



Massachusetts Institute of Technology
Sloan School of Management

Simulating Epidemics Using VensimPLE

For use with VensimPLE

Originally developed by Professor Nelson Repenning, MIT System Dynamics Group.
Revised by John Sterman (MIT) and Professor Hazhir Rahmandad (Virginia Tech, ISE Department).

Last revision: October 2010 by John Sterman

I. Getting Started

The purpose of this tutorial is to help you develop some familiarity with building and analyzing system dynamics models using the VensimPLE software. To do so you are going to build a simple model that captures the dynamics of an infectious disease like SARS.


In this tutorial we assume you are already familiar with the use of VensimPLE for mapping system structure (see the Vensim tutorial “Mapping the Structure of an Epidemic Using VensimPLE”).

This tutorial takes you to the next level: starting from a map showing the stock-flow and feedback structure of an epidemic, this tutorial shows you how to enter equations, including units of measure for each variable, how to check your model for dimensional consistency, how to simulate it and compare the behavior of the model against historical data, and how to get graphical and tabular output.


The tutorial should give you the basic skill you need to begin to map, model, simulate and analyze your own models.

Background: SARS, a corona-virus, emerged in Asia in 2003. To focus our modeling effort, we will build a model to capture the course of the epidemic in Taiwan. The graphs below show the incidence (rate at which new cases were reported, measured in people/day) and cumulative prevalence (cumulative number of cases reported, measured in people) for SARS in Taiwan.

New Reported Cases

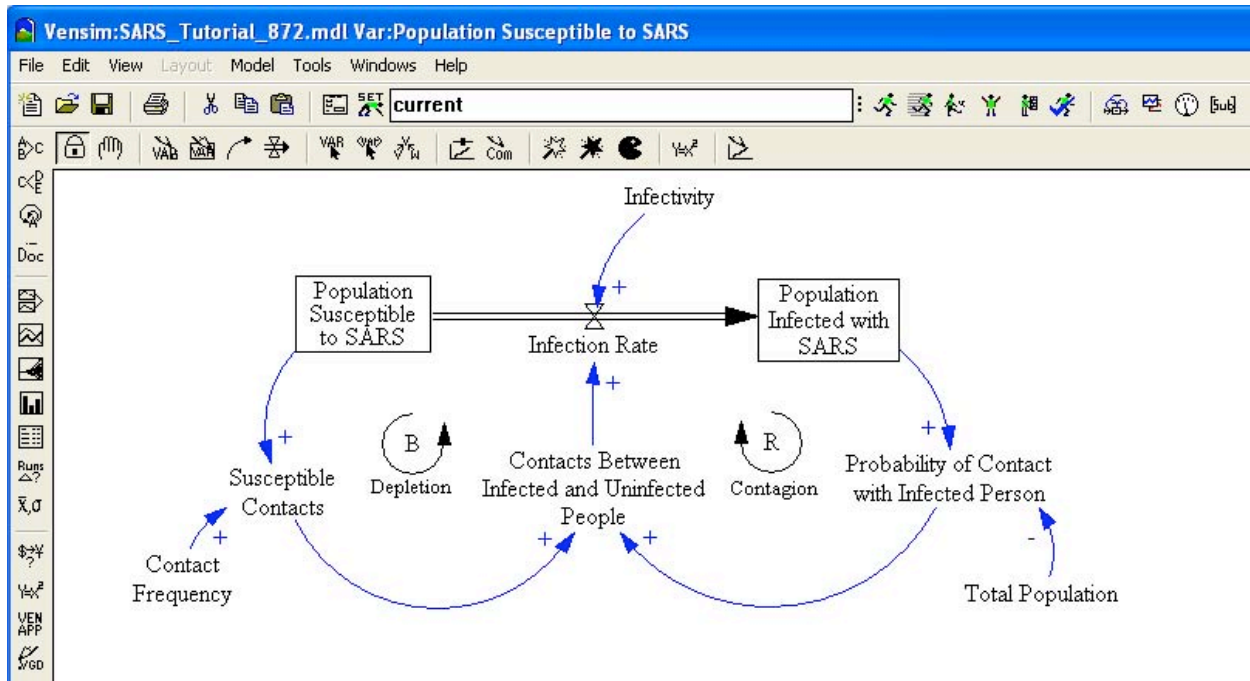


Cumulative Reported Cases



To begin, make sure you have the most current version of vensim, available for free download from <http://www.vensim.com>. Download the model “SARS_Tutorial_872.mdl” (the site containing this model will be provided by your instructor). Also download the file containing the historical data shown in the graphs above “SARSDATA.vdf”. Put these two files in the same folder.

Now open the SARS_Tutorial_872.mdl model (double click on its icon). You should see the following screen:



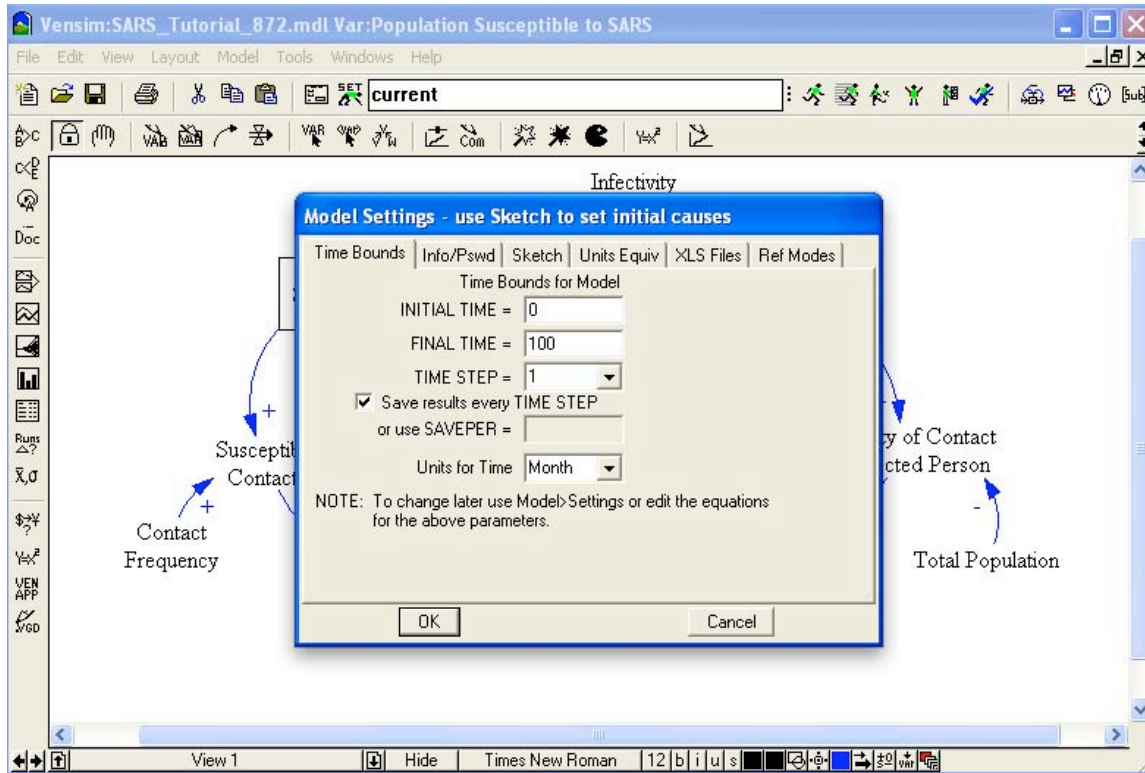
Vensim software copyright 1996-2011 Ventana Systems, Inc. Used with permission.

The diagram you see is similar to the map of an epidemic you created in the “Mapping the Structure of an Epidemic Using VensimPLE” tutorial. There are two stocks, the Population Susceptible to SARS and the Population Infected with SARS. When people contract SARS they flow from the susceptible to infected compartments via the Infection Rate.

The Infection rate depends on the number of contacts between infected and uninfected people each day and the fraction of those contacts resulting in the transmission of the virus, denoted “infectivity” (infectivity is a fraction between zero and one, representing the probability that a contact between an infected and uninfected individual results in transmitting the virus to the uninfected individual. That probability is less than one because not all contacts result in an infection. The number of contacts between infected and uninfected individuals depends on the total number of contacts among the susceptible population per day and the probability of contact with an infected individual. The probability of contact with an infected individual depends on the proportion of infected individuals in the total population. The total number of contacts the susceptible population has is the contact frequency (number of contacts each susceptible person has with others each day) and the total number of susceptible people.


II. Choosing Model Settings and Loading Data

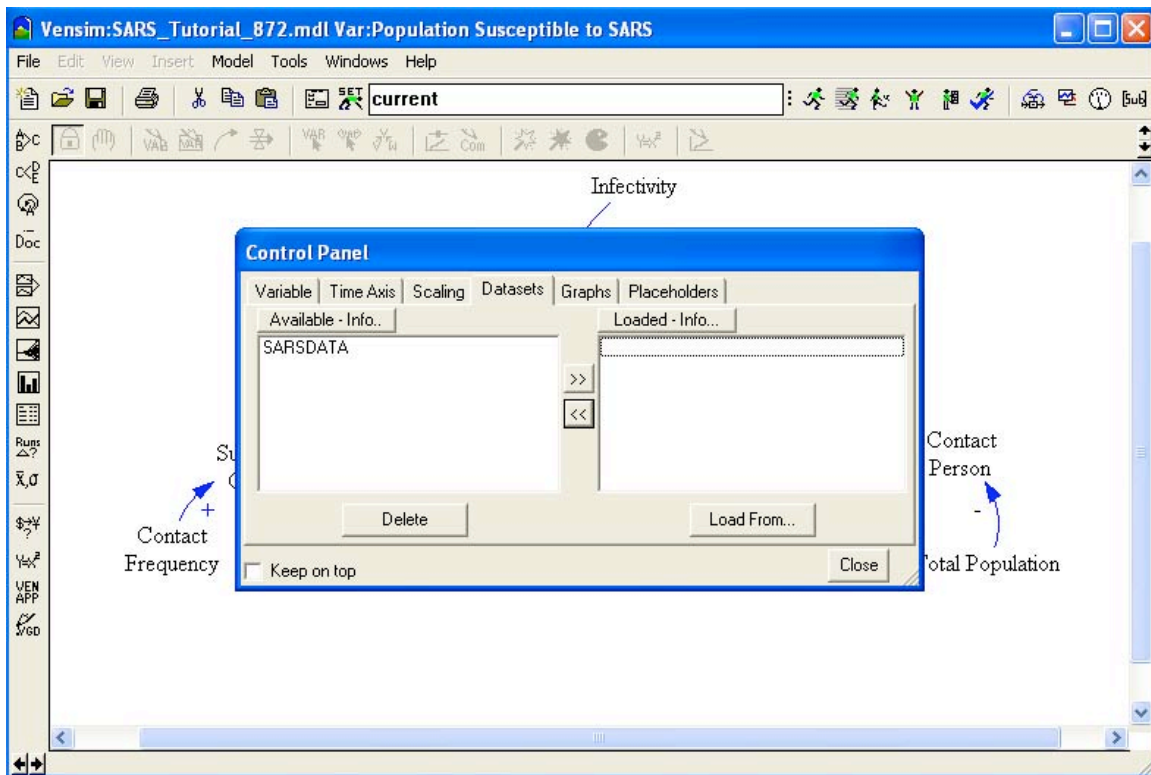
Before you enter the equations, you should specify when the simulation will start, when it will end, and other simulation control parameters. To do so, select **Settings...** from the **Model** menu. You should see:



Vensim software copyright 1996-2011 Ventana Systems, Inc. Used with permission.

Start your model of the Taiwanese SARS epidemic at day 0 (enter 0 in the **INITIAL TIME** box) and simulate for 120 days (enter 120 the **Final Time** box). Select a time step of 0.25 days. Check the box saying **Save results every TIME STEP**. Also select Day as the **Units for Time**. Click OK.

Finally, make sure the file with the historical data is loaded. To do so, first click on the  icon to access the control panel (which should look like this):




Vensim software copyright 1996-2011 Ventana Systems, Inc. Used with permission.

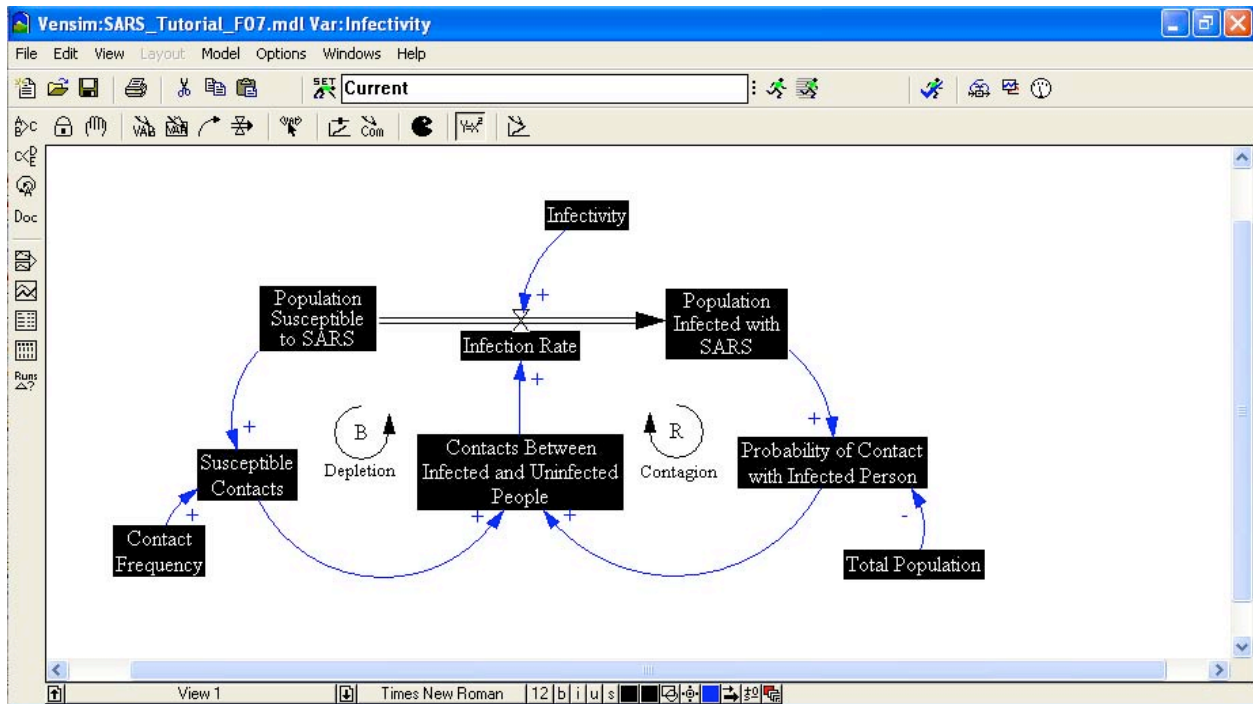
If you see something different, click on the **Datasets** tab. To load the data file, **SARSDATA**, select it by clicking and then hit the >> button which will move it to the **Loaded** column. The data set is now available.

Now you are ready to enter the equations for your model.

III. Specifying Equations for Your Model

You are now ready to write equations for each of the variables in your model.

To begin, click on the  button on the *top horizontal toolbar*. The variables in your diagram become highlighted.

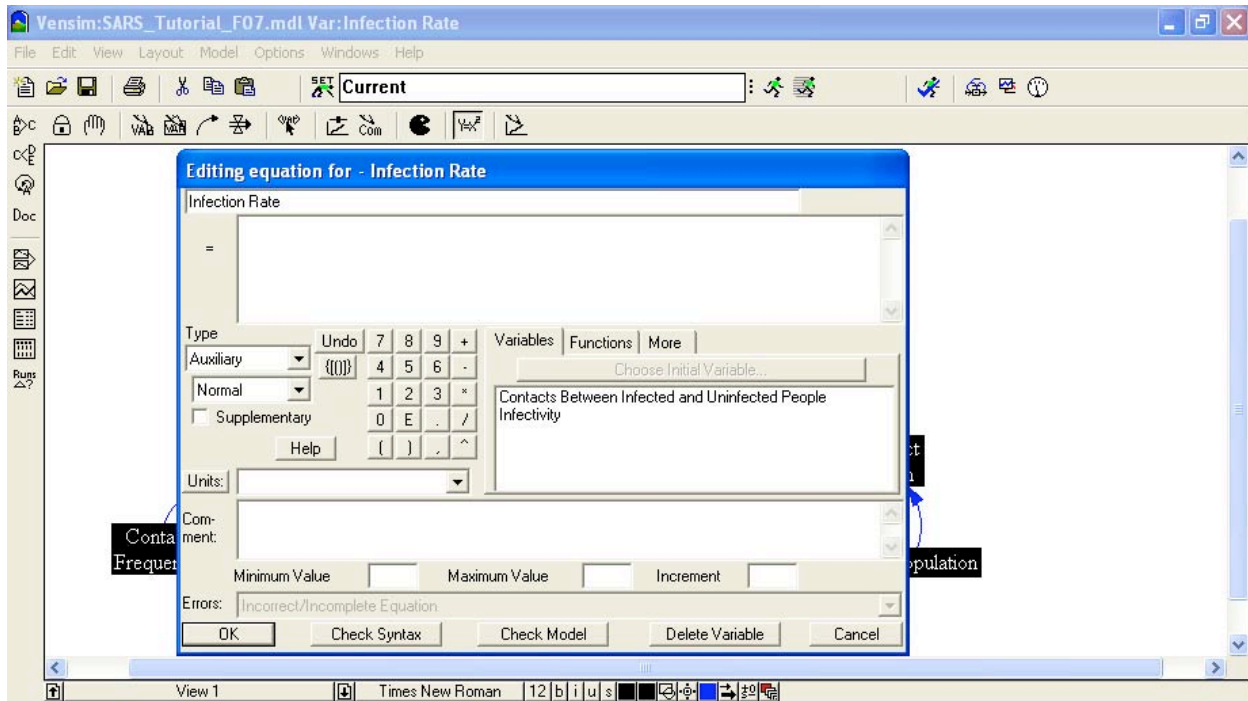


Vensim software copyright 1996-2011 Ventana Systems, Inc. Used with permission.

A highlighted variable indicates that the equation for that variable is incomplete.

Variables in system dynamics models are classified as either exogenous or endogenous. Exogenous variables are those that are defined independent of other variables of the model. They are functions of time (i.e., Exogenous Variable = $f(t)$). Of course the exogenous variables may be constants, in which case they are called parameters. Endogenous variables are members of determined based on other variables (Endogenous Variable = $f(x, y, z)$, where x, y, z are other variables in the model). Your model has three parameters (constants)—**Infectivity**, **Contact Frequency**, and the **Total Population**. There are no exogenous variables, and six endogenous variables—**Population Susceptible to SARS**, **Infection Rate**, **Population Infected with SARS**, **Probability of Contact with Infected Person**, **Susceptible Contacts**, and **Contacts Between Infected and Uninfected People**.

Start by writing the equations for the endogenous variables. Begin by clicking on the **Infection Rate**. You then see the following dialog box:

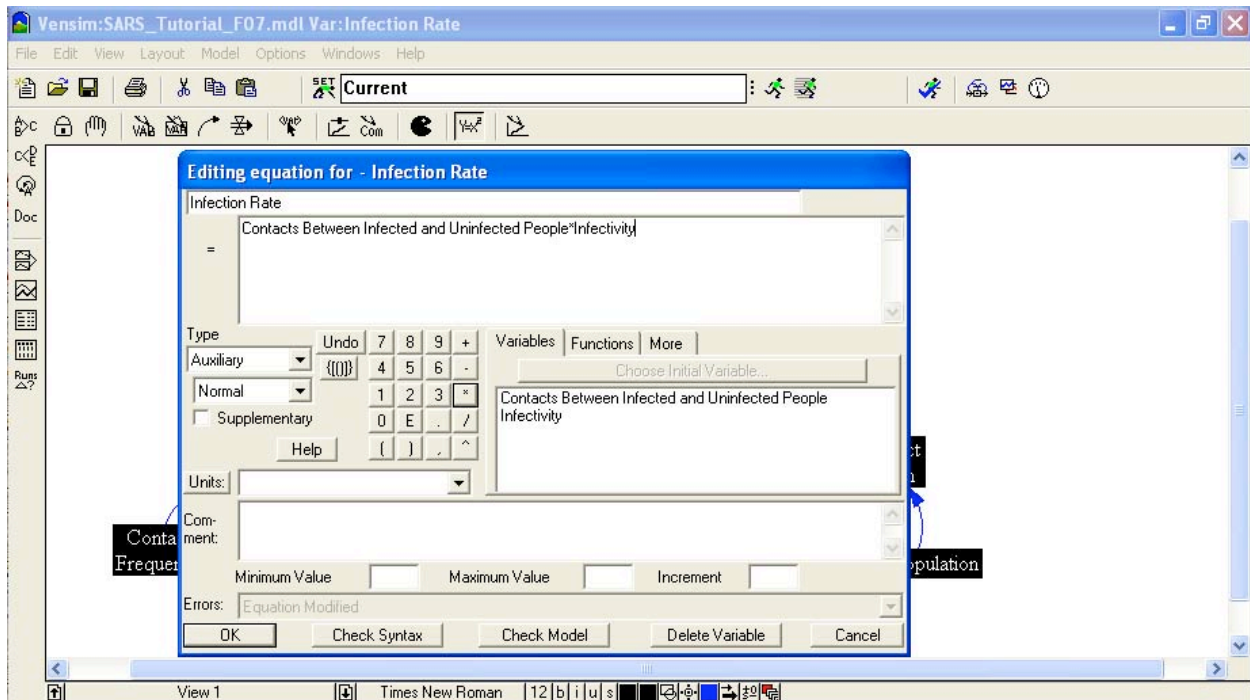


Vensim software copyright 1996-2011 Ventana Systems, Inc. Used with permission.

Good modeling practice requires that each equation in a model have three elements: (1) *the equation itself*; (2) *specified units of measure*; and (3) *complete documentation*.

You enter the equation in the box to the right of the = sign. You enter the units of measure in the text field to the right of the word **Units**. Equation documentation is entered in the box to the right of the word **Comment**.

Write the equation first. Begin by clicking in the box to the right of the equal sign. Now, look in the box in the middle right. Here Vensim has provided you with the two variables that your model diagram shows as influencing the infection rate, **Contacts Between Infected and Uninfected People** and **Infectivity**. A simple and plausible formulation for the infection rate is simply to multiply the number of contacts by the infectivity (which, by definition, is the fraction of contacts that result in transmission of the disease, or equivalently, the conditional probability of transmission given a contact between an infected and susceptible individual). To enter this equation first click on the **Contacts Between Infected and Uninfected People** variable, which will now appear in the equation box to the right of the = sign. Now either click on or type “*” to represent multiplication. Complete the equation by clicking on **Infectivity**. Your dialog box should now look like this:



Vensim software copyright 1996-2011 Ventana Systems, Inc. Used with permission.

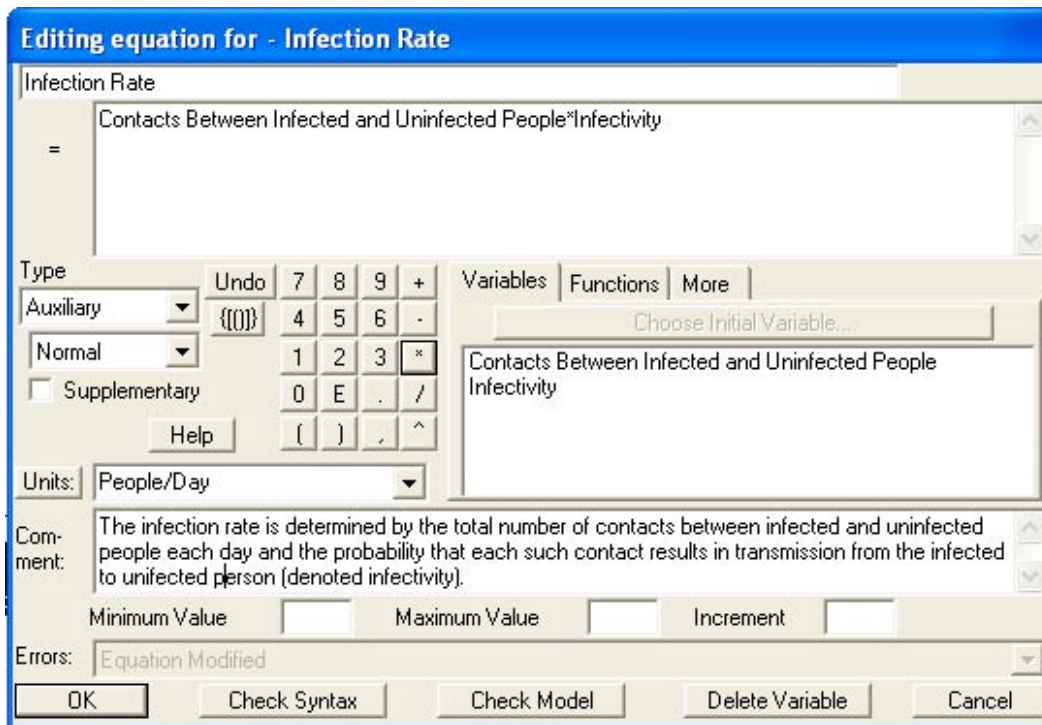
Now fill in the units. **Infection Rate** represents flow of people through time, so the appropriate unit of measure for this equation is people/time unit. Since you already chose to run the model in time steps of 1 day, the appropriate unit is people/day. Type **People/Day** in the units field.

(The next time you specify the units for a variable in this model, People/Day will appear in the units pull-down menu. You can click on the arrowhead to the right of the units field to see units already specified for other variables in the model, and then select units you've already defined from that list as appropriate.)

Finally, provide a description of this equation in the comment field. A good comment will be brief but also give the reader some sense of the logic behind the equation as well as state the key assumptions. For example, one might write for this equation:

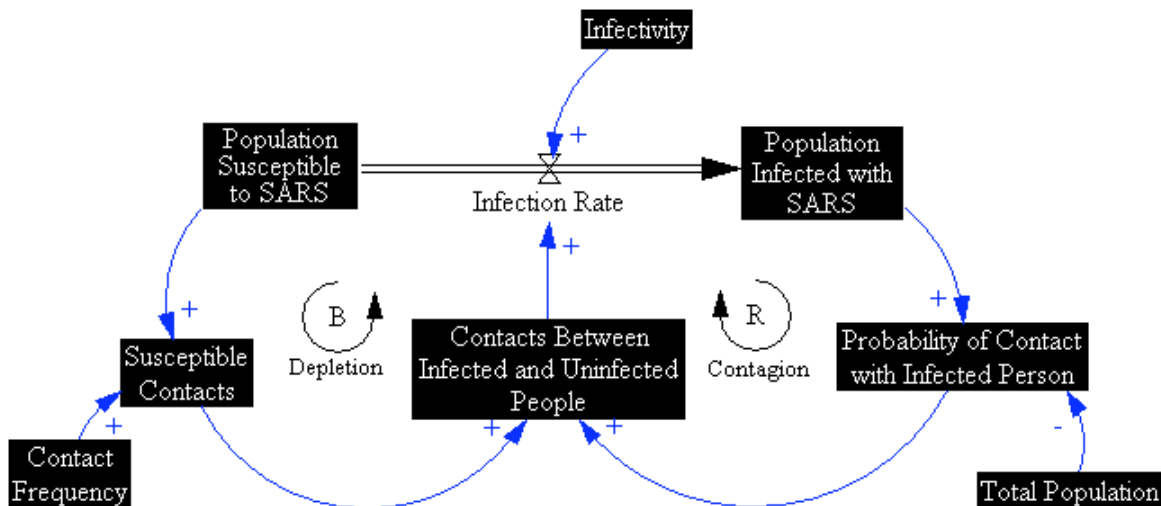
The infection rate is determined by the total number of contacts between infected and uninfected people each day and the probability that each such contact results in transmission from the infected to uninfected person (denoted infectivity).

Your dialog box should now look like this:



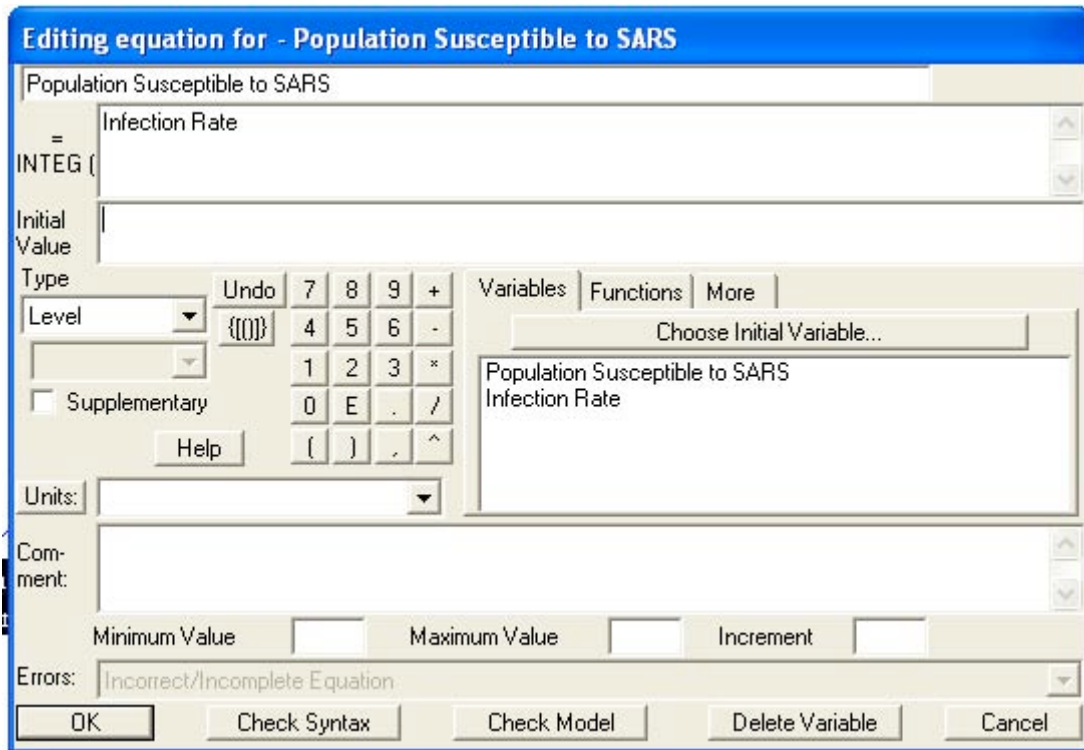
Vensim software copyright 1996-2011 Ventana Systems, Inc. Used with permission.

Click on **OK** or hit return and your diagram will look like this:



The **Infection Rate** is no longer highlighted, indicating that you have specified its equation.

Now write the equation for the stock of susceptible people. Begin by clicking on the susceptible stock. You should get a dialog box that looks like the following:



Vensim software copyright 1996-2011 Ventana Systems, Inc. Used with permission.

You enter the equation for the stock in the box to the right of the word **Integ**. **Integ** stands for “integrate” and simply means that the value of a stock at any moment in time is equal to the sum of all that stock’s inflows less all of its outflows from the start time of the simulation plus the initial value.

When you created the stock, flow, and feedback diagram, you made the **Infection Rate** an outflow from the stock of **People Susceptible to SARS**. VensimPLE captures this stock-flow relationship by showing the equation for the susceptible stock as **Integ(-Infection Rate)**. This means that at any point in time the susceptible stock is equal to its initial condition less all the people that have flowed out of it up to that point in time.

Now we need to specify the initial value for the susceptible stock. You enter it in the box just to the right of the words **Initial Value**. Initial conditions can either be numbers or other variables in a model. It is usually better to initialize a stock in terms of a parameter, because you can then alter that parameter easily for sensitivity and policy testing. In this case it is a good idea to

* Stock = INTEG(Inflow – Outflow) with an Initial Value of Stock(t_0) is equivalent to:

$$Stock(T) = \int_{t_0}^T (Inflow(t) - Outflow(t)) dt + Stock(t_0)$$

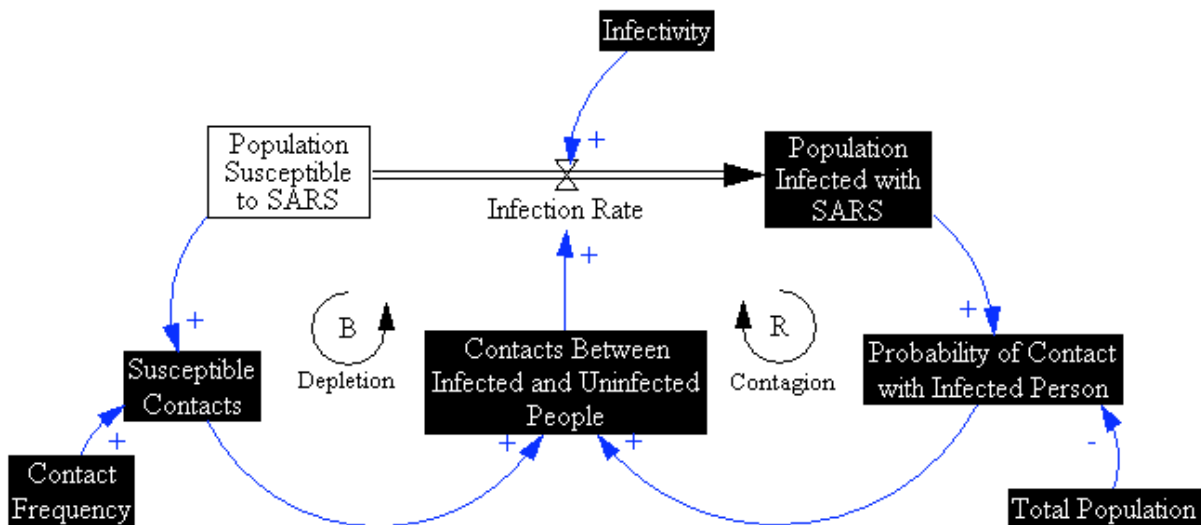
initialize the stock based on another variable, specifically the **Total Population**. Enter **Total Population** in the **Initial Value** box.

Your equation for the **Population Susceptible to SARS** indicates that the susceptible population is equal to the initial value at the start of the simulation less all of those people who have subsequently contracted the disease.

You still need to specify the unit of measure and document your equation in the comment field. The **Population Susceptible to SARS** is a stock and is measured in units of people. A sample comment for **Population Susceptible to SARS** is:

The Population Susceptible to SARS is the equal to the population susceptible prior to the onset of the disease less all of those that have contracted it. It is initialized to the Total Population, which assumes that all individuals are initially susceptible (no prior natural or vaccine-conferred immunity).

After completing the dialog box, click **OK** or press return. Your model diagram should now look like this:



Following a similar process to the one outlined above, you should now be able to specify the equations for the remaining endogenous variables, **Susceptible Contacts**, **Contacts Between Infected and Uninfected People**, and **Probability of Contact with Infected Person**:

$$\text{Susceptible Contacts} = \text{Population Susceptible to SARS} * \text{Contact Frequency}$$

(People/Day)

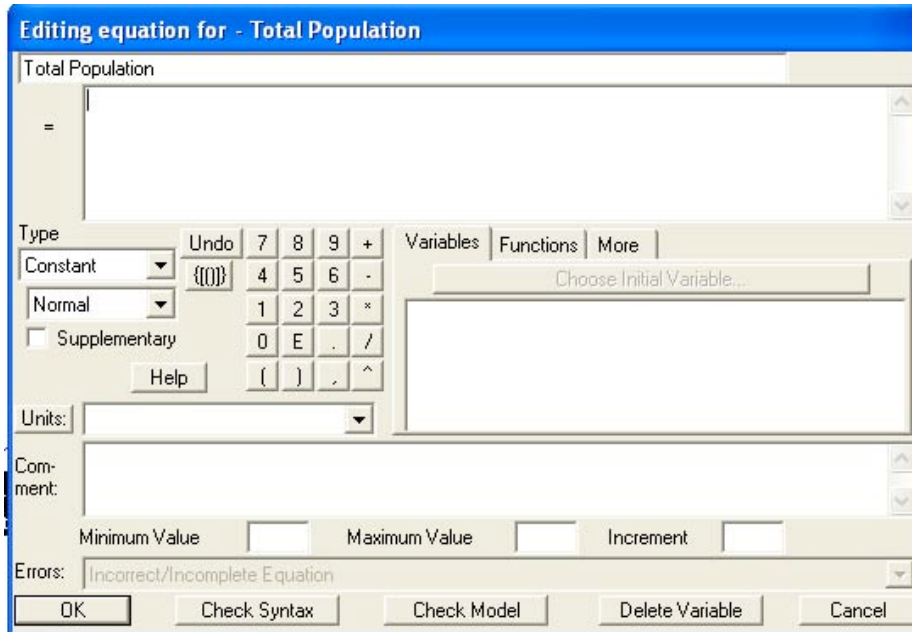
$$\text{Contacts Between Infected and Uninfected People} = \text{Susceptible Contacts} * \text{Probability of Contact with Infected Person}$$

(People/Day)

$$\text{Probability of Contact with Infected Person} = \text{Population Infected with SARS} / \text{Total Population}$$

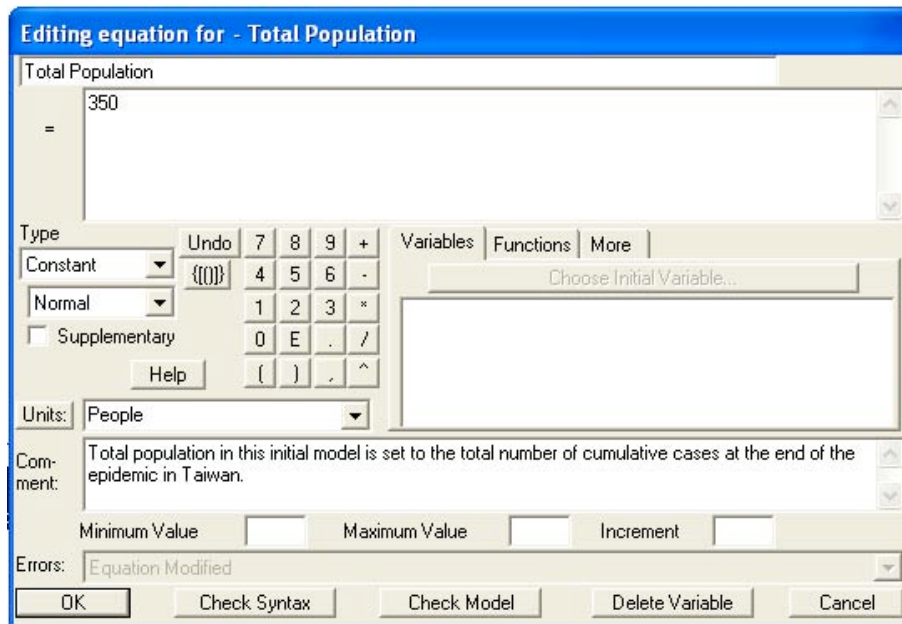
(Dimensionless)

Finally, you need to specify equations for the exogenous variables, **Infectivity**, **Contact Frequency**, and **Total Population**. Begin with **Total Population**. Clicking on it should bring up the following dialog box:



Vensim software copyright 1996-2011 Ventana Systems, Inc. Used with permission.



Note first that there are no inputs to Total Population, just as you specified in the diagram, indicating that Total Population will be a constant. The data suggest that in Taiwan, the cumulative number of cases eventually rose to approximately 350, so start with that value for **Total Population** (you may want to revisit this assumption later). You also need to specify units of measure and provide a comment. When you are finished your dialog box will look like this:

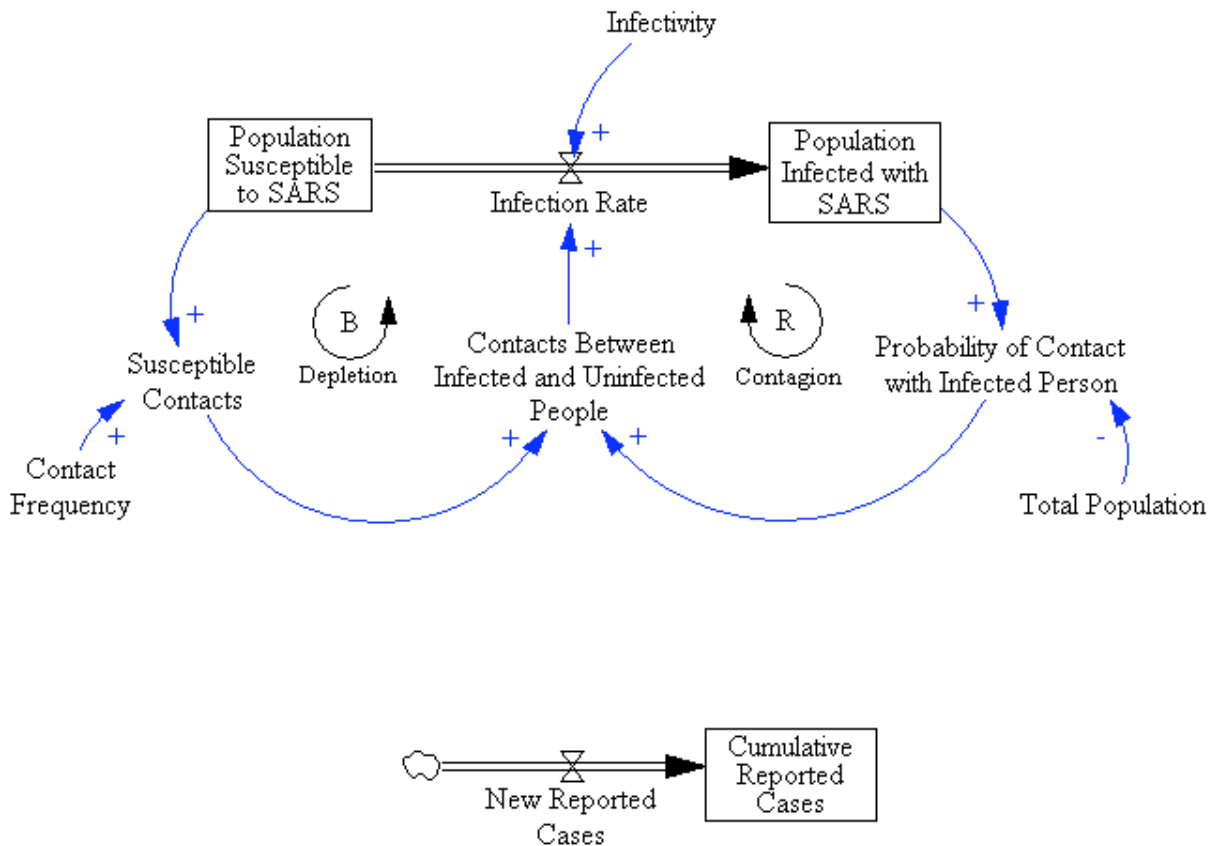


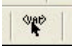
Vensim software copyright 1996-2011 Ventana Systems, Inc. Used with permission.

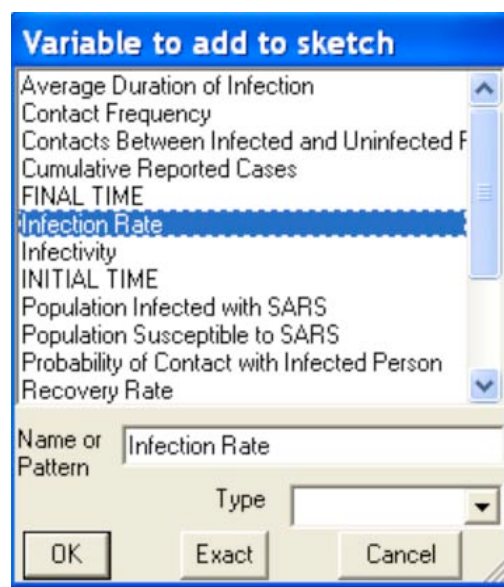
Following this example, complete the equations for the other two parameters, **Contact Frequency** (people/day) and **Infectivity** (dimensionless). You do not have hard data for the values of these parameters so you will need to use your judgment to determine appropriate values. Later we will discuss how you can test the model's sensitivity to these assumptions. Finally, set an appropriate initial value, units, and comment for the **Population Infected with SARS**.

Now we will add the structure to allow us to compare the behavior of the model to the data for the SARS epidemic in Taiwan. We have created a file containing these data (the file SARSDATA.vdf). To compare the model against the data we need to create variables that correspond to the incidence of new cases, denoted **New Reported Cases**, and the prevalence of SARS, denoted **Cumulative Reported Cases**. Because the data file uses these variable names you must create variables in your model with exactly the same names.

Using the stock variable tool  create a new stock called **Cumulative Reported Cases**, placed a few inches below the current structure. Then use the flow tool  to create an inflow to that stock called **New Reported Cases**. This structure captures the relationship between incidence and prevalence: As new cases are reported the cumulative number of cases increases. Your diagram should look like this:


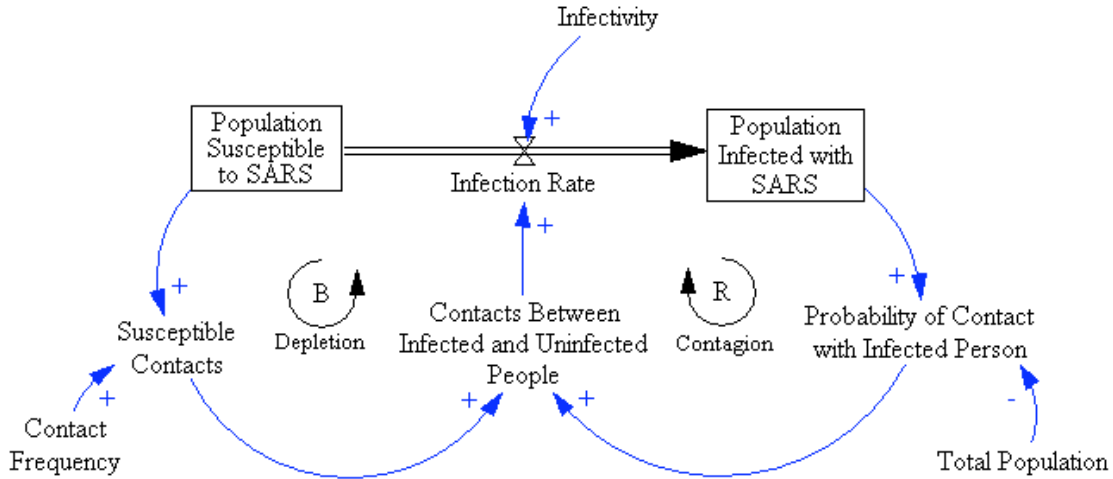


Now you need to define **New Reported Cases**. For simplicity we will assume that new cases (the infection rate) are immediately and correctly diagnosed and reported. In reality, of course, there is a delay between the infection rate, diagnosis, and reporting through the public health system; in addition, diagnosis is imperfect (some actual SARS cases were missed; other respiratory diseases might be incorrectly classified as SARS). Modeling both diagnostic errors and reporting delays is straightforward, but we will not do so in this tutorial. To define **New Reported Cases** as equal to the infection rate you could simply draw a causal arrow from the **Infection Rate** to **New Reported Cases**. However, doing so would clutter the diagram. Instead, we will use a *Shadow Variable*. Select the Shadow Variable tool  from the top toolbar. Click about an inch to the left of **New Reported Cases**. The following dialog will pop up, listing all variables in the model. Select the **Infection Rate**, then click OK.



Vensim software copyright 1996-2011 Ventana Systems, Inc. Used with permission.

Vensim creates a copy (Shadow) of the **Infection Rate**. Shadow variables are gray and placed within braces: **<Infection Rate>**. Your diagram should now look like this:



Shadow variables enable you to keep your model diagrams clean and readable. Now draw a causal arrow from the shadow **<Infection Rate>** to **New Reported Cases**:



Use the equation tool to define **New Reported Cases** as equal to the **Infection Rate**. The units for **New Reported Cases** are People/Day. Add an appropriate comment (point out that you are assuming that new cases are reported immediately). Now define the stock of **Cumulative Reported Cases**. You will see that **Cumulative Reported Cases** is already defined as the integral (accumulation) of **New Reported Cases**. You need to supply the initial number of reported cases, which is zero. Specify the units as People and add an appropriate comment.

Note that in this initial model Cumulative Reported Cases will equal the actual population infected with SARS, because you have assumed that New Reported Cases equals the actual Infection Rate. Why then bother to define Cumulative Reported Cases separately from the Population Infected with SARS? There are two reasons: (1) They are conceptually different: the actual number of people infected with SARS differs from the number reported through the public health surveillance system (there may be delays and measurement errors you might want to capture later); (2) you might later alter the model structure so that the equivalence of the two stocks no longer holds (for example, capturing the rates of recovery and mortality from SARS).

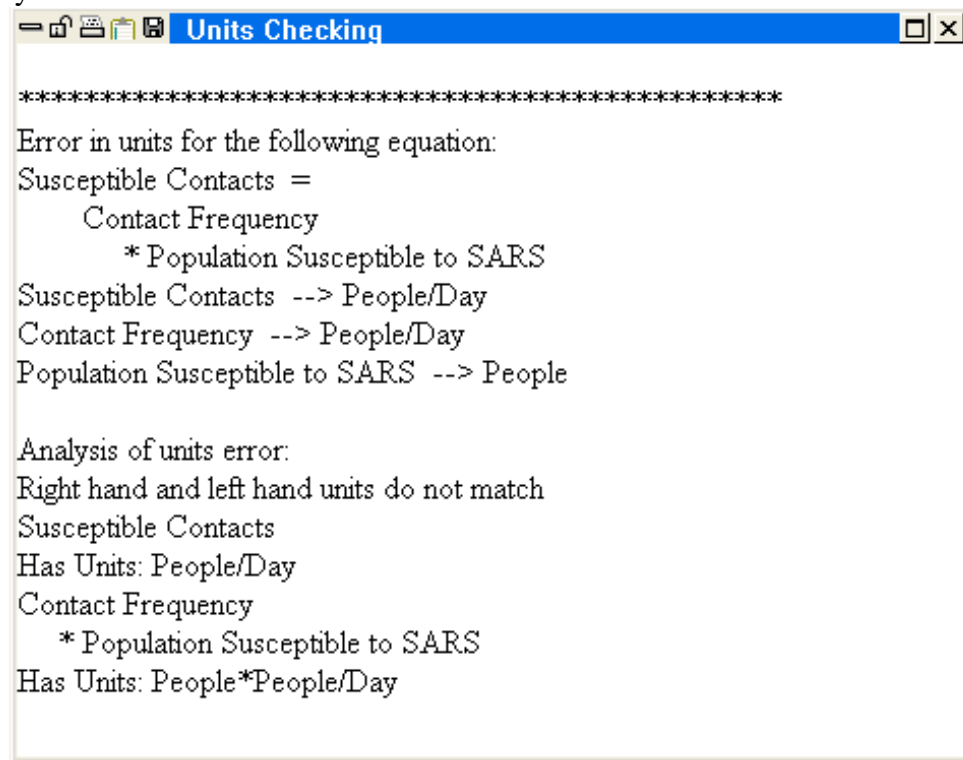
IV. Using the Model Structure Analysis Tools

You are now ready to analyze your model. VensimPLE provides several tools for analyzing and understanding the structure of your model.

By far the most important of these is the unit-checking function. To access this tool, go to the **Model** menu and select **Units Check**. An important feature of any system dynamics equation is *dimensional consistency*. This is just a fancy way of saying that the units of measure must be the same on both the left and right sides of the equation. For example, if you have followed the instructions in the tutorial so far and you run the units check, you will get the following message:




Followed by:



Vensim software copyright 1996-2011 Ventana Systems, Inc. Used with permission.

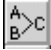


The problem is that the equation for **Susceptible Contacts** is not dimensionally consistent: The right and left sides of the equation have different units. **Susceptible Contacts** is measured in contacts per day, but, in the current formulation, multiplying the **Population Susceptible to SARS** by the **Contact Frequency** yields a value with units of people*people/day.


The cause of the problem is that unit of measure for **Contact Frequency** is incorrect. The **Contact Frequency** represents the number of contacts *each* person makes per day, not the total number of contacts (which is captured by the variable **Susceptible Contacts**). To fix this problem, change the units of measure for contact frequency to be people/person/day. After you do this, run the units check again. When you are done click **OK**. Now when you run the units check you should get the following:



Vensim software copyright 1996-2011 Ventana Systems, Inc. Used with permission.


The units in your model now balance.

The other tools are also useful. The  and the  buttons create “causes” and “uses” trees for the selected variable. To use these tools, you need to first “select” a variable. Do this by first clicking on the  button and then clicking on the variable you wish to select. Clicking on a variable selects that variable to the workbench. You can tell which variable is selected by looking at the top border of the VensimPLE window. If you select the variable **Infection Rate**, the top border will look like this.






Vensim software copyright 1996-2011 Ventana Systems, Inc. Used with permission.

Having selected **Infection Rate**, clicking on the two “causes” and “uses” buttons in sequence gives you:



and



The  tool identifies all the feedback loops in which the selected variable is a member, and the  button creates a documented listing of the equations in your model.

V. Simulating Your Model

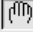
Once all the equations are entered and your model passes the dimensional consistency test, you are ready to simulate the model.


To specify the simulation control parameters, go to the

Before doing this, however, you must actually simulate the model so you have some behavior to analyze. And, since we have historical data against which to compare the behavior of your model, it is also necessary to load the historical data file.

To run a simulation, you first need to choose a name for this particular model run. It is helpful to choose names that suggest some idea of what is being tested rather than using names like SIM1, SIM2, etc. Since this is the base case run for your model, you might choose to call the run **BASE**.^{*} The name of the simulation is entered in the white box in the analysis tool bar. The default name for any simulation run is **Current**. Change this to **Base**.


To run a simulation, click the  button and your model will start simulating.


Once the simulation run is completed, you can look at the results of your simulation. VensimPLE provides many tools with which to view simulation output. The most basic of these is the graph. To create a graph of the **Cumulative Reported Cases**, first click on the  tool and then click on the desired variable.


To see a strip graph, click on the  button in the left toolbar. You will then see the output of your model, compared against the data. The behavior of your model will depend on the parameters you have chosen. My parameter values and initial conditions generated the following graph:


Vensim software copyright 1996-2011 Ventana Systems, Inc. Used with permission.

^{*} Advanced Tip: VensimPLE also offers you the choice of two numerical integration methods, Euler and Runge-Kutta 4. Runge-Kutta 4 is a more accurate simulation method, but it is also more computationally intensive. Generally it is better to use the Euler method and change only if you believe you are seeing integration error. Read *Business Dynamics*, Appendix A, for further details.



Besides the graph, VensimPLE provides many other ways to examine simulation output. The  button displays a strip graph of the currently selected variable, along with graphs of all the variables that determine the selected variable (the causes). Clicking this button gives you:

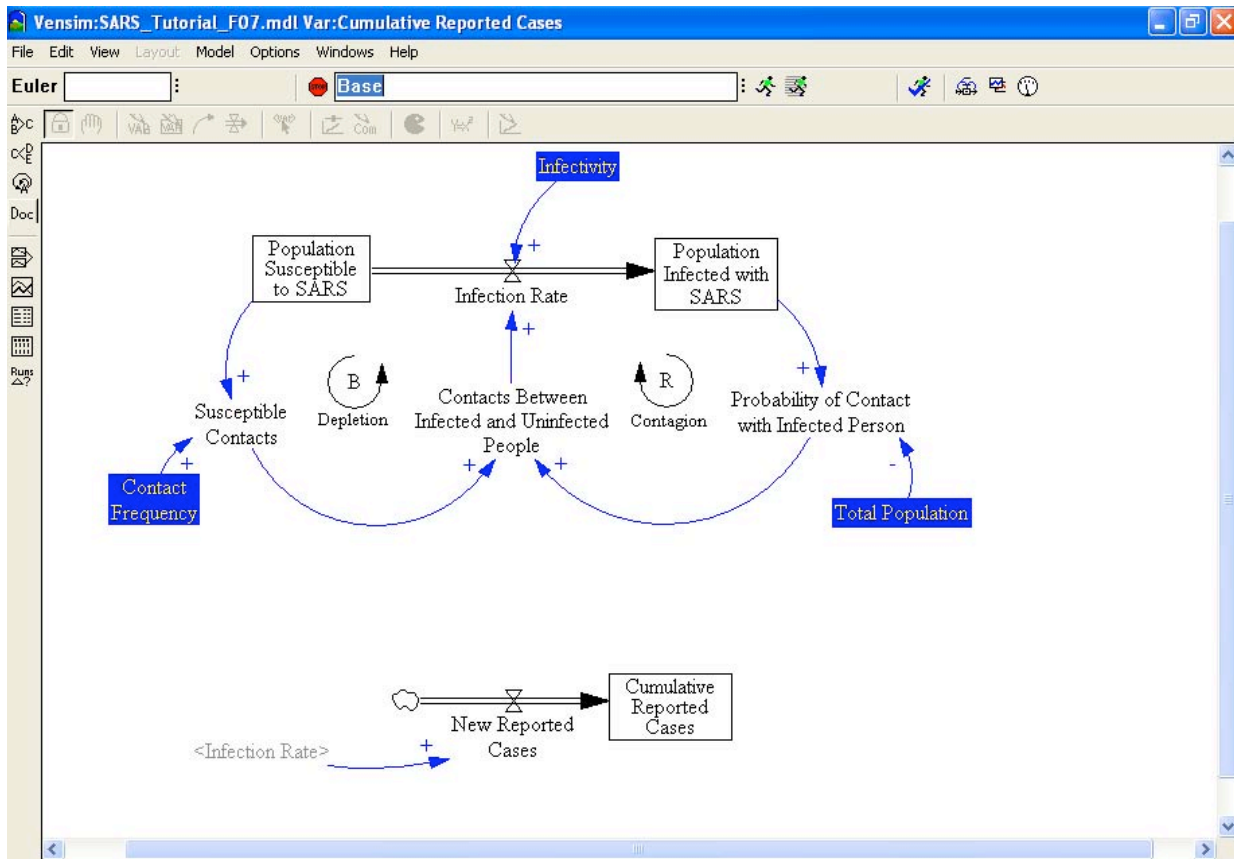


VensimPLE also can present the output in the form of a table rather than a graph. To see a table simply click on the  button.

Having examined the output of this simulation, you may wish to run additional simulations under different assumptions. For example, what would happen if you cut the contact frequency in half?


One way to do this is simply to change the model equation in the variable's equation box. This is time-consuming and error-prone, however, since you must remember to go back and change the parameter back to its original value before testing another assumption. If you change a variable's formulation in its equation box, it remains as you specify until you change it again. So, if you want to alter a variable's value only once, to observe its effect on the simulation, and then have the variable revert to its original value, you can use the Parameter Changes button, accessible on the analysis tool bar.

Test the effect of a change in the **Contact Frequency** by clicking the  button. Your screen should now look like this:




Vensim software copyright 1996-2011 Ventana Systems, Inc. Used with permission.

To change the contact frequency for this simulation, click on **Contact Frequency** and its current value will appear in a box. Type in a smaller value (such as a 50% reduction) and hit **return**. Before simulating, make sure you the name of the simulation so that you can compare the two different runs. Select a name that helps describe the changes you made, for example, “Low Contact Frequency”.


To run the simulation with your new value for the contact frequency just click on the  button. Note that you have changed the contact frequency for that run only, and have not changed it in the underlying model equations.

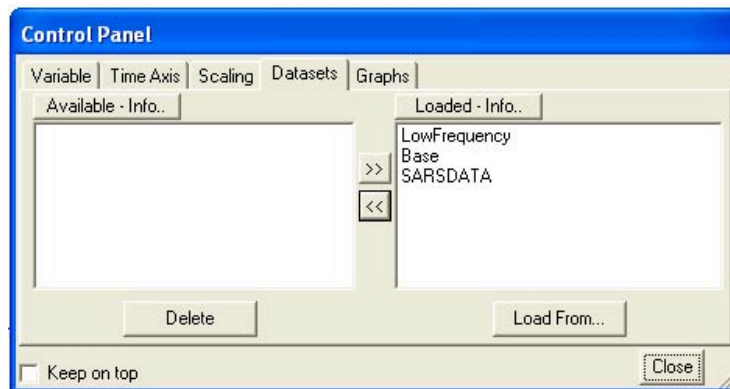
Now when you select any of the model output tools, they show the results from both simulations (along with the data). For example, the graph tool displays:



Vensim software copyright 1996-2011 Ventana Systems, Inc. Used with permission.

If you do not wish to see the previous (base) run displayed with the new simulation run, then

click on the  button. This is the **Control Panel**. Click on the **Datasets** tab and the dialog box will display the three data sets you have created so far:




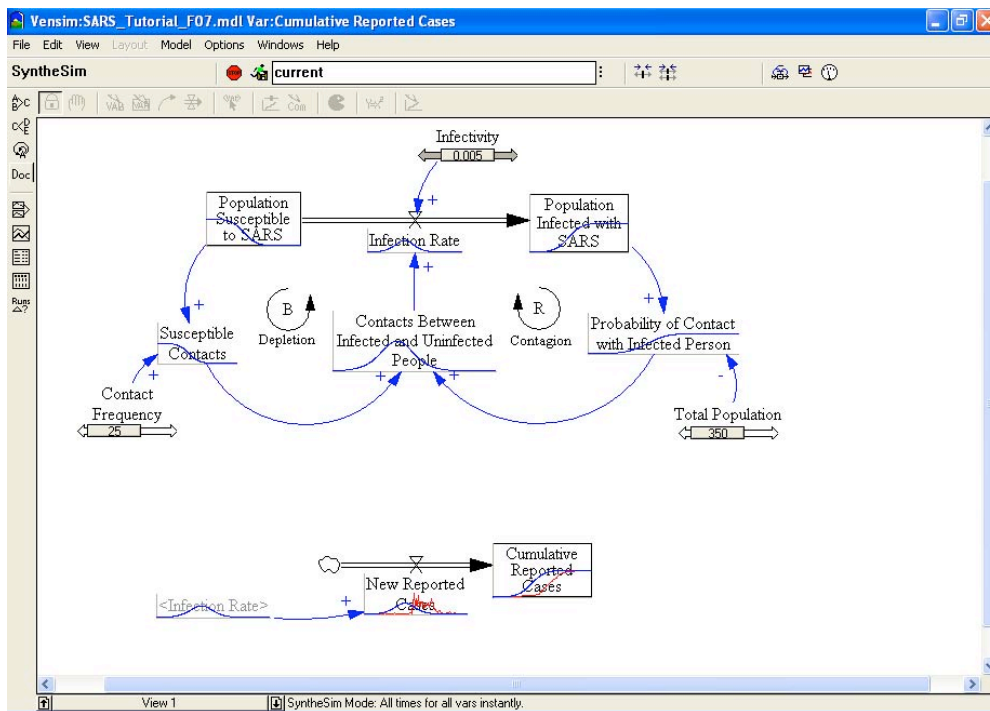
Vensim software copyright 1996-2011 Ventana Systems, Inc. Used with permission.

To de-select a run, simply double-click on it. The run you de-select will now appear on the left side of the dialog, in the “available” column. It remains stored on your hard disk and can be re-loaded again any time you wish to display it. Clicking on a run in the “loaded” column will also move it to the top of the list of loaded runs so you can control the order in which runs are listed in output such as graphs.

The control panel also allows you to select a given variable (choose the **Variable** tab). Modify the horizontal axis of your graphs (the **Time Axis** tab), and create custom graphs (the **Graph** tab). See the help menu for more information on all of these features.

Vensim also offers another simulation mode, called SyntheSim, to help you analyze your model.

To activate Synthesim click on the  button at the top of your screen. When you do this, SyntheSim is activated and your screen should now look like this:



Vensim software copyright 1996-2011 Ventana Systems, Inc. Used with permission.

All the constants in your model will now appear with sliders while the other variables will have small graphs superimposed on them showing the behavior of each variable under the current slider settings. With Synthesim activated you can now change any of the constants in your model using the sliders and Vensim will automatically re-simulate the model and provide updated output.

For small models such as this one, the response is almost instantaneous. Synthesim can be extremely useful in developing intuition as to how your model’s behavior changes as its underlying parameters are adjusted. You may wish to use synthesim now to find parameter values that provide a good fit of your model output to the data for cumulative reported cases.

MIT OpenCourseWare
<http://ocw.mit.edu>

15.872 System Dynamics II
Fall 2010

For information about citing these materials or our Terms of Use, visit: <http://ocw.mit.edu/terms>.