

Using LLMs to Support Teachers Infusing System Dynamics Concepts in High School Curricula
By Diana M. Fisher, Portland State University, fisherd@pdx.edu

Abstract

Large language models (LLMs) bring exciting support capabilities to the entire spectrum of education. The focus of this paper is how LLMs could support teachers who are trying to infuse SD modeling in their instruction. Three examples of (free versions of) LLM features are given. The first example uses an LLM to help describe the sequence of activities, and questions to ask students at each step, when a teacher wants to include an SD modeling task to study a certain topic in a high school curriculum. Then this instructional sequence is regenerated for students who might learn a little about the topic in middle school. Secondly, LLM access is now included in all versions of the Stella SD modeling software, including free Stella Online. Within the software code from a textual description of a dynamic problem can be automatically translated (via the build feature) into a stock-flow model. Several examples are given. Moreover, Stella provides a “discuss” feature (Seldon) that can take the Loops That Matter analysis and produce a detailed description of why the model behaves the way it does. Finally, LLMs can provide targeted feedback for student open ended question responses. Two different approaches are shown. LLMs can provide support for some of the most difficult planning and assessing that tends to accompany infusion of SD analysis in high school instruction.

Key Words: LLMs, artificial intelligence, AI , system dynamics modeling, providing feedback, text to SD models

Introduction

Like the graph of an exponential growth function where there appears to be very little change at the beginning of the time axis then shortly the function grows in an explosive fashion, artificial intelligence (AI) exploded on the world scene in November 2022 with the release of OpenAI’s ChatGPT. Yet work on developing adaptive machine learning had taken decades to get to its 2022 capability. Beginning with the Turing test in 1950, the goal was to have computing machine responses indistinguishable from human responses. Work to improve computer learning had steadily progressed, regressed, remained seemingly dormant and progressed some more, largely due to the irregularity of funding (due to early exaggerated promises of AI capabilities) and limitations of computer hardware (Mitchell, 2019). In the 1980s and 1990s computers relied on primarily “rules-based/brute-force” algorithms (expert systems) to try to compete against human intelligence in tasks such as winning a chess game, but this approach was not adaptable across applications (Mitchell, 2019). It was not until the refinement of multi-layered neural network configurations and back propagation that machine learning started to have the structure necessary to learn from its mistakes on a large scale (Li, 2023). The neural network deep learning approach and the digitization of vast quantities of information and images upon which to test recursive/adaptable learning overcame the hurdles needed to bring us to the Large Language Models (LLMs) that emerged in 2022 (Li, 2023).

Now, AI is progressing faster than anticipated. There are some LLM test benchmarks called ARC-AGI (Abstraction and Reasoning Corpus for Artificial General Intelligence) that are

designed to test an AI's ability to solve difficult problems it has not seen before. There have been incredible improvements in OpenAI ChatGPT's capability shown on this test with ChatGPT4o scoring just 5% on the ARC-AGI assessment in early 2024 but subsequent versions scoring 88% by December 2024 (Jones, 2025, Klein, 2025), a feat previously not expected to happen for at least several years.

It seems hard to imagine how precollege education could catch up, adapting curricular innovations incorporating the 2022 LLM versions of AI, let alone prepare for the versions that will emerge in 2025. Shah (2023) identified recommended curricular adaptations of AI for precollege teachers. Some adaptations include:

- creating lesson plans adaptable to more differentiated (ability/cultural/native spoken language) student groups in a class, including making more student-centered, active lessons,
- creating more diverse assessment questions, including creating answers and rubrics
- providing individual feedback for students on open-ended questions,
- allowing teachers to provide students with one-on-one tutoring (Khanmigo),
- gamifying instruction for certain educational concepts, and
- creating slide decks for presentations.

Precollege education tends to be slow to adapt to significant changes in instructional methods (Schank, 2004; Stigler & Hiebert, 2009), but student access to free AI apps will, no doubt, require teachers to adapt their instruction more quickly if not only to allow teachers to provide accurate assessments of student learning. Students will learn to harness AI quickly.

AI can assist precollege teachers in their efforts to infuse System Dynamics (SD) modeling into their curriculum. This paper identifies three capabilities of AI that can assist a (hypothetical) high school algebra teacher as she attempts to infuse SD in her instruction. The examples include:

- differentiating instructional lessons in an algebra class where students are to build an SD model of the spread of an epidemic (using free ChatGPT 4o),
- steps to help the instructor describe, with text, a simple dynamic system, and have the LLM built into the Stella software produce the stock-flow model for the system. This feature is also available in (free) Stella Online. Moreover, the LLM in Stella can use the Loops That Matter tool to provide detailed explanations of the behavior of the model, and
- producing individualized feedback for student open-ended explanations for SD model behavior for an epidemic model students have built and executed (using free ChatGPT4o).

Some LLM Support for Precollege Teachers using SD

I. Designing Differentiated SD Enhanced Lessons in High School Algebra (Using free version of ChatGPT4o)

Training ChatGPT4o

To train ChatGPT with a basic understanding of the tools and concepts for educational applications of System Dynamics a subset (about one third) of the information contained in

Emre Göktepe's book "Systems Thinking in Education"¹ was uploaded into the free version of ChatGPT4o. (The free version of ChatGPT4o restricts the total number of pages of text that can be uploaded as training material so this subset of Göktepe's book was broken into two file. The files can be downloaded at:

https://bit.ly/Part1_TeachAI_SD_Concepts and https://bit.ly/Part2_TeachAI_SD_Concepts.)

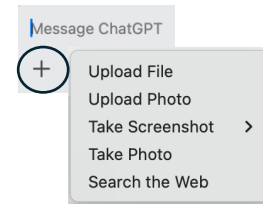


Figure 1: How to locate the icon to press to allow the user to upload a file into the free ChatGPT4o.

AI can be useful in describing a potential lesson for a given application. For example, given the logistic behavior for the growth of people who become infected with a particular disease (in a closed population), an algebra teacher might want to devise a lesson to help students determine the important variables to consider in this system. Note: It is assumed the students have already built and analyzed small system scenarios to replicate linear, parabolic, exponential and convergent behavior, by building SD models by hand, before this assignment is given.

High School Algebra Class

Prompt: I am a high school algebra teacher who knows how to build system dynamics (stock-flow) models. I want to eventually have students build the stock-flow diagram for the spread of an epidemic but I want you to help me determine a sequence of activities to use with my 16 year old students to help them determine the important components needed to understand the dynamics of the simple spread of a disease in a closed population and the discussion questions to use in this process.

(free) ChatGPT4o Response: (See complete response in Appendix 1)

1. Introduction to Epidemic Dynamics

- *Objective:* Familiarize students with basic concepts such as “healthy,” “infected,” and “recovered.”
- *Activity:* Show a simple simulation (e.g., a video or animation) of disease spread in a closed population.
- *Discussion Questions:*
 - What factors might influence the spread of a disease?
 - How does the size of the population affect the spread?
 - What happens when individuals recover? What about immunity?

2. Behavior Over Time Graphs

- *Objective:* Help students think about how the number of healthy, infected, and recovered individuals changes over time.
- *Activity:* Students sketch behavior over time graphs for each group based on observations or given scenarios.
- *Discussion Questions:*
 - How do the numbers of infected individuals change at the start of an outbreak? Why?

¹ A digital copy of this book is available for free on the SDS Precollege Education webpage: <https://systemdynamics.org/special-interest-groups/pre-college-education-sig/>

- What do you notice about the healthy population as infections increase?
- How does the recovered group grow over time?

3. Introducing Stock and Flow Concepts

- *Objective:* Distinguish between stocks (e.g., healthy, infected, recovered) and flows (e.g., infection, recovery).
- *Activity:* Define stock and flow concepts and brainstorm variables for an epidemic model.
- *Discussion Questions:*
 - What is the “stock” of healthy individuals, and how can it decrease or increase?
 - What factors (flows) might increase or decrease the number of infected individuals?
 - Could there be inflows or outflows that influence recovery?

4. Building a Simple Stock-Flow Diagram

...

5. Exploring Feedback Loops

...

6. Experimenting with Model Variables

...

7. Policy and Real-World Reflection

...

Final Project

- **Objective:** Have students independently construct and present a stock-flow model for a specific disease scenario.

AI can provide the same basic lesson concepts but at different levels of sophistication for students, helping high school teachers provide more differentiated instruction. Or it can provide the same basic lesson concepts for different ages of children, as is demonstrated by using the same idea/prompt regarding the spread of an epidemic, but for a 7th grade math class. It is assumed that the students have already built and analyzed small system scenarios, by building SD models, to replicate linear, exponential, and convergent behavior.

Middle School Math Class

Prompt: I am a middle school math teacher who knows how to build system dynamics (stock-flow) models. I want to eventually have students build the stock-flow diagram for the spread of an epidemic but I want you to help me determine a sequence of activities to use with my 13 year old students to help them determine the important components needed to understand the dynamics of the simple spread of a disease in a closed population and the discussion questions to use in this process.

(free) ChatGPT4o Response: (See complete response in Appendix 2)

1. Introducing the Problem: How Do Diseases Spread?

- *Objective:* Engage curiosity and connect to real-world situations.
- *Activity:* Share a relatable story or video showing disease spread (e.g., a common cold in a classroom or COVID-19).
- *Discussion Questions:*

- How do people get sick from one another?
- What might stop the disease from spreading to everyone?
- What does “recover” mean in this context?

2. Visualizing the Disease Spread

- *Objective:* Observe and record how groups of individuals (healthy, sick, recovered) change.
- *Activity:* Conduct a simple hands-on simulation:
 - Use a dice-roll or colored tokens to simulate disease transmission over several “days.”
 - Track the number of healthy, sick, and recovered individuals each day.
- *Discussion Questions:*
 - What happens to the number of sick people as the disease spreads?
 - Why do healthy people decrease over time?
 - How does recovery affect the spread?

3. Building Behavior Over Time Graphs

- *Objective:* Represent the changes in the healthy, sick, and recovered populations visually.
- *Activity:* Students plot data from the simulation to create behavior-over-time graphs for each group.
- *Discussion Questions:*
 - How do the numbers of sick and recovered individuals change over time?
 - What do you notice about the healthy population as the disease spreads?
 - Why does the graph eventually level off?

...

8. Culminating Project: Create and Explain

- *Objective:* Synthesize understanding into a student-built model.
- *Activity:*
 - Students design their own stock-flow diagram of an epidemic.
 - Present the diagram, explaining how stocks and flows interact.
- *Discussion Questions:*
 - Why did you place the stocks and flows where you did?
 - How does your model explain what might happen during an outbreak?

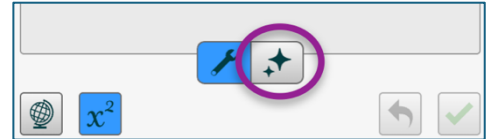
The teacher can ask for more details on any or all parts of the outlined plan, especially adding suggested answers to some of the discussion questions (for teachers who are more novice users of SD). The teacher can ask AI to produce handouts for students for any part of these plans.

For precollege teachers who are relatively new to SD it may be difficult to come up with stock-flow models to use for initial guided model-building lessons. Or for teachers who want to help students get started designing their own (relatively simple) SD models based on a news article they have read, the following sequence of steps shows how AI can provide assistance. Note: both teachers and students need to have enough experience building small SD models to be able to debug any problems in the models that AI produces.

II. From Text Descriptions of a Simple Dynamic System to LLM Computer Generated Stock-Flow Models (using the LLM feature of the free Stella Online SD software)

In July 2025 isee systems will release Stella 4.0 which will contain AI (LLMs) for both building stock-flow simulation models from a text description of a systems problem and a specially trained chatbot extension called Seldon for discussing problems or providing analysis of SD model behavior. A quick summary of each AI based support tool will be provided here. In keeping with the effort to discuss only free AI options I will talk about these features in the free Stella Online SD modeling software (exchange.iseesystems.com)

In Stella Online you can locate the AI property panel by clicking on the button with the two stars at the bottom of the Model Settings property panel.



A. AI Build: The “build” segment of the Virtual Assistant (AI) property panel in all the Stella SD software products produced by isee systems looks like this:

- Be sure to select the “Build” choice for the mode;
- Type a title or one sentence description of the problem in the Problem Statement section;
- Copy and paste any background information the LLM might need to understand in the Background Knowledge section (isee systems has already built in information about SD models and basic SD concepts into the AI tools);
- Type in your prompt in the Virtual Assistant section;
- The AI reasoning about your model will be displayed in the AI Reasoning section, along with a copy of the prompt you used.

Isee systems has preselected the LLM that will be used to build the stock-flow diagrams. At this time there is no charge for this access to the LLM they have chosen to provide. However, if you have a commercial version of a Google LLM you can provide the API key for it at the top of the second section of this panel.

Figure 2: The “Build” AI Virtual Assistant property panel in Stella Online.

Billy Schoenberg, the principal software engineer at isee systems, has written a paper “Building and Learning with Models Using AI” that he will present at the ISDC 2025 in Boston. It contains the rationale behind the “open source” design of the AI SD modelbuilding capability he created in Stella. The paper can give you more detail about how he taught the LLM to understand the core concepts in SD so the LLM could produce mostly correct SD models from the prompts describing a systems problem input by the user.

The examples that follow are simple. The more complicated the system, the more detail is needed in the prompt for AI to produce the code needed to allow Stella to produce the actual stock-flow model. Moreover, if the system is more complicated there is a greater chance that the model produced will contain some errors, so it is recommended that the teacher/student have basic knowledge of SD modeling to be able to debug the models.

Bank account models:

Prompt: I want you to create a system dynamics simulation of a checking account which has \$1000 in it and is earning an annual interest rate of 5% compounded monthly. Also, I would like you to have a monthly deposit of \$2000 for this account.

Initial model output

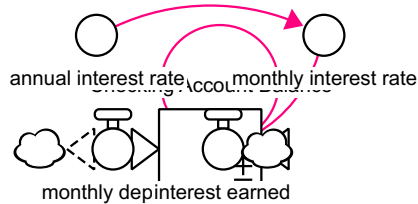


Figure 3: The compressed SD model produced by the LLM in Stella Online.

Reorganizing the Stock-Flow model diagram:

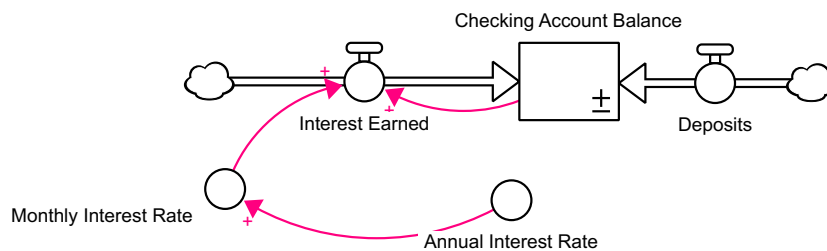


Figure 4: The reorganized SD model from Figure 3.

Notice that the LLM created a converter to translate the annual interest rate to a monthly interest rate. The diagram is created correctly and the values are correct in each icon, as are the units and documentation (for each icon), and the link polarity for each link.

A slightly more sophisticated banking model, extending the previous model

Prompt: I want you to create a system dynamics simulation that extends the previous checking account model. Send 10% of the amount in the checking account to a savings account. The savings account earns 10% annual interest compounded monthly. Also, I am saving to purchase a car. Add a car account and have 5% of savings transferred to the car account each month.

Unscrambled stock-flow output from Stella Online

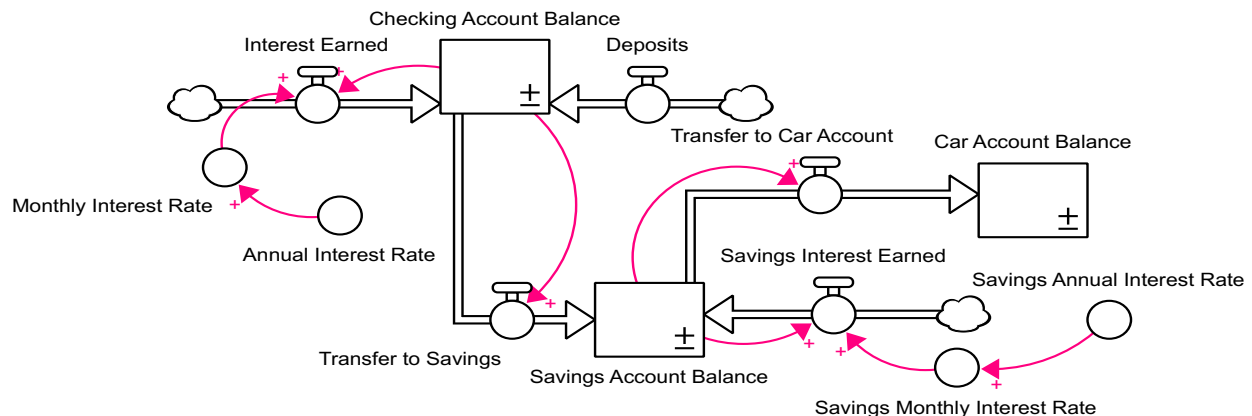


Figure 5: The expanded stock-flow model for the second bank account model produced by the LLM in Stella Online.

When spread out notice the diagram is created correctly, except for the fact that the Car Account Balance should contain some kind of interest added inflow (which is missing). (Interestingly, the first time I used this prompt Stella produced an interest inflow for the Car Account Balance of 8% annually. The second time the same thing happened but the interest was 3% annually.) The values are correct in each icon. The Savings stock was correctly set to an initial value of zero. Units and documentation for each icon were correct, as was each link polarity. Note: It is still important, especially as models get larger, to check that each icon is defined correctly, that units and documentation are correct, and that polarity on each link is correct. The software does a good job, but it is not perfect.

Non-Renewable Resource Model

A more complicated scenario (that I have my high school algebra students build from a guided model building lesson) involves a non-renewable resource (vegetation - that is not renewing quickly enough for the given simulation time) for a deer population. I had already built this model by hand many years ago. I wanted to see if I could describe the scenario in detail and see what the Stella Online LLM would produce.

Prompt: I want you to create a system dynamics simulation of a population (stock) that starts at 29 deer that has an inflow called births that is defined as the product of the population and the birth fraction (0.4) and has an outflow called deaths that is defined as the product of the population and the death fraction (0.1). This system dynamics simulation should also contain a vegetation (stock) that starts at 25000 veg units and only has an outflow called consumption. Each deer consumes one unit of vegetation per year. The deer consume the vegetation. A converter, which depends on the deer population stock and vegetation stock, is called effect of vegetation per deer on death fraction and should be constructed as a graphical converter. The graph in this graphical converter has a horizontal axis that goes from 0 to 1 and a vertical axis that goes from 1 to 15. Assume the graph in this graphical converter is concave down and connects the points (0,15) and (1,1). The effect of vegetation per deer on death fraction value should influence (multiply by) the death fraction. The AI should build the correct stock flow diagram with appropriate values within each icon. The simulation should run for 25 simulated years.

The unscrambled stock-flow model:

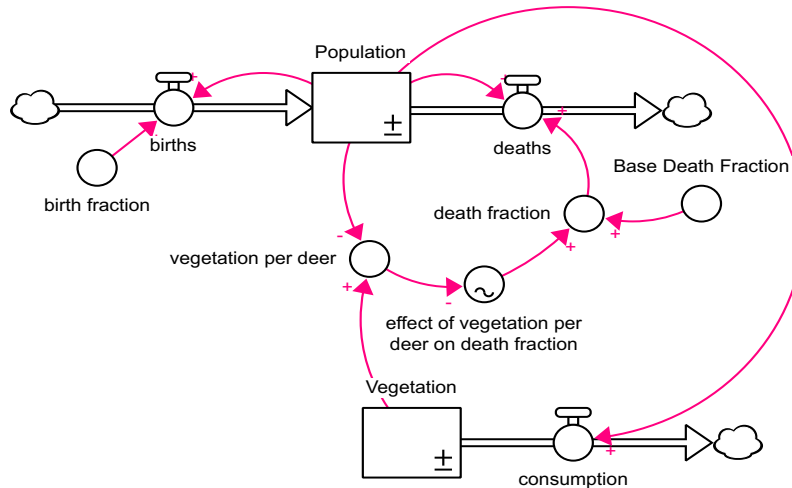


Figure 6: The expanded stock-flow model for the non-renewable resource model produced by the LLM in Stella Online.

The dimensionless multiplier “effect of ...” definition was mostly correct, but it set the wrong concavity for the graph. I fixed the concavity and created a graph of deer population. The graph was reasonably correct. The units, documentation, and polarity for each connector were correct.

Appendix 3 shows the AI Response provided by the Stella Online “Build” feature for this model.

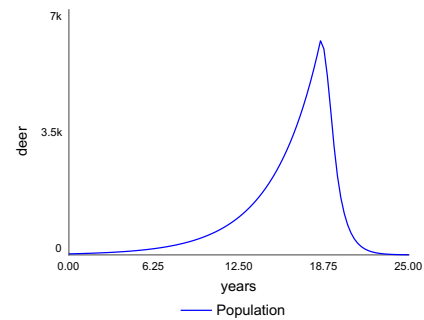


Figure 7: The graphical output for the deer population shown in Figure 6. This graph was set up by the teacher, not AI.

B. AI Discuss: Now let's take a moment to look at the "Discuss" section of the LLM capability in the Virtual Assistant (AI) property panel. This feature is called Seldon (after Hari Seldon of Foundation (book) fame, written by Asimov).

To access the conversation feature of the LLM in Stella Online you must select "Discuss" in the Mode section of the Virtual Assistant (AI) property panel.

At this time, Stella will allow you to choose among several LLMs to use for this discussion. That may change soon.

As with the "Build" mode you will want to create a title or one sentence description of the discussion you want to have with Seldon.

You can copy and paste additional background information into the "Background Knowledge" section.

Now, a really interesting part. If you had executed the model simulation with the Loops That Matter (LTM) feature turned on, Seldon will know all the feedback analysis that LTM produces. Moreover, if you created a table to include important variable output you can copy that output and paste it into the Behavioral Description section. Seldon will use the output from LTM and the tabular output to increase its understanding of your model behavior.

Then you can ask Seldon to explain the behavior of your model and you will have a much improved explanation with less chance of hallucination because the LLM is not doing pattern matching but using actual correct information for this analysis. The explanation from Seldon will be given in the AI Reasoning section of the panel.

I executed the non-renewable resource for the deer model, explained above, setting the $DT = 1$ for 25 years with LTM turned on. I selected the Discuss part of the Virtual Assistant and copied and pasted the tabular output for the population and vegetation stock variables into the Behavioral Description section.

Prompt: Please explain the behavior of the population and vegetation variable over time and how the feedback control the behavior of this model output

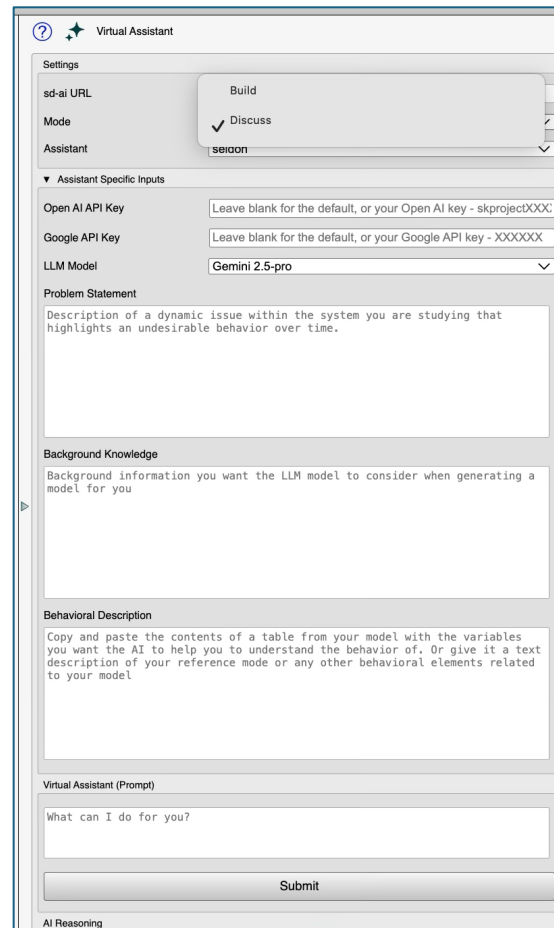
The image shows a web-based interface for a 'Virtual Assistant'. At the top, there's a header with a question mark icon and the text 'Virtual Assistant'. Below this is a 'Settings' section. In the 'Settings' section, there are three main areas: 'sd-ai URL', 'Mode', and 'Assistant'. The 'Mode' dropdown menu is open, showing 'Build' and 'Discuss' (which is selected with a checkmark). Below 'Mode' is the 'Assistant' dropdown, which is set to 'seldon'. Underneath the 'Settings' section is a '▼ Assistant Specific Inputs' section. This section contains three input fields: 'Open AI API Key' (with a placeholder text 'Leave blank for the default, or your Open AI key - skprojectXXXX'), 'Google API Key' (with a placeholder text 'Leave blank for the default, or your Google API key - XXXXXX'), and 'LLM Model' (a dropdown menu set to 'Gemini 2.5-pro'). Below these inputs is a 'Problem Statement' section with a text area containing the text: 'Description of a dynamic issue within the system you are studying that highlights an undesirable behavior over time.' This is followed by a 'Background Knowledge' section with a text area containing the text: 'Background information you want the LLM model to consider when generating a model for you'. Below that is a 'Behavioral Description' section with a text area containing the text: 'Copy and paste the contents of a table from your model with the variables you want the AI to help you to understand the behavior of. Or give it a text description of your reference mode or any other behavioral elements related to your model'. At the bottom of the panel is a 'Virtual Assistant (Prompt)' section with a text area containing the text: 'What can I do for you?'. Below the prompt area is a 'Submit' button. At the very bottom of the panel, there is a section labeled 'AI Reasoning'.

Figure 8: The "Discuss" (Seldon) AI Virtual Assistant property panel in Stella Online.

Seldon response: (I will give an overview here. The complete response is in appendix 4.)

The model describes the dynamics of a deer population and its primary food source, vegetation. The simulation results show a classic "overshoot and collapse" behavior, driven by the interplay of reinforcing and balancing feedback loops.

Behavior of Population and Vegetation Over Time

0. **Initial Growth Phase (Time 0 to approximately Time 20):**
 - **Population:** Starts at 29 deer and exhibits rapid, exponential growth, reaching over 7,000 deer by Time 21.
 - **Vegetation:** Starts at 25,000 vegetation units...
1. **Overshoot and Collapse Phase (Time 21 to Time 22):**
 - **Population:** Peaks around 7,160 ...
 - **Vegetation:** Drops sharply from 1,210 units at Time 21 ...
2. **Unphysical Collapse State (Time 23 to Time 25):**
 - **Population:** Becomes negative (-232 deer) at Time 23 ...
 - **Vegetation:** Continues to be negative (-8,270 units) at Time 23 ...
 - This final phase represents a complete system collapse...

How Feedback Loops Control Model Behavior

Let's analyze the influence of each feedback loop during the different phases:

Feedback Loops in the Model:

- **R1 (Reinforcing Loop): Population -> Births -> Population (+)**
 - This loop represents ...
- **B1 (Balancing Loop): Population -> Deaths -> Population (-)**
 - This loop represents ...
- **B2 (Balancing Loop): Population -> Consumption -> Vegetation (-) -> Vegetation per deer (-) -> Effect of vegetation per deer on death fraction (+) -> Deaths (+) -> Population (-)**
 - This is the "limits to growth" loop...
- **U1 (Unidentified/Partial Loop):** This loop appears to be a subset of B2...

Dominance and Control of Behavior:

1. **Initial Growth Phase (Time 0 to ~Time 20):**
 - **Dominant Loop(s):** R1 (80% of behavior explained)
 - **Explanation:** ...
2. **Overshoot and Collapse Phase (Time 21 to Time 22):**
 - **Dominant Loop(s):** B2 (64.82% at Time 21), B1 (66.45% at Time 22)
 - **Explanation:** ...
3. **Unphysical Collapse State (Time 23 to Time 25):**

- **Dominant Loop(s): R1 (80% at Time 23), B1 (20% at Time 23), then 0% for all loops.**
- **Explanation: ...**

In summary, the model's behavior is initially dominated by the **reinforcing loop R1**, leading to exponential population growth. This growth, however, is unsustainable and eventually triggers the powerful **balancing loop B2**. B2, by linking population to its finite resource (vegetation) and increasing mortality when resources are scarce, becomes dominant and causes the population to overshoot its carrying capacity and then collapse. The model illustrates a classic ecological principle of limits to growth and the consequences of exceeding those limits.

There is much more that Seldon can do, such as help the student develop ideas for a model, recommend stock variables for a given system, suggest feedback for a given system, and more, typical of many LLMs. But the response explained above (using LTM and model numerical output of stocks) to explain model behavior is superior to what a regular LLM could do to explain SD model output.

Using AI in this capacity has the possibility of assisting teachers relatively new to SD with models they may wish to use with their students. It also has the capacity to assist students who, after working with guided SD model-building lessons, want to transition to reading articles and developing their own initial stock-flow models and explaining them to the class – indicating what AI gave them as correct and incorrect in their initial efforts.

As a final application of AI to support infusion of SD as part of instruction, in this paper, we will look at providing feedback to students regarding their explanation of model behavior. Providing individual, targeted feedback to students on their open-ended responses to questions about models is extremely important but very time consuming. AI can provide useful support for this process.

III. Using AI to Provide Individual, Targeted Feedback on Assignments for Students

An extremely valuable application of LLMs for instruction is the chance to provide more detailed individual feedback to students on assignments.

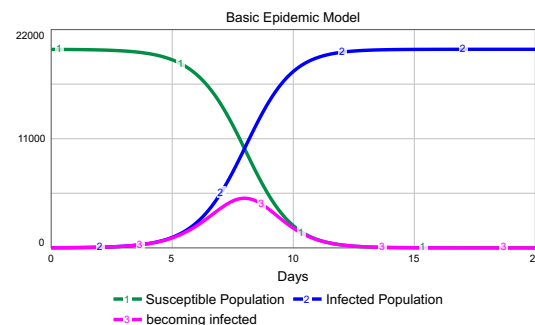
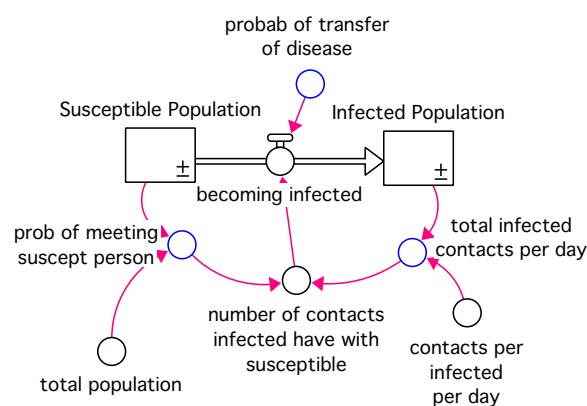
In a high school SD modeling course students played the handshake game to gain a more conceptual understanding of the dynamics of the spread of an epidemic. Then, via a classroom conversation, students developed a stock-flow model diagram on the board. Each student was then to build the stock-flow model in Stella, anticipate model behavior, then run the model and record actual model behavior. (This was before AI was available to them.) The next step was for each student to explain how the model structure produced the behavior of the susceptible population, infected population, and becoming infected population shown when the model was executed. They were to use the feedback in the model to support their explanations, including the idea of transfer of loop dominance.

Now, using LLMs, the student explanations for how the model structure produced the behavior shown in the graphs could be graded relatively quickly producing individual feedback for each student. Students would have to submit their answers to this question to the teacher in digital form.

First it would be necessary to upload the correct/expected answer file into the AI software. The following text would be placed in a file and uploaded into the AI software.

Suggested answer from the teacher:

This is the stock-flow model diagram for the spread of an epidemic in a closed community (left) and the graphical behavior of the Susceptible Population stock, the Infected Population stock, and the becoming infected flow (right).



There are two feedback loops in this model:

Reinforcing: More Infected Population increases total infected contacts per day, which increases the number of contacts infected have with susceptible, which causes more Susceptible People to become infected, increasing the Infected Population.

Balancing: More Susceptible people means there is a greater probability of meeting a susceptible person on a given contact, which means the number of contacts infected have with susceptible will increase, which increases the number of susceptible people becoming infected each day. But if the number of susceptible people becoming infected increases that decreases the total number in the Susceptible Population.

Susceptible Population Graph: Since there are few infected people for the first 4 days the susceptible population is mostly horizontal at its initial maximum level during that time. The reinforcing feedback loop is dominant and the balancing loop is weak. As more people become infected the rate of infection increases, causing the susceptible population to start to decrease more quickly, until the probability of meeting a susceptible person starts to drop significantly (after the inflection point – ie, the maximum of the becoming infected graph – about day 8.5). Now the balancing loop is dominant and slows the decline of the Susceptible Population, causing it to converge to zero.

Infected Population Graph: There are few infected people for the first 4 days so there are few people shown on this graph at this time. But the reinforcing feedback is growing stronger with each newly infected person so the rate of infection (becoming infected) grows quickly, until the number of susceptible people starts to decrease enough that it is difficult for an infected person to meet (and infect) a susceptible person. During this time, after about day 8.5, the balancing loop is becoming dominant, constricting the number of contacts infected have with susceptible, so fewer and fewer susceptible people are becoming infected (infection rate is decreasing). Adding fewer susceptible people over time causes the infected population numbers to reach a steady state.

becoming infected graph: The graph of becoming infected displays when each feedback is dominant in the simulation. At the beginning, as the epidemic is increasing, slowly at first, then more quickly, the reinforcing feedback is dominant. More people are becoming infected at an increasing rate. Then, about day 8.5, the susceptible population has dropped enough to make the probability of meeting a susceptible person small enough that it constricts the growth of the infection. As the probability value continues to decrease the rate of becoming infected decreases until there are no susceptible people left to meet, let alone infect.

After the suggested answer is uploaded to the AI software, then it is possible to use AI to help with grading student responses.

Prompt: I am a high school System Dynamics modeling teacher. I want you to analyze and comment on each student response to the statement: “Explain how the model structure for the spread of an epidemic produced the graphical behavior of the susceptible population, infected population, and becoming infected population shown in its execution.” Students should use the feedback in the model to support their explanations, including the idea of transfer of loop dominance. Identify the strengths (make word “strengths” in bold type) of the student answers, the weaknesses (make word “weaknesses” in bold type) of the student answers, and how the student could improve (make word “could improve” in bold type) their explanation. Be encouraging. In your sentences do not use “the student” or “they” to refer to the student. Instead use “you” to make the response more personal.

Possible student answer:

The susceptible graph stays high at the beginning because there are not many infected people then drops quickly as more people are sick, then levels off because there are no people left to infect. The infected graph starts low because there are very few infected people at the beginning of the epidemic but grows as more people become sick, then levels off because there are no people left to infect. The becoming infected graph starts low, few people are getting sick because there are not many infected people, but grows as more people become sick. Then when all the graphs cross the rate of infection starts to decrease because there are not many people left to infect.

(free) ChatGPT4o Response:

Strengths:

- *Accurate Graph Descriptions:* You clearly describe the general behavior of the graphs, including when the susceptible population decreases and the infected population increases. This shows a solid understanding of how the epidemic progresses over time.

- *Logical Progression of Ideas:* The explanation follows a logical sequence, describing the initial conditions and how the disease spreads and eventually levels off.

Weaknesses:

- *Lack of Feedback Loop Explanation:* The response misses the critical role of the reinforcing loop (driving rapid infection growth) and the balancing loop (slowing the spread as susceptible individuals dwindle).
- *Insufficient Focus on Loop Dominance:* The concept of how the model transitions from reinforcing to balancing feedback dominance isn't included, which is essential for explaining the shapes of the graphs.
- *Vague Terminology:* Phrases like "all the graphs cross" are imprecise and could be replaced with terms like "inflection point" or "steady state."

Could Improve:

- *Incorporate Feedback:* Make the connection between the reinforcing loop at the start of the epidemic (as infections grow rapidly) and the balancing loop that slows the spread as fewer susceptible individuals remain.
- *Use Precise Terms:* Replace "all the graphs cross" with "the inflection point of the becoming infected graph," marking where the rate of infection starts to decline. Clearly identify steady-state behavior for the infected population.
- *Explain Loop Dominance:* Highlight how the dominance of the reinforcing loop transitions to the balancing loop, shaping the behavior of all three graphs.

Encouragement: [Before giving this response to the student I would remove this title but still include the paragraph below]

You've done a great job describing what the graphs show—this is a strong foundation! Now, take it further by explaining *why* the graphs behave this way using feedback loops and loop dominance. These concepts will deepen your understanding and help you connect the model's structure to its outcomes. Keep going—you've got this!

Some high school teachers indicate that this general type of feedback (the ChatGPT4o response above) is not as motivating for students as they hoped. Another type of feedback, called "targeted feedback" has been more motivating for students. A commercial LLM, Brisk, has this feature built into its LLM. I was able to "mimic" this targeted feedback process somewhat in the free ChatGPT4o using this prompt:

Prompt: Using the correct answer for explaining the graphical output for the spread of an epidemic and using the student incomplete response that was provided, produce targeted feedback to the student by highlighting student sentences that were not complete explanations by providing comments in italics with square braces around it recommending to the student what concepts were missing. Make the comment right after the sentence to which it applies. Give a title within your comment in bold referring to the general idea that is missing. Be somewhat general in your feedback selecting the 4 or 5 most important feedback comments to make to help them understand what is missing in their response, pointing the student to what is missing but not telling the student exactly what to say to fix it.

(Originally I had not restricted the LLM to providing 4 or 5 feedback comments. The LLM provided feedback on almost every sentence in the student response, so I placed this comment (4 or 5) in the prompt so students would not become too discouraged.)

(free) ChatGPT4o Response: (Student sentence is in regular type, AI response is in italics. You may want to review the “possible student answer” above.)

The susceptible graph stays high at the beginning because there are not many infected people.

[Missing Loop Structure: This part is missing a reference to the feedback structure that explains why the graph eventually begins to fall.]

Then drops quickly as more people are sick,

[Transfer of Loop Dominance: What causes this drop to speed up and then slow down? Consider how one loop becomes more dominant than the other over time.]

The infected graph starts low because there are very few infected people at the beginning of the epidemic.

[Reinforcing Feedback Loop: How does the presence of more infected people affect the spread of the disease? Think about whether this is part of a reinforcing pattern.]

Then when all the graphs cross the rate of infection starts to decrease because there are not many people left to infect.

[Unclear System Behavior: This statement is vague. Try to be more precise about what is happening in the model's structure that causes the infection rate to fall.]

There are, no doubt, many more useful applications of LLMs to infusing SD into classrooms.

Discussion:

The ability of even free versions of LLMs to support pre-college instruction is exciting and it has the potential to make some of the more difficult SD infusion tasks easier for teachers to incorporate into their teaching (making it easier to create differentiated SD lessons for students, providing individualized feedback for student explanations of Behavior Over Time Graphs and feedback analysis, and especially for helping teachers move to the place where they can have students start to design their own SD models).

Of course students will know LLM features, so will probably use LLMs to construct their responses explaining how feedback works for a given model, or using LLMs when writing their papers, etc. If teachers do not redesign their curriculum with LLMs in mind teacher assessment will lose much of its validity as an accurate assessment of student learning. Bower, et al (2024) analyzed the results of a survey about how teachers felt assessments would have to change given student access to LLMs. The survey had 318 valid responses from teachers (elementary through university, globally) and indicated there was a general consensus that assessments would have to shift towards more in-person and teacher supervised tasks that required higher-order thinking.

A possible assessment adjustment for teachers using LLMs is to have students provide digital responses to open ended questions where the teacher just uses an LLM to provide feedback to the students and does not provide a grade. Or a rubric could be uploaded to the LLM and a temporary grade, with the feedback, could be provided to the students. Then teachers could provide subsequent, similar, open-ended questions in an in-class quiz/test where students have to provide hand-written responses. At least these hand-written responses should be an improvement over the initial responses students submitted (where LLM assisted feedback was given), for those students who actually read the feedback comments and internalize it when studying for the quiz. And the points for the final grade on the in-class assessment could be modified based on whether the student actually took the comments on the intermediate assessment into account when responding to their in-class assessment questions.

Another consideration for teacher assessments, when students have access to LLMs, will be making sure instructional strategies still require students to think hard about responses or their critical thinking skills will suffer. Bower, et al. (2024) indicated that, with regard to how their curriculum should change due to the advent of LLMs, the teachers surveyed felt teachers needed to focus on critical thinking, ethical values (of AI use), creativity, collaboration, and having students attend to “learning processes rather than learning products.”

The types of instructional strategies will need to change. For example, the teacher could ask students to get into groups and use LLMs to develop different stakeholder perspectives on a systemic problem that the group is studying (homelessness, congestion pricing for traffic flow, etc.) and ask each group to comment/expand on information the LLM provided. Or a teacher could suggest student teams ask LLMs to list somewhat simple dynamic feedback system topic ideas and have students sketch possible stock-flow diagrams, etc.

Conclusion:

AI technology, especially LLMs, will have a significant impact on education at all levels. As with most innovations, the delay in making the best and most timely use of the technology will be impacted by how quickly teachers become effective in their use of LLMs in the redesign of their instructional approaches and in how they assess student learning. Of course, state and national assessments of core content in math, science, social science, etc. will also need to adapt, and these organizations change more slowly than classroom instruction.

Yet this fast-moving train may provide opportunities for infusing SD into the middle and high school curriculum because some difficult insertion points could be made easier with the help of LLMs. Using LLMs to:

- identify places in their curriculum that might be most accessible to an SD approach;
- design lesson plans that incorporate BOTGs, stock-flow diagrams, and/or causal loop diagrams. LLMs can help teachers see a possible progression of activities to use and help identify questions to ask students (even providing suggested answers to those questions for the teacher);
- create handouts for students to use as part of the lessons;

- create different types of assessments (on which teachers may decide to provide LLM targeted feedback);
- provide lists of dynamic feedback topics to use when designing lessons;
- create initial stock-flow models for dynamic feedback topics (for the teacher or for the students) that could be expanded upon as a lesson or an assignment;
- determine how to gamify some topics using SD.

Adapting to the availability of LLMs inside and outside the classroom should change what is taught and how it is taught. If this transition is anything like teaching high school computer science courses in the past two decades, where interested students spent hours outside of class creating computer games on their own, having teachers develop a partnership with the students in sharing insights about how to use LLM technology will help teachers progress faster than any professional development schedule administrators could design.

References:

- Bower, M., Torrington, J., Lai, J.W.M., Petocz, P., Alfano, M. (2024). How should we change teaching and assessment in response to increasingly powerful generative Artificial Intelligence? Outcomes of the ChatGPT teacher survey. *Educ Inf Technol* **29**, 15403–15439. Retrieved from: <https://doi.org/10.1007/s10639-023-12405-0>
- Jones, N. (Jan. 14, 2025). “How should we test AI for human-level intelligence? OpenAI’s o3 electrifies quest.” *Nature*. Retrieved at: <https://www.nature.com/articles/d41586-025-00110-6>
- Klein, E. (Jan. 12, 2025). “Now Is the Time of Monsters.” *New York Times*. Retrieved at: https://www.nytimes.com/2025/01/12/opinion/ai-climate-change-low-birth-rates.html?campaign_id=2&emc=edit_th_20250113&instance_id=144577&nl=today%27s-headlines®i_id=57111603&segment_id=188090&user_id=7f6a2e8b3bbeab3aacdf0af8ba12e1f
- Li, F. F. (2023). *The Worlds I See: Curiosity, exploration, and discovery at the dawn of AI*. Flatiron books, New York.
- Mitchell, M. (2019). *Artificial intelligence: A guide for thinking humans*. Farrar, Straus and Giroux Publishers. New York.
- Schank, R. C. 2004. *Making minds less well educated than our own*. Mahwah, NJ: Lawrence Erlbaum Associates, Inc.
- Shah, P. (2023). *AI and the Future of Education: Teaching in the Age of Artificial Intelligence*. John Wiley & Sons.
- Stigler, J. W., & Hiebert, J. 2009. Closing the teaching gap. *Phi Delta Kappan*, 91(3), 32.

Appendix 1: (Full ChatGPT4o response for sequence of lessons/topics for introducing SD into a high school 2nd year algebra class regarding dynamics involved in the spread of an epidemic)

1. Introduction to Epidemic Dynamics

- **Objective:** Familiarize students with basic concepts such as “healthy,” “infected,” and “recovered.”
- **Activity:** Show a simple simulation (e.g., a video or animation) of disease spread in a closed population.
- **Discussion Questions:**
 - What factors might influence the spread of a disease?
 - How does the size of the population affect the spread?
 - What happens when individuals recover? What about immunity?

2. Behavior Over Time Graphs

- **Objective:** Help students think about how the number of healthy, infected, and recovered individuals changes over time.
- **Activity:** Students sketch behavior over time graphs for each group based on observations or given scenarios.
- **Discussion Questions:**
 - How do the numbers of infected individuals change at the start of an outbreak? Why?
 - What do you notice about the healthy population as infections increase?
 - How does the recovered group grow over time?

3. Introducing Stock and Flow Concepts

- **Objective:** Distinguish between stocks (e.g., healthy, infected, recovered) and flows (e.g., infection, recovery).
- **Activity:** Define stock and flow concepts and brainstorm variables for an epidemic model.
- **Discussion Questions:**
 - What is the “stock” of healthy individuals, and how can it decrease or increase?
 - What factors (flows) might increase or decrease the number of infected individuals?
 - Could there be inflows or outflows that influence recovery?

4. Building a Simple Stock-Flow Diagram

- **Objective:** Create a basic stock-flow diagram for the spread of disease.
- **Activity:** Collaboratively identify stocks, flows, and factors to build the diagram together on paper or digitally.
- **Discussion Questions:**
 - How can you represent the healthy population in a diagram?
 - What is the role of infection rate in the model? How might it connect healthy and infected populations?
 - Why might recovery depend on the number of infected individuals?

5. Exploring Feedback Loops

- **Objective:** Understand the feedback loops in the system.
- **Activity:** Identify reinforcing (spread of infection) and balancing (recovery) loops in the model.
- **Discussion Questions:**
 - What happens if there are no balancing loops? Can a disease infect everyone indefinitely?
 - How does immunity or recovery slow down the spread of infection?

6. Experimenting with Model Variables

- **Objective:** Test how different parameters affect the dynamics (e.g., infection rate, recovery time).
- **Activity:** Use a simple simulation tool (e.g., InsightMaker, Stella) to explore different scenarios.
- **Discussion Questions:**
 - How does increasing the infection rate change the spread?
 - What happens if the recovery rate is very slow or very fast?
 - Can you find conditions that lead to the disease dying out?

7. Policy and Real-World Reflection

- **Objective:** Reflect on interventions and their impacts on disease dynamics.
- **Activity:** Discuss and simulate the effects of policies like vaccination or quarantine.
- **Discussion Questions:**
 - How would isolating infected individuals affect the system?
 - What role could vaccination play in altering the stocks and flows?
 - In the real world, what trade-offs exist between these interventions?

Final Project

- **Objective:** Have students independently construct and present a stock-flow model for a specific disease scenario.
- **Prompt:** Build a stock-flow diagram and behavior-over-time graphs for the spread of an epidemic in a closed population, and analyze the results of varying parameters.

Appendix 2: (Full ChatGPT4o response for sequence of lessons/topics for introducing SD into a middle school algebra class regarding dynamics involved in the spread of an epidemic)

1. Introducing the Problem: How Do Diseases Spread?

- **Objective:** Engage curiosity and connect to real-world situations.
- **Activity:** Share a relatable story or video showing disease spread (e.g., a common cold in a classroom or COVID-19).
- **Discussion Questions:**
 - How do people get sick from one another?
 - What might stop the disease from spreading to everyone?
 - What does “recover” mean in this context?

2. Visualizing the Disease Spread

- **Objective:** Observe and record how groups of individuals (healthy, sick, recovered) change.
- **Activity:** Conduct a simple hands-on simulation:
 - Use a dice-roll or colored tokens to simulate disease transmission over several “days.”
 - Track the number of healthy, sick, and recovered individuals each day.
- **Discussion Questions:**
 - What happens to the number of sick people as the disease spreads?
 - Why do healthy people decrease over time?
 - How does recovery affect the spread?

3. Building Behavior Over Time Graphs

- **Objective:** Represent the changes in the healthy, sick, and recovered populations visually.
- **Activity:** Students plot data from the simulation to create behavior-over-time graphs for each group.
- **Discussion Questions:**
 - How do the numbers of sick and recovered individuals change over time?
 - What do you notice about the healthy population as the disease spreads?
 - Why does the graph eventually level off?

4. Introducing Stocks and Flows

- **Objective:** Distinguish between accumulations (stocks) and rates of change (flows).
- **Activity:** Explain:
 - Stocks: Healthy, Sick, Recovered individuals.
 - Flows: Infection rate and recovery rate.
 - Use relatable examples like a bathtub (water filling = flow, water amount = stock).
- **Discussion Questions:**
 - What can fill or empty the “stock” of healthy people?
 - How does recovery “add” people to the recovered stock?
 - What causes infection rates to increase or decrease?

5. Drawing a Basic Stock-Flow Diagram

- **Objective:** Visualize how stocks and flows are connected.
- **Activity:**
 - Students help create a simple stock-flow diagram as a class.
 - Include stocks: Healthy, Sick, Recovered.
 - Flows: Infection (Healthy → Sick), Recovery (Sick → Recovered).
- **Discussion Questions:**
 - How does the number of healthy people affect the number of sick people?
 - What happens to the sick group as recovery occurs?
 - Can sick people influence how many get infected?

6. Exploring Feedback Loops with Hands-On Demonstrations

- **Objective:** Understand reinforcing and balancing loops.
- **Activity:**
 - Reinforcing Loop: Show how one sick person leads to more infections.
 - Balancing Loop: Introduce recovery slowing down the infection spread.
- **Discussion Questions:**
 - What happens if there's no recovery? Can everyone get sick?
 - How does recovery help balance the spread?

7. Exploring Parameters with Simplified Data

- **Objective:** Test how different factors (infection rate, recovery time) affect outcomes.
- **Activity:** Use a simple spreadsheet or paper scenarios:
 - Vary infection rates and track outcomes (e.g., fast recovery vs. slow recovery).
- **Discussion Questions:**
 - What happens when people recover faster? How does it affect others getting sick?
 - What if the disease spreads very quickly?
 - What might keep the infection rate low in real life?

8. Culminating Project: Create and Explain

- **Objective:** Synthesize understanding into a student-built model.
- **Activity:**
 - Students design their own stock-flow diagram of an epidemic.
 - Present the diagram, explaining how stocks and flows interact.
- **Discussion Questions:**
 - Why did you place the stocks and flows where you did?
 - How does your model explain what might happen during an outbreak?

Appendix 3: (The AI reasoning output provide by the LLM in Stella Online using the “Build” LLM feature to produce a stock-flow model from the text description of a non-renewable resource with a deer population)

You instructed me to...

I want you to create a system dynamics simulation of a population (stock) that starts at 29 deer that has an inflow called births that is defined as the product of the population and the birth fraction (0.4) and has an outflow called deaths that is defined as the product of the population and the death fraction (0.1). This system dynamics simulation should also contain a vegetation (stock) that starts at 25000 veg units and only has an outflow called consumption. Each deer

consumes one unit of vegetation per year. The deer consume the vegetation. A converter, which depends on the deer population stock and vegetation stock, is called effect of vegetation per deer on death fraction and should be constructed as a graphical converter. The graph in this graphical converter has a horizontal axis that goes from 0 to 1 and a vertical axis that goes from 1 to 15. Assume the graph in this graphical converter is concave down and connects the points (0,15) and (1,1). The effect of vegetation per deer on death fraction value should influence (multiply by) the death fraction. The AI should build the correct stock flow diagram with appropriate values within each icon. The simulation should run for 25 simulated years.

Here is what I have done...

Deer Population and Vegetation Dynamics

The model now incorporates a 'Limits to Growth' archetype. The deer population grows, but this growth is eventually constrained by the availability of vegetation. As the deer population increases, it consumes more vegetation, leading to a decline in the vegetation stock. This reduction in vegetation per deer then negatively impacts the deer population by increasing the death rate, thus creating a balancing feedback loop that limits population growth. This loop is closed by the relationships from population to consumption, consumption to vegetation, vegetation to vegetation per deer, vegetation per deer to the effect on death fraction, and finally back to deaths and population.

1. population -->(+) births

Births are directly proportional to the current population size.

2. birth fraction -->(+) births

The birth fraction is a direct multiplier for calculating the birth rate.

3. population -->(+) deaths

Deaths are directly proportional to the current population size.

4. death fraction -->(+) deaths

The base death fraction is a direct multiplier for calculating the death rate.

5. effect of vegetation per deer on death fraction -->(+) deaths

This variable directly multiplies the base death fraction to determine the overall death rate.

6. population -->(+) consumption

Each deer consumes a certain amount of vegetation, so more deer lead to more total consumption.

7. vegetation -->(+) vegetation per deer

Vegetation per deer is calculated as total vegetation divided by the deer population.

8. population -->(-) vegetation per deer

Vegetation per deer is calculated as total vegetation divided by the deer population, so an increase in the denominator leads to a decrease in the result.

9. vegetation per deer -->(-) effect of vegetation per deer on death fraction

The availability of vegetation per deer directly influences the multiplier for the death fraction, as defined by the graphical converter.

Appendix 4: (Stella Online Seldon (LLM) output explaining the behavior of the non-renewable resource and deer model)

My [Seldon] reply...

The model describes the dynamics of a deer population and its primary food source, vegetation. The simulation results show a classic "overshoot and collapse" behavior, driven by the interplay of reinforcing and balancing feedback loops.

Behavior of Population and Vegetation Over Time

1. Initial Growth Phase (Time 0 to approximately Time 20):

- **Population:** Starts at 29 deer and exhibits rapid, exponential growth, reaching over 7,000 deer by Time 21.
- **Vegetation:** Starts at 25,000 vegetation units. During this phase, vegetation decreases, but relatively slowly at first (e.g., only dropping to 23.8k by Time 10). As the population grows, the rate of vegetation consumption increases, leading to a more noticeable decline in later years of this phase (e.g., dropping to 1.21k by Time 21). However, vegetation remains positive and sufficient to support the population's growth.

2. Overshoot and Collapse Phase (Time 21 to Time 22):

- **Population:** Peaks around 7,160 deer at Time 21. Immediately after, it experiences a dramatic crash, falling to 2,320 deer by Time 22.
- **Vegetation:** Drops sharply from 1,210 units at Time 21 to a negative value of -5,950 units by Time 22. This indicates that the vegetation has been severely depleted, reaching a point where it can no longer sustain the population.

3. Unphysical Collapse State (Time 23 to Time 25):

- **Population:** Becomes negative (-232 deer) at Time 23 and remains constant at this unphysical value for the rest of the simulation.
- **Vegetation:** Continues to be negative (-8,270 units) at Time 23 and remains constant.
- This final phase represents a complete system collapse, where the model's variables have gone beyond realistic bounds, indicating the unsustainability of the initial growth trajectory.

How Feedback Loops Control Model Behavior

Let's analyze the influence of each feedback loop during the different phases:

Feedback Loops in the Model:

- **R1 (Reinforcing Loop): Population -> Births -> Population (+)**
 - This loop represents the inherent growth potential of the population. More deer lead to more births, which leads to even more deer.
- **B1 (Balancing Loop): Population -> Deaths -> Population (-)**
 - This loop represents the basic mortality rate. More deer lead to more deaths, which tends to reduce the population.
- **B2 (Balancing Loop): Population -> Consumption -> Vegetation (-) -> Vegetation per deer (-) -> Effect of vegetation per deer on death fraction (+) -> Deaths (+) -> Population (-)**
 - This is the "limits to growth" loop. As the population grows, it consumes more vegetation, depleting the resource. Less vegetation per deer leads to a higher death rate, which then reduces the population.
- **U1 (Unidentified/Partial Loop):** This loop appears to be a subset of B2, focusing on the impact of vegetation per deer on deaths. Given B2's comprehensive nature, we will primarily focus on B2 for the resource-dependent balancing effect.

Dominance and Control of Behavior:

4. **Initial Growth Phase (Time 0 to ~Time 20):**
 - **Dominant Loop(s): R1 (80% of behavior explained)**
 - **Explanation:** During this phase, the **Reinforcing Loop R1** is overwhelmingly dominant. The birth rate (birth fraction = 0.4) is significantly higher than the base death rate (death fraction = 0.1). Crucially, the effect of vegetation per deer on death fraction is close to 1 because vegetation is abundant, meaning vegetation per deer is high. This keeps the death rate relatively low. The strong positive feedback from R1 drives the exponential growth of the deer population. The simple balancing loop B1 (deaths proportional to population) is present but only accounts for 20% of the behavior, not strong enough to counteract R1's growth. The resource-dependent balancing loop B2 has virtually no influence (0%) because vegetation is not yet a limiting factor.
5. **Overshoot and Collapse Phase (Time 21 to Time 22):**
 - **Dominant Loop(s): B2 (64.82% at Time 21), B1 (66.45% at Time 22)**
 - **Explanation:** As the population grows exponentially due to R1, it consumes vegetation at an ever-increasing rate. By Time 21, the vegetation stock has been severely depleted (down to 1.21k). This causes vegetation per deer to drop significantly (e.g., at Time 21, 1210 units / 7160 deer \approx 0.17 units/deer).
 - When vegetation per deer falls to such low levels, the **effect of vegetation per deer on death fraction** (as defined by its graphical function) dramatically increases the death multiplier (e.g., at 0.2 units/deer, the multiplier is 10; at 0 units/deer, it's 15).
 - This activation of the effect of vegetation per deer on death fraction significantly amplifies the deaths rate.
 - Consequently, the **Balancing Loop B2** becomes highly active and dominant at Time 21, explaining nearly 65% of the model's behavior. It effectively imposes a "limit to growth" by increasing mortality due to resource scarcity.

- At Time 22, the general **Balancing Loop B1** (which is now heavily influenced by the high death rate driven by B2) becomes the most dominant factor, explaining over 66% of the behavior. The combined effect of B1 and B2 (which directly impacts B1's strength) causes the population to crash rapidly, as deaths far exceed births. The reinforcing loop R1's influence drops sharply during this period.

6. **Unphysical Collapse State (Time 23 to Time 25):**

- **Dominant Loop(s):** R1 (80% at Time 23), B1 (20% at Time 23), then 0% for all loops.
- **Explanation:** The model has entered an unphysical state where both population and vegetation are negative. The loop dominance percentages for Time 23 might be an artifact of the calculation in this extreme state, as the system has already crashed. For Time 24 and 25, all loops show 0% influence, suggesting that the model has reached a numerical equilibrium where no further changes are occurring, or the dynamics are no longer meaningful.

In summary, the model's behavior is initially dominated by the **reinforcing loop R1**, leading to exponential population growth. This growth, however, is unsustainable and eventually triggers the powerful **balancing loop B2**. B2, by linking population to its finite resource (vegetation) and increasing mortality when resources are scarce, becomes dominant and causes the population to overshoot its carrying capacity and then collapse. The model illustrates a classic ecological principle of limits to growth and the consequences of exceeding those limits.