

S-I-R Model of Epidemics

Part 1

Basic Model and Examples

Revised January 13, 2009

`sysid`

Mathematica 6.0.3, DynPac 11.01, 1/13/2009

1. Introduction

■ Description of the Model

In this notebook, we develop in detail the standard S-I-R model for epidemics. For the integration of the nonlinear differential equations, we use the package DynPac. Although some familiarity with DynPac is assumed, brief descriptions of some DynPac commands are given as they are used. The integrations in this notebook also can be done easily with the *Mathematica* function NDSolve.

Much of our presentation is tied to the specific example given in section 3, an influenza epidemic in a British Boarding School, and for that example we follow the treatment given by J.D. Murray (**Mathematical Biology I, An Introduction**, p. 325-326, Springer-Verlag, 2002).

The S-I-R model was introduced by W.O. Kermack and A.G. McKendrick ("A Contribution to the Mathematical Theory of Epidemics," Proc. Roy. Soc. London **A 115**, 700-721, 1927), and has played a major role in mathematical epidemiology. A summary of the model and its uses is given by Murray. In the model, a population is divided into three groups: the susceptibles S , the infectives I , and the recovered R , with numbers s , i , and r respectively. The total population is

$$n = s + i + r . \quad (1)$$

The susceptibles are those who are not infected and not immune, the infectives are those who are infected and can transmit the disease, and the recovered are those who have been infected and are immune (recovered or dead). In the simple form of the model considered here, we assume that the recovered are permanently immune. We also assume that the whole event is of sufficiently short duration that we may ignore natural births and deaths during the epidemic. Finally, we ignore any subdivisions of the population by age, sex, mobility, or other factors, although such distinctions are obviously of importance.

■ Differential Equations and Parameters of the Model

The basic differential equations are as follows:

$$\frac{ds}{dt} = -\alpha si, \quad (2)$$

$$\frac{di}{dt} = \alpha si - \beta i, \quad (3)$$

$$\frac{dr}{dt} = \beta i, \quad (4)$$

These equations describe the transitions of individuals from S to I to R. By adding the three equations, we show easily that the total population n is constant. We will call the parameter α the transmissivity and the parameter β the recovery rate. The term αsi is a standard kinetic term, based on the idea that the number of encounters per unit time between susceptibles and infectives will be proportional to the numbers of each. The transmissivity α is determined by both the encounter frequency and the efficiency with which the disease is transmitted per encounter. In later notebooks, we will look more closely at the problem of estimating α . We will see later that there are advantages to formulating the problem in terms of population fractions rather than total population. For example, it is obvious that if a given population is placed into a much larger area, the encounter rate will decrease and α will decrease. A proper formulation in terms of fractions will give a transmissivity independent of total population in the first approximation. For the present, we continue with the formulation in terms of total populations. We will not be concerned further in this notebook with the estimation of α from first principles.

The interpretation and estimation of β is more straightforward. We now show that β^{-1} is the average duration of the infection. Consider the special case of no new infections, so that equation (3) reduces to $di/dt = -\beta i$. Then $i(t) = i(0)e^{-\beta t}$. The number of infectives recovering in time $(t, t + dt)$ is $|di| = |(di/dt)dt| = i(t)\beta dt = i(0)\beta e^{-\beta t} dt$, and this is the number of infectives with infection durations in the range $(t, t + dt)$. The average duration t_{dur} is then the average of t with this as a weighting function, hence

$$t_{\text{dur}} = \frac{\int_0^{\infty} i(0) \beta e^{-\beta t} t dt}{\int_0^{\infty} i(0) \beta e^{-\beta t} dt} = \frac{1}{\beta}. \quad (5)$$

We may treat the three equations given by (2) - (4) as a set of two equations in s and i (equations (2) and (3)). After these equations are solved for s and i , we may calculate r from equation (1).

■ Conditions for an Epidemic

In epidemiology, the word "epidemic" has a technical meaning: it is a situation in which the number of infectives increases from the initial value. The condition for an epidemic follows directly from equation (3): di/dt will be positive whenever

$$s > \beta/\alpha. \quad (6)$$

Thus for given α and β there is a critical number of susceptibles for an epidemic, and if s_0 is the initial number of susceptibles, then the condition for an epidemic is $s_0 > \beta/\alpha$. There is a simple interpretation of equation (6). A given infective will, on the average, be infectious for a time $1/\beta$. The number of susceptibles infected by one infective per unit time is αs . Hence the total number of infections produced by one infective is $\alpha s/\beta$, and for an event to qualify as an epidemic, this number must exceed 1, which gives condition (6). The parameter appearing here,

$$R_0 = \alpha s_0 / \beta, \quad (7)$$

is called the reproduction ratio of the epidemic.

It is somewhat surprising that the number of initial infectives does not play a role in determining whether there is an epidemic. Of course other aspects of the epidemic will depend on the number of initial infectives -- especially the time of the peak of the epidemic. The peak of the epidemic is the time of maximum number of infectives. We see from equation (3) that this will occur when $s = \beta/\alpha$ -- that is, when the reproduction ratio is 1.

■ Equilibrium and Stability

It is easy to show that the equilibrium states of equations (2) - (4) are those states in which $i = 0$, and $s = s_*$, where s_* is any positive constant. Thus the equilibria are non-isolated. A formal stability analysis by linearization yields two eigenvalues: 0, and $\alpha s_* - \beta$. If the condition for an epidemic is satisfied, this second eigenvalue is positive and the equilibrium is unstable. If the second eigenvalue is negative, there is no conclusion about the stability from the linearization because of the zero eigenvalue. However, it is obvious that the equilibrium is not strictly stable even in that case, because any perturbation with non-zero i will lead to a situation in which s decreases and therefore does not return to s_* .

Stability isn't really the primary issue in this model. We are much more interested in the time course caused by the introduction of some infectives, and in the final asymptotic state of the system. Some typical questions of interest are: (1) How many susceptibles become infected? (2) What is the peak number of infectives? (3) When does the epidemic peak? (4) How does the time course depend on the initial number of infectives? Before attempting to answer these questions in any generality, we look at an example, in which the model results are compared with observations of a flu epidemic. Before doing that, we define the system for DynPac so that we may carry out the integrations in the example.

2. Defining the Equations for DynPac

In much of what follows, we will be integrating equations (2) - (4) numerically. In this section, we define the equations for DynPac. We start by defining the state vector, the parameter vector and the slope vector.

```
setstate[{s, i}];
```

This tells DynPac that the two state variables are s and i .

```
setparm[{α, β}];
```

This tells DynPac that the "official" system parameters are α and β .

```
slopevec = {-α * s * i, α * s * i - β * i};
```

This defines the slope vector of the equations for DynPac.

We also define an initial time t_0 , a time step h , and the number of time steps $nsteps$. For the events studied here, we use days as the time unit. These parameters may be adjusted later as we work through other examples. The choices below take us out to 25 days with 20 time steps per day.

```
t0 = 0.0;
h = 0.05;
nsteps = 500;
plotreset; imsize = 250;
intreset;
```

The DynPac commands `plotreset` and `intreset` perform resets to default options for plotting and integrating. Other DynPac commands will be used later when we graph the results of our integrations.

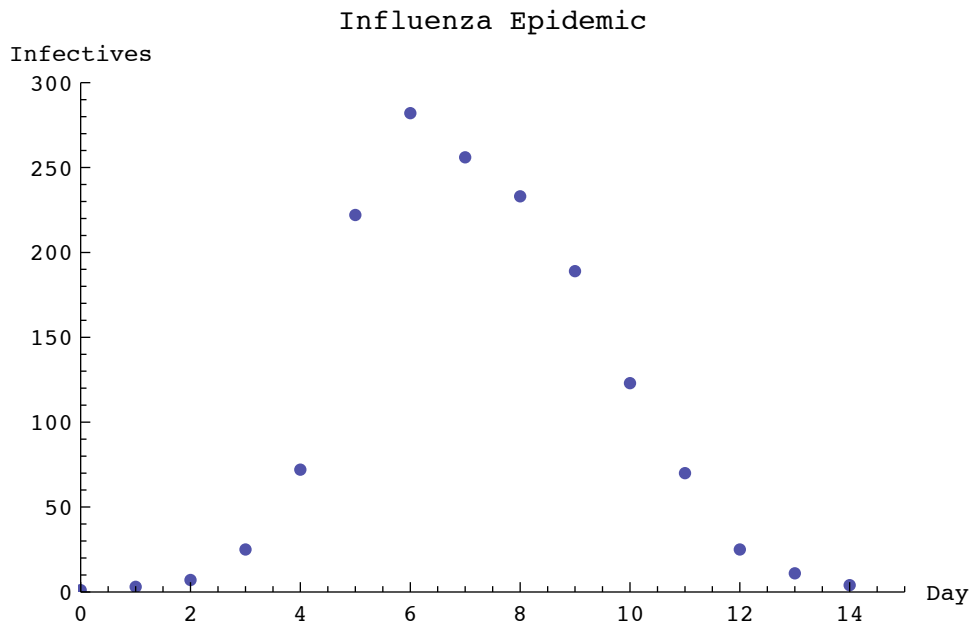
3. Example - An Influenza Epidemic

This example was given by Murray, based on data compiled by the British Communicable Disease Surveillance Centre (British Medical Journal, March 4 1978, p. 587). The event was a flu epidemic in a boys boarding school in the north of England. There were 763 resident boys, including one initial infective. The data (given in graphical form only) is for residents confined to bed, and, following Murray, we will assume those to be the infectives. The data for the two-week epidemic are given in the list below, constructed by reading values from the graph in the original publication.

```
flulist = {{0, 1}, {1, 3}, {2, 7}, {3, 25}, {4, 72}, {5, 222}, {6, 282}, {7, 256},
           {8, 233}, {9, 189}, {10, 123}, {11, 70}, {12, 25}, {13, 11}, {14, 4}};
```

We plot these values:

```
fludatagraph =
ListPlot[flulist, AxesLabel -> {"Day", "Infectives"}, PlotLabel -> "Influenza Epidemic",
PlotRange -> {{0, 15}, {0, 300}}, PlotStyle -> PointSize[0.015]]
```



In this example, one looks for values of the transmissivity α and the duration β which give a good fit of the model to the data. The initial conditions are

```
s0 = 762; i0 = 1;
```

According to Murray, the data are best fit with $\alpha = 2.18 \times 10^{-3} \text{day}^{-1}$ and $\beta = 0.441 \text{day}^{-1}$. Let's try these values and see how good the fit is. We specify these values for DynPac as values of the parameter vector:

```
parmval = {2.18 * 10-3, 0.441};
```

Now we are ready to integrate the equations and to compare the results with the data. We use the DynPac command `integrate`, which takes as arguments the initial conditions, the initial time, the time step, and the number of time steps.

```
flusol = integrate[{s0, i0}, t0, h, nsteps];
```

First we look at the overall solution, and then we compare it with the data. We set the plot range and the aspect ratio. We use `sysname` to give the system a name to be used in plot labels. We set `sciflag = True`, which causes parameter values in the graph label to be printed in scientific form. We produce a plot of the solution versus time with the DynPac command `timeplot`, which takes as arguments the name of the solution, and the indices of the components to be plotted.


```
plrange = {{0, 20}, {0, 800}};
```

```
asprat = 0.5;
```

```
sysname = "Flu Epidemic";
```

```
sciflag = True;
```

```
timeplot[flusol, {1, 2}]
```




We also construct a graph of the incidence for the epidemic, which is defined to be the rate at which new cases appear -- in this model equal to αsi .

```
inc = transol[flusol, {\alpha * s * i}];
```

```
axlabon = "incidence";
```

```
plrange = {{0, 20}, {0, 250}};
```

```
timeplot[inc, 1]
```




```
axlabon = True;
```

Now we look at a comparison of theory and data for the number of infectives. We change to a better graph scale, and combine the theory and data graphs without displaying the theory graph separately.

```
plrange = {{0, 15}, {0, 300}};
```

```
infgraph = timeplot[flusol, 2];
```

```
show[infgraph, fludatagraph]
```



The fit is not bad, and it is clear that the model predicts the right general features. You might want to test the robustness of this fit by repeating the comparison for different values of α and β .

The average recovery time is $1/\beta = 1/0.441 = 2.27$ days. This time scale is not particularly obvious in the raw data, so it is interesting that the fit of data to model gives us a value for it. From the data given in the original article in the British Journal of Medicine, one may infer that this figure is in fact consistent with the time patients were confined to bed.

Let's see what the model predicts about the total number who eventually get the flu. We look at the last data point in the solution:

```
lasttx
```

```
{25., 19.3234, 0.32234}
```

Thus at 25 days, we have 19 susceptibles left, and 0.3 infectives. Although with only 0.3 infectives, we don't expect significant further change, let's make sure by running the epidemic longer.

```
nsteps = 1000;
```

```
flusol2 = integrate[{s0, i0}, t0, h, nsteps];
```

```
lasttx
```

```
{50., 19.2894, 0.0000150265}
```

No significant change. Thus the model predicts that all but 19 of the residents will get the flu. This raises the question of whether we can make such final state predictions in general, without necessarily running through the entire numerical integration. In dealing with that question, it is helpful first to develop the phase-plane view of this process.

4. Phase Plane Picture

We start by looking at a phase plane plot of our last solution. We fix the aspect ratio at 1, set an appropriate plot range, ask for two arrows, at the $1/3$ and $2/3$ points, and use the command phaser to produce the plot.

```
asprat = 1.0;
```


```
plrange = {{0, 800}, {0, 800}};
```

```
arrowflag = True;
```

```
arrowvec = {1 / 3, 2 / 3};
```

```
phasegraph = phaser[flusol2]
```

Flu Epidemic $\{\alpha, \beta\} = \{2.18 \times 10^{-3}, 4.41 \times 10^{-1}\}$



As the graph shows, the number of infectives increases until s has dropped to about 200. The exact value at which this happens is the value which makes the reproductive ratio R equal to 1, which is $s = \beta/\alpha$:

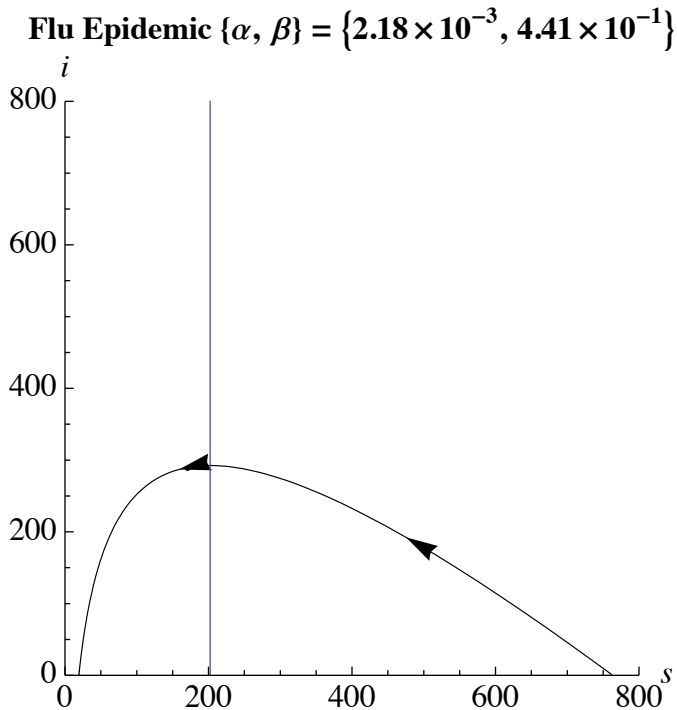
```
evalparm[ $\beta / \alpha$ ]
```

```
202.294
```

Here we have used the DynPac command `evalparm`, which replaces parameter symbols by the current parameter values stored in `parmval`. Let's add a vertical line to our phase plot at this critical value of s . We do that with a parametric plot which we then combine with the phase plot.

```
linegr = ParametricPlot[{evalparm[ $\beta / \alpha$ ], q}, {q, 0, 800}];
```

show[phasegraph, linegr]



Any initial point to the right of the line leads to an epidemic. Any initial point to the left of the line gives a monotonic decay of number of infectives. Clearly all possible final states are on the s -axis, to the left of the line. We would like to be able to predict the final state from a knowledge of the initial state. This turns out to be possible, because there is an analytical expression for the orbits of this system. We consider that next.

5. Predicting the Final State

One of the interesting predictions of the S-I-R model is that the epidemic stops for lack of infectives, not lack of susceptibles. At the end of the epidemic, there may be only a few susceptibles left, as in the example above, or there may be many. The S-I-R equations are sufficiently simple that we may obtain by analytical means some useful results about the final state.

We start with equations (2) and (3), and we combine them to get a first-order equation for the orbits in the i - s phase space:

$$\frac{di}{ds} = \frac{\beta}{\alpha s} - 1, \quad (8)$$

which we integrate to get

$$i = i_0 + (\beta/\alpha)\log(s/s_0) - (s - s_0) \quad (9)$$

As $t \rightarrow \infty, i \rightarrow 0$, so

$$f(s_\infty) = i_0,$$

$$\text{where } f(s_\infty) = (s_\infty - s_0) - (\beta/\alpha)\log(s_\infty/s_0) \quad , \quad (10)$$

and this determines s_∞ , the final value of the number of susceptibles. It is easy to establish the following properties of the function $f(s_\infty)$: (1) $f(s_0) = 0$; (2) $f(0) = \infty$; (3) if $R_0 > 1$, f takes on negative values; (4) f is monotonically increasing when it is positive. Thus for any given s_0 and i_0 there is a unique value of s_∞ satisfying equation (10). We can use *Mathematica* to find that value:

```
fifunc[s0_, i0_, sinf_] := evalparm[(β / α) * Log[s0 / sinf] - (s0 - sinf) - i0]
FindRoot[fifunc[762, 1, sinf], {sinf, 10}]
{sinf -> 19.2894}
```

Thus in any given case, we can easily find the final state from a knowledge of the initial state.

For a strong epidemic, defined as one in which R_0 is large, we can obtain a simple analytical approximation for s_∞ . We begin by rewriting equation (10) as

$$(1/R_0)\log(s_0/s_\infty) = (1 - s_\infty/s_0) + i_0/s_0 \quad . \quad (11)$$

For R_0 large, s_∞/s_0 will be small, and we can drop it in comparison with 1 on the right hand side of (11). This gives

$$s_\infty = s_0 e^{-R_0(1+i_0/s_0)} = s_0 e^{-(\alpha/\beta)(s_0+i_0)} = s_0 e^{-(\alpha/\beta)n} \quad . \quad (12)$$

Let's see what this gives for the flu example:

```
evalparm[762 * Exp[-(α / β) * 763]]
17.535
```

Quite close to the actual value of 19.3, especially considering that the present value of R_0 is not all that large:

```
evalparm[762 * α / β ]
3.7668
```

In addition to knowing the final state, it is desirable to have simple estimates for peak values -- peak number of infectives (maximum number of beds needed), and the time of the peak (time when maximum number of beds are needed). We consider that next.

6. Predicting Epidemic Peaks

■ Calculation of Peak Values

In general, the prediction of peaks can be accomplished by numerical integration of the model equations. As an alternative, the analytical formulas developed above also can be used. As we have noted before, the value of s at the epidemic peak corresponds to $R = 1$, and is

```
speak := evalparm[β / α ]
```

For our present example:

```
speak
202.294
```

From the orbit equation (9) we get

```
ipeak[s0_, i0_] := evalparm[i0 + (β / α) * Log[speak / s0] - (speak - s0)]
```

For the present case:

```
ipeak[762, 1]
292.419
```

Thus we can easily get exact predictions of the magnitude of the peak values of both s and i .

The time of the peak is also of interest. We start again with the orbit equation (9).

```
inf[s_, s0_, i0_] := i0 + (β / α) * Log[s / s0] - s + s0
```

Now we integrate equation (2) to get

$$t_{\text{peak}} = - \int_{s_0}^{s_{\text{peak}}} (1 / (\alpha * s * i)) ds . \quad (13)$$

We carry this out numerically in *Mathematica*, using NIntegrate.

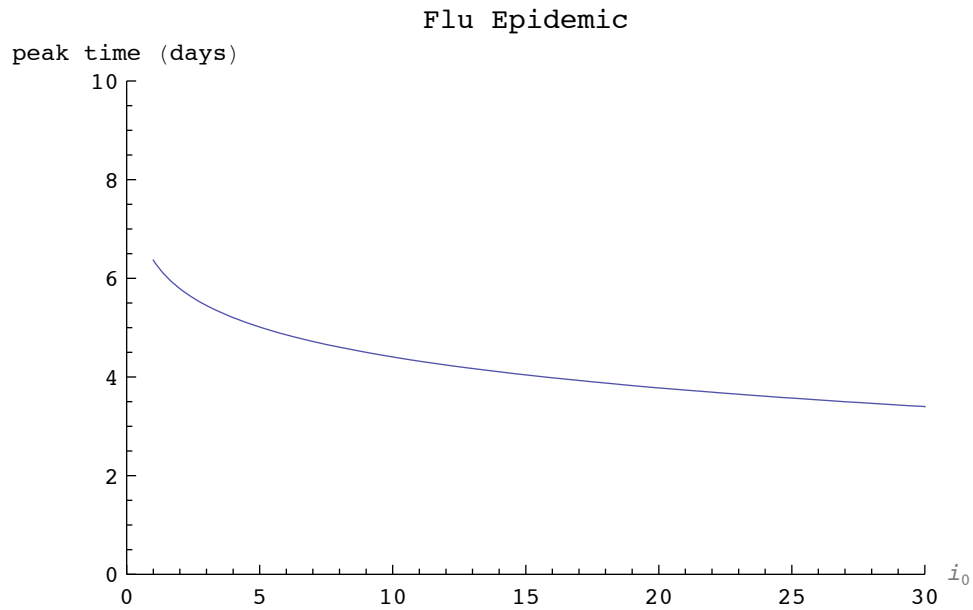
```
tpeak[s0_, i0_] := Module[{f, alph, bet, uplim, lowlim, ans},
  f[s_, ds0_, di0_] = evalparm[inf[s, ds0, di0]]; alph = evalparm[α];
  bet = evalparm[β]; uplim = bet / alph; lowlim = s0;
  ans = -NIntegrate[1.0 / (alph * s * f[s, s0, i0]), {s, lowlim, uplim}]; ans]
tpeak[762, 1]
6.36898
```

This last result is in agreement with the graphs obtained earlier.

■ Dependence of Peak Values on Initial Infectives

We take a quick look at the dependence of the time peak on the initial number of infectives by plotting the peak versus i_0 .

```
tpeakgraph = Plot[tpeak[762, i0], {i0, 1, 30},
  AxesLabel -> {"i0", "peak time (days)"}, PlotRange -> {{0, 30}, {0, 10}},
  PlotLabel -> "Flu Epidemic", AxesOrigin -> {0, 0}, Axes -> True]
```



Finally, we look at time courses for three different values of initial infectives: 1, 5, and 10. We integrate out to 20 days.

```
sol1 = integrate[{762, 1}, 0, 0.05, 400];
sol2 = integrate[{762, 5}, 0, 0.05, 400];
sol3 = integrate[{762, 10}, 0, 0.05, 400];


asprat = 0.7;

plrange = {{0, 20}, {0, 400}};

arrowflag = False;
```

```
timeplot[{sol1, sol2, sol3}, 2]
```

Flu Epidemic $\{\alpha, \beta\} = \{2.18 \times 10^{-3}, 4.41 \times 10^{-1}\}$



The curves are very similar in shape, and are close to being just time shifts of one another, so that the principal difference is in the timing of the peak. Let's repeat this using a different method of comparison, in which we keep the total population fixed as we vary the number of initial infectives.


```
sol1 = integrate[{762, 1}, 0, 0.05, 400];
```

```
sol2 = integrate[{758, 5}, 0, 0.05, 400];
```

```
sol3 = integrate[{753, 10}, 0, 0.05, 400];
```

```
timeplot[{sol1, sol2, sol3}, 2]
```

Flu Epidemic $\{\alpha, \beta\} = \{2.18 \times 10^{-3}, 4.41 \times 10^{-1}\}$



Very similar, but with even less variation in peak amplitude.

■ Summary of Peak Values

The peak of the epidemic is the state of maximum number of infectives. At the peak, the number of susceptibles is $s_p = \beta/\alpha$. The number of infectives is $i_{\text{peak}} = i_0 + \Delta i(s_p)$, where $\Delta i(s) = (\beta/\alpha) \cdot \text{Log}(s/s_0) + (s_0 - s)$. The time of the peak is given by equation (13) with i given by equation (9).

7. Another Example

This example was developed as a homework problem in a course in ordinary differential equations. Though pure fiction, it gives us one more chance to practice our newly developed techniques. Here is the problem:

An Army base has a total staff of 8342. An individual who has just returned from leave becomes ill and is diagnosed with Jade fever -- an exotic, dangerous and highly contagious variety of flu. All individuals getting this flu must be hospitalized. The base hospital has 240 beds. The transmission parameter for this flu at this base is $\alpha = 5 \times 10^{-5} \text{ day}^{-1}$, and the recovery parameter is $\beta = 0.32 \text{ day}^{-1}$.

- (a) Is the condition for an epidemic satisfied?
- (b) Construct a plot of $s(t)$, $i(t)$ and the incidence over the course of the epidemic.
- (c) Does the base hospital have enough beds?
- (d) How many of the base staff will get the flu?

We start by defining the parameter values for DynPac.

```
parmval = {5.0 * 10-5, 0.32};
```

The initial values are

```
s0 = 8341; i0 = 1;
```

■ Part (a)

To answer part (a), we compute the initial reproduction ratio for the epidemic:

```
R0 =  $\alpha * s0 / \beta$  /. Thread[parmvec -> parmval]
```

```
1.30328
```

Because this is greater than 1, there will be an epidemic.

■ Part (b)

Before integrating the equations, we look at the time for the peak, so that we will know how long to integrate.

```
tpeak[s0, i0]
```

```
72.1701
```

The peak occurs at 72 days, and this epidemic is a much longer event than our earlier measles epidemic. We will integrate out 150 days, using a time step of 0.2 days. We also look at the peak values first, as an aid in plotting the results.

```
speak
```

```
6400.
```

```
ipeak[s0, i0]
```

```
246.735
```

```
episol = integrate[{s0, i0}, 0.0, 0.2, 750];
```

We set some plotting parameters and then plot.

```
sysname = "Army Base Epidemic";
```


```
sciflag = True;
```

```
asprat = 0.7;
```

```
plrange = {{0, 150}, {0, 9000}};
```

```
sgraph = timeplot[episol, 1]
```

Army Base Epidemic $\{\alpha, \beta\} = \{5. \times 10^{-5}, 3.2 \times 10^{-1}\}$



```
plrange = {{0, 150}, {0, 250}};
```

```
igraph = timeplot[episol, 2];
```

We also construct a graph of the incidence for the epidemic, which is defined to be the rate at which new cases appear -- in this model equal to αsi .

```
inc = transol[episol, { $\alpha * s * i$ }];  
axlabon = "incidence";  
plrange = {{0, 150}, {0, 100}};  
timeplot[inc, 1];
```

■ Part (c)

We already answered this in part (b) above, where we calculated the peak number of infectives to be 247, so the hospital will need a few more beds.

■ Part (d)

To find how many got the flu, we examine the state of the system at the end of the integration. That state is stored in the variable `lasttx`.

```
lasttx  
{150., 4788.16, 1.60252}
```

We see that there is still 1.6 infectives, so the epidemic is almost, but not quite, over. We run it a little further.

```
episol2 = integrate[{s0, i0}, 0.0, 0.2, 1000];  
lasttx  
{200., 4783.49, 0.0282399}
```

Now the epidemic is finished, and there are 4783 susceptibles, meaning that the number who had the flu is

```
s0 + i0 - 4783  
3559
```