t}{S_0} \right) \right)$$
- Or also as
$$I_t = 1 - R_0 - S_t + \mathcal{R}_0 \left(\log \left(\frac{S_t}{S_0} \right) \right)$$
- This gives the path for infections at all time. Note the peak of infections occurs when $S_t = S^* = \frac{1}{\mathcal{R}_0}$ so the peak level is given by
$$I_t^{\max} = 1 - R_0 - \frac{1}{\mathcal{R}_0} - \mathcal{R}_0 \left(\log \left(\mathcal{R}_0 S_0 \right) \right)$$
- One can also drive an implicit formula for the amount of people still susceptible after the epidemic is over, i.e. when $I_t = 0$.

Starting With Lower Numbers Susceptible

- Ben Moll's graph assumed the epidemic starts with almost no infections $I_0 \approx 0$ and almost everyone susceptible so at the start, we have $S_0 \approx 1, R_0 \approx 0$.
- We can vary these assumptions to get a greater sense of the model's dynamics.
- For example, we could continue to assume that everyone starts out either susceptible or infected, but raise the number initially infected to various levels.
- The program `SIR.m` has code to make phase diagrams for the model using simple finite-difference approximations to the derivatives. It varies the initial amount infected from almost none to just over 0.6. As with Ben Moll's programme, herd immunity is set at 0.4.
- The code is on the next page and the phase diagram is on the next page.
- It shows, unsurprisingly, that the more people that are initially infected, the higher will be the eventual number infected. That said, despite large differences in the number infected (ranging from $I_0 = 0.001$ to $I_0 = 0.601$), the differences in final number infected are fairly modest.
- The programme uses Matlab's `quiver` command which draws arrows showing the direction the system is going. The larger the arrows, the faster the speed.

SIR.m

```
t     = 400;    % Length of simulation
dt    = 0.5;    % Time step

% Initialising Variables
S = zeros(1,t);
I = zeros(1,t);
R = zeros(1,t);

% Epidemiological Parameters, taken from Ben Moll's programme
Rnot     = 2.5;            % Number of transmissions of an infectious (I) person
Tinf     = 7;              % Duration of infectious period
beta     = Rnot/Tinf;      % Transmission rate of an infectious (I) person
gamma      = 1/Tinf;       % Exit rate from I state into recovery (R)
S_herd = 1./Rnot;
R_herd = (1-1./Rnot);

% Simulating Various Initial Conditions
for Iinit = 0.001:0.1:0.601
    disp(Iinit);
I(1) = Iinit;
S(1) = 1-Iinit;
R(1) = 1 - S(1) - I(1);

for n=2:t
        S(n) = S(n-1) - dt*beta*I(n-1)*S(n-1) ;
        I(n) = I(n-1) + dt*(beta*S(n-1)*I(n-1) - gamma*I(n-1));
        R(n) = 1 - I(n) - S(n);
end; %for n

dI = diff(I);
dS = diff(S);

   quiver(S(1:t-1),I(1:t-1),dS,dI)
   xlabel('Susceptibles $S$ (\% of Population)','Fontsize',20)
   ylabel('Infectious $I$ (\% of Population)','Fontsize',20)
   line([S_herd S_herd],[0 1],'Color','k','LineWidth',2)
   set(gca,'FontSize',15);
   hold on
end; %for Iinit
```

Varying the Initial Amount Infected

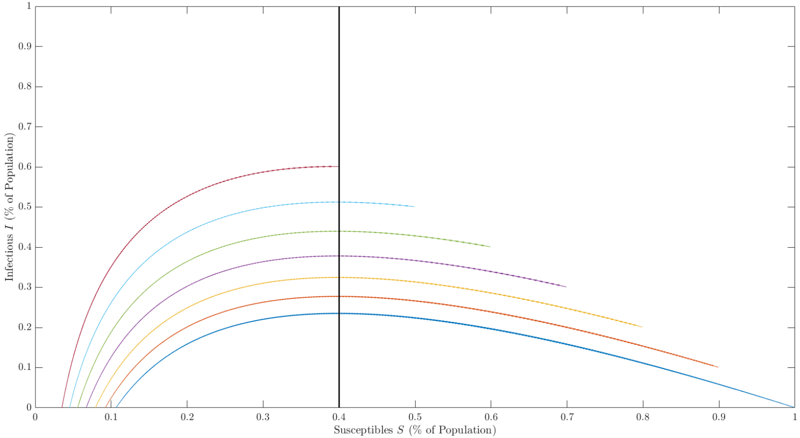

What If Not Everyone Is Susceptible?

- In the program `SIR.m`, a lower level of initial susceptibility came from a large number of people being infected at the start of the epidemic.
- A more likely reason to start out with lower susceptibility is that some people are immune to catching the infection, either because they have had it before or been vaccinated or simply aren't genetically susceptible.
- We can model this situation by having a higher value of R_0, placing those who are immune in the "recovered" category but still starting with a minimal amount of infected people. This is done in the programme `SIRImmune.m`
- The phase diagram this produces are shown on the next page. It illustrates the role of herd immunity.
- The lower the initial number of susceptible people, the weaker the epidemic is, both in terms of peak infections and eventual numbers who get infected.
- If you can get to herd immunity (in this case $S_t = 0.4$) and you have very low infection numbers, then this outbreak will quickly disappear. But if you arrive at the herd immunity threshold and there are still lots of people infected, then having reached the level doesn't help much – you are still going to end up with most of initially susceptible people getting the infection.

Allowing Some People To Not Be Susceptible

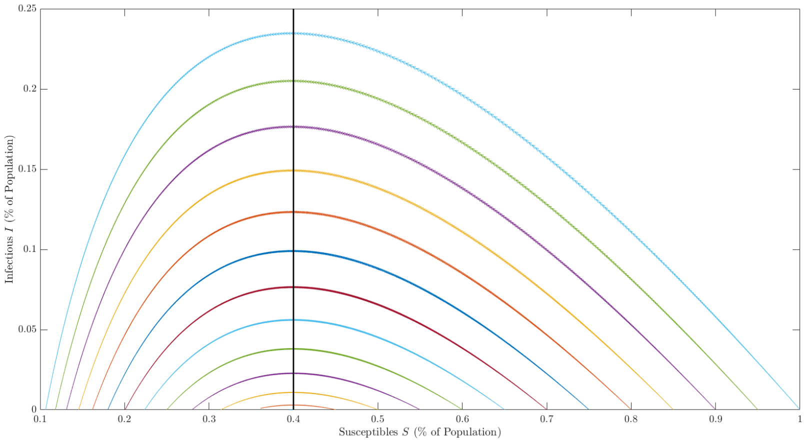

Modelling Lockdown

- Now let's introduce a lockdown policy into the model.
- Recall that the parameter β depends on both the intensity with which people meet others and the infectiousness of the disease. Governments cannot influence the infectiousness of the disease but they can reduce the intensity with which people meet each other using lockdowns.
- The programme `SIRLockdown.m` takes the standard case of $I_0 \approx 0$ and $S_0 \approx 1$ and allows a 50-day lockdown from period 25 onwards.
- We simulate lockdowns for five different cases, reducing β by 20%, 40%, 60%, 80% and 100%.
- The chart on the next page shows the outcomes from the five different lockdowns: Can you guess which lockdown policy produces the best outcome in terms of lowering the peak level of infections and the total of people to get infected?

50 Day Lockdowns Of Various Intensities

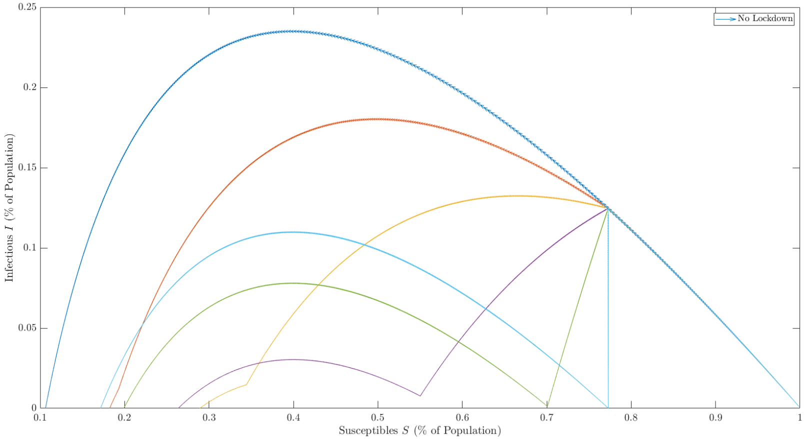

Which Is the Most Effective Lockdown?

- The lines are as follows
  - 20 % Reduction in β: Orange line
  - 40 % Reduction in β: Yellow line
  - 60 % Reduction in β: Purple line
  - 80 % Reduction in β: Green line
  - 100 % Reduction in β: Light blue line
- The most effective lockdown (in this very specific case) is a 40% reduction in activity for 50 days (the yellow line).
- The 20% reduction is not enough to stop cases rising at a fairly fast pace. On the other hand, the more stringent reductions all fail to get below the herd immunity level of susceptibility when lockdown is removed, leading to second waves of various sizes.
- The temporary full lockdown ends up with more infected people than all the other lockdowns.
- The graphs on the next page compare the time-paths of S_t, I_t and R_t without any lockdown and with the 50-day 40% lockdown.

Time Path for 50 Period 40% Reduction in Activity

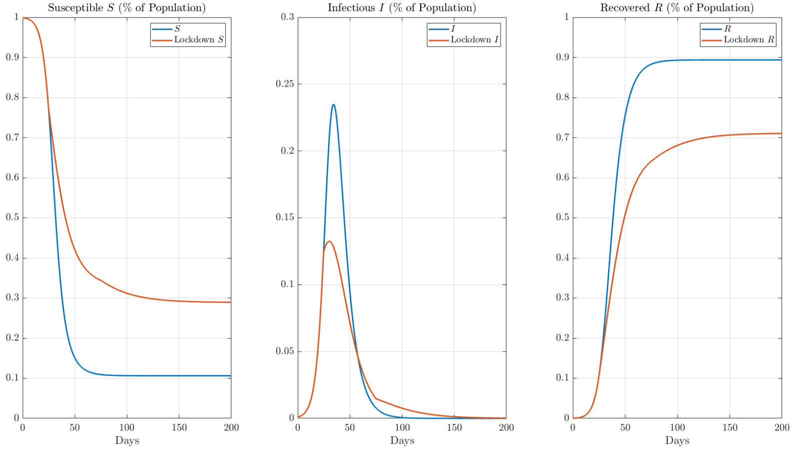

Caveats

- I am not an epidemiologist. Don't take any of this as proof that "loose lockdowns work best" or that this is my personal opinion.
- There are lots of caveats one could apply to these results:
  - A strict lockdown allows time for finding better ways to treat the infection, for building up health system capacity and for progress to be made on finding a vaccine. This was clearly relevant for Covid-19.
  - Strict lockdowns may allow time to implement a testing system that allows governments to change the dynamics and suppress the virus after the lockdown is over.
  - The model assumes that people don't change their behaviour in response to learning about the infection and its progress. Endogenising behaviour could produce different predictions.

Part III

Example: The Economics of Easter Island

A Malthusian Model of Population and Resources

- For another example of using differential equations and phase diagrams to answer interesting economic questions, we will discuss the model in Brander and Taylor's 1998 *American Economic Review* paper "The Simple Economics of Easter Island: A Ricardo-Malthus Model of Renewable Resource Use."
- The model that combines a Malthusian approach to population dynamics with modelling changes in a renewable resource base.

The World's Most Remote Place

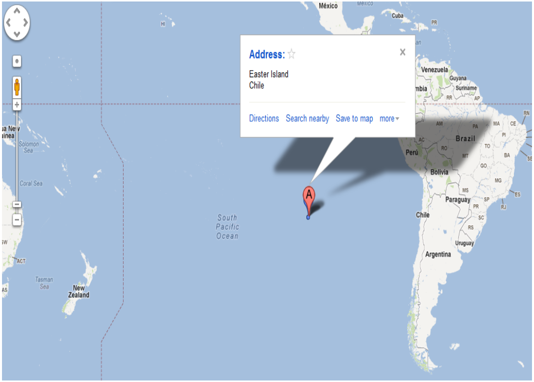

Easter Island Statues

Standing and Toppled Statues

Model of Resources: Population and Harvests

- The model economy has N_t people.
- They sustain themselves by collecting a harvest, H_t from a renewable resource stock denoted by S_t.
- The model consists of three elements:
- **The Change in Population**: This depends positively on the amount of harvest per person and on an exogenous factor $d > 0$ (without a harvest, there is a certain percentage reduction in population).

$$\frac{dN_t}{dt} = -dN_t + \theta H_t$$

- **The Harvest**: The harvest reaped per person is a positive function of the size of the resource stock.

$$\frac{H_t}{N_t} = \gamma S_t$$

Model of Resources: Stock of Resources

- The final element in the model is **the change in the resource stock**.
- We are describing a resource stock that is renewable. It doesn't simply decline when harvested until it is all gone.
- Instead, it has its own capacity to increase. For example, stocks of fish can be depleted but will increase naturally again if fishing is cut back.
- So, our equation for the change in resources is

$$\frac{dS_t}{dt} = G\left(S_t\right) - H_t$$

- The second term on the right-hand-side captures that the resource stock is reduced by the amount that is harvested.
- The first element describes the ability of the resource to grow.

Renewal of the Stock of Resources

- Brander and Taylor use a logistic function to describe how the resource stock renews itself
$$G(S_t) = rS_t(1 - S_t)$$
- The maximum level of resources is $S_t = 1$: At this level, there can no further increase in S_t.
- If $S_t = 0$ so the resource base has disappeared, then it cannot be regenerated.
- For all levels in between zero and one, we can note that
$$\frac{G(S_t)}{S_t} = r(1 - S_t)$$
- So the amount of natural renewal as a fraction of the stock decreases steadily as the stock reaches its maximum value of one. If the stock gets very low, it can grow at a fast rate if there is limited harvesting.

Dynamics of Population

- We are going to use a phase diagram to describe the joint dynamics of N_t and S_t.
- Inserting the equation for the harvest into the equation for the change in population we get
$$\frac{dN_t}{dt} = -dN_t + \theta\gamma S_t N_t$$
- This equation shows us that population growth is a positive function of the resource stock.
- This means there is a particular value of the resource stock, S^*, for which population growth is zero. When resources are higher than S^* population increases and when it is lower than S^* population declines.
- The value of S^* can be calculated as
$$-dN_t + \theta\gamma S^* N_t = 0 \Rightarrow S^* = \frac{d}{\theta\gamma}$$

Phase Diagram: Population Dynamics

- The figure on the next slide shows how we illustrate the dynamics with a phase diagram.
- We put population on the x-axis and the stock of resources on the y-axis.
- Population dynamics can then be described as follows:
    1. Unchanged population corresponds to a straight line at S^*.
    2. For all values of resources above S^* population is increasing: Thus in the area above the line, we show an arrow pointing right, meaning population is increasing.
    3. In the area below this line, there is an arrow pointing left, meaning population is falling.

Population Dynamics

Resource Dynamics

- Combining logistic renewal with harvest equation, resource stock dynamics are

$$\frac{dS_t}{dt} = rS_t\left(1 - S_t\right) - \gamma N_t S_t$$

- The stock of resources will be unchanged for all combinations of S_t and N_t that satisfy

$$rS_t\left(1 - S_t\right) - \gamma N_t S_t = 0 \Rightarrow N_t = \frac{r\left(1 - S_t\right)}{\gamma}$$

- This means that there is downward sloping line in $N - S$ space along each point of which the change in resources is zero.

Phase Diagram: Resource Dynamics

- The upper point crossing the S axis corresponds to no change because $S = 1$ and there are no people.
- As we move down the line we get points that correspond to no change in the stock of resources because while there are progressively larger numbers of people, the stock gets smaller and so can renew itself at a faster pace.
- Growth rate of resources depends negatively on the level of the stock:
  1. Every point that lies above the downward-sloping $\frac{dS}{dt} = 0$ line has a higher level of resources than the point on line below it. So S_t is declining for every point above the line and increasing for every point below it, hence arrow pointing down.
  2. In the area below this line, there is an arrow pointing up, meaning the stock of resources is increasing.

Resource Dynamics

Joint Dynamics of Population and Resources

- In the next figure, we put together the four arrows drawn in the two previous figures.
- The joint dynamics of population and resources can be divided up into four different quadrants.
- We can also see that there is one point at which both population and resources are unchanged.
- We know already that the level of the resource stock at this point is $S^* = \frac{d}{\theta\gamma}$.
- The level of population associated with this point is:

$$N^* = \frac{r\left(1 - \frac{d}{\theta\gamma}\right)}{\gamma} = \frac{r\left(\theta\gamma - d\right)}{\theta\gamma^2}$$

Combining Population and Resource Dynamics

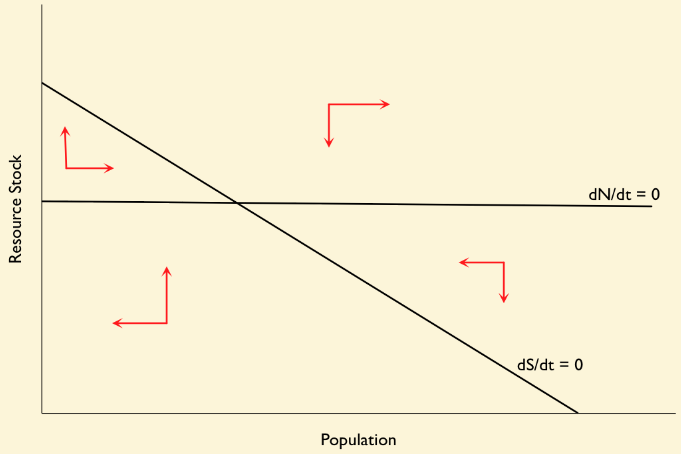

Joint Dynamics of Population and Resources

- This point is clearly some kind of "equilibrium" in the sense that once the economy reaches this point, it tends to stay there.
- But is the economy actually likely to end up at this point?
- Yes. By calculating the slopes of the trajectories in each quadrant, Brander and Taylor find that from any interior point (i.e. a point in which there is a non-zero population and resource stock) the economy eventually ends up at (N^*, S^*).
- You can show that

$$
\begin{aligned}
\frac{1}{N_t}\frac{dN_t}{dt} &= \theta\gamma\left(S_t - S^*\right) \\
\frac{1}{S_t}\frac{dS_t}{dt} &= \gamma\left(N^* - N_t\right) + r\left(S^* - S_t\right)
\end{aligned}
$$

so the dynamics of both population and the resource stock are both driven by how far the economy is from this equilibrium point.

Harvesting and Long-Run Population

- How does more intensive harvesting (higher γ) affect the long-run equilibrium level of population N^*?
- You can show that

$$\frac{dN^*}{d\gamma} = \frac{r}{\gamma^2}\left(2S^* - 1\right)$$

- The right-hand side here may be greater than or less than zero.
- Whether an increase in γ raises or reduces the equilibrium population depends on the size of the equilibrium level of resources.
  1. If $S^* > 0.5$ then a more intensive rate of harvesting raises the population even though it reduces the total amount of resources.
  2. If $S^* < 0.5$ then a more intensive rate of harvesting reduces the population because it reduces the total amount of resources.
- Easter Island devastation scenario (ending with very low resource stock) more like the latter case.
- Note I've simplified things here a bit. S^* itself depends upon γ but the calculation above doesn't take that into account.

Back to Easter Island

- Let's go back to Easter Island and imagine the island in its early days with a full stock of resources and very few residents. What happens next?
- For many years, the population expands and resources decline.
- Then, when it moves into the bottom right quadrant, population falls and resources keep declining.
- It moves though the quadrants and ends up at equilibrium with $S = S^*$ and $N = N^*$
- Our theoretical Easter Island sees its population far overshoot its long-run equilibrium level before collapsing below this level and then oscillating around the long-run level and then finally settling down.

Illustrative Dynamics Starting from Low Population and High Resources

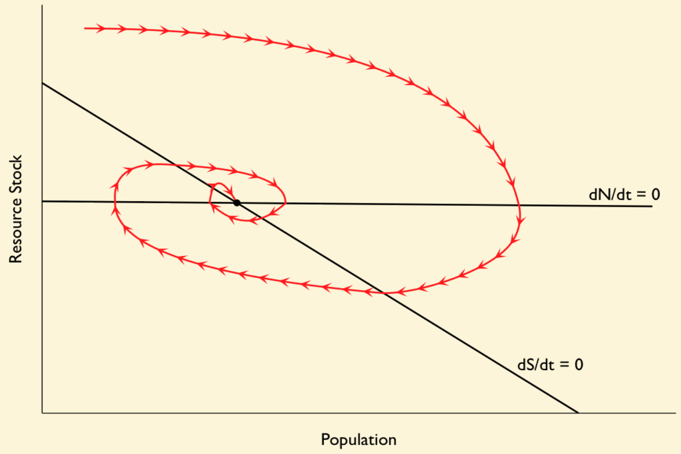

Matlab Programme Output

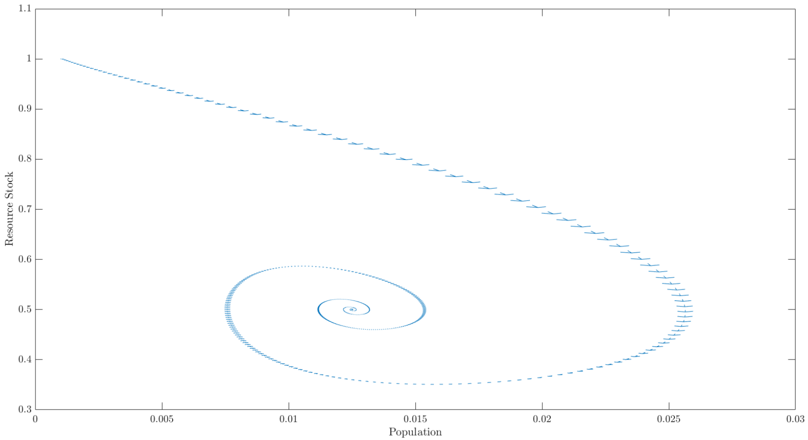

Dynamics Over Time

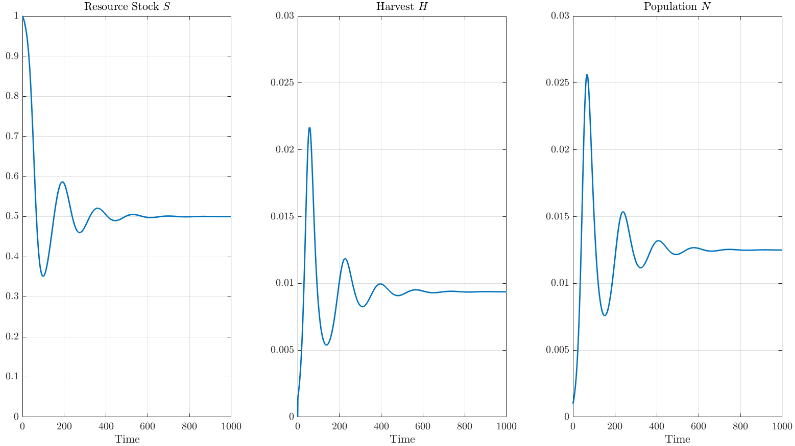

Dynamics Over Time with Less Harvesting

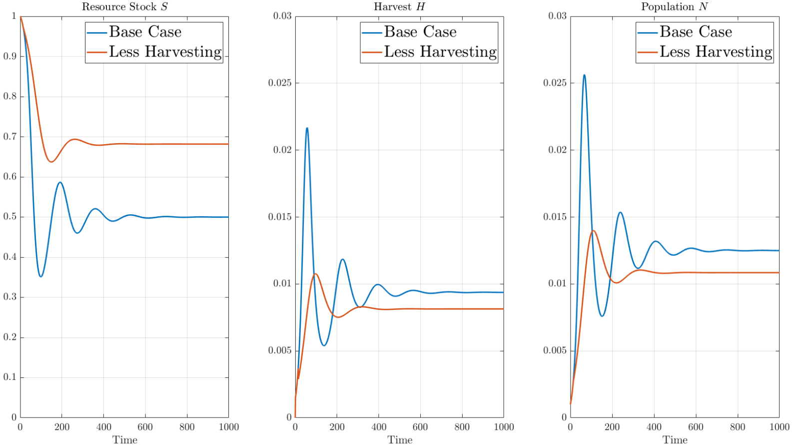

Why Doesn't Someone Shout Stop?

- In his book, *Collapse*, Jared Diamond discusses Easter Island and a number of other cases in which societies saw dramatic collapses, many triggered by long-term environmental damage.
- Diamond points to a number of potential explanations for why societies can let environmental damage occur up to the point where they trigger disasters.
  - **The Tragedy of the Commons**: It may simply never be in anyone's individual interests at any point in time to prevent environmental degradation. Need political institutions to take into account externalities associated with self-interested behaviour.
  - **Failure to Anticipate**: Societies may not realise exactly how much damage they are doing to their environment or what its long-term consequences will be.
  - **Failure to Perceive, Until Too Late**: Diamond notes that environmental change often occurs at such a slow pace that people fail to notice it and plan to deal with it.
- Analogies with current debate about climate change are clear.