Cloudy skies: assessing public understanding of global warming

John D. Sterman* and Linda Booth Sweeney

John D. Sterman is the Standish Professor of Management and Director of the System Dynamics Group at the MIT Sloan School of Management.

Linda Booth Sweeney (Linda_Booth_Sweeney @harvard.edu) is a doctoral candidate at the Harvard Graduate School of Education. Her current research involves developing and testing a Systems Thinking Inventory to identify folk theories underlying conceptions of complex system behaviors, and investigating whether understanding of ecosystems translates into improved reasoning about social systems.

Abstract

Surveys show that most Americans believe global warming is real. But many advocate delaying action until there is more evidence that warming is harmful. The stock and flow structure of the climate, however, means ‘‘wait and see’’ policies guarantee further warming. Atmospheric CO_2 concentration is now higher than any time in the last 420,000 years, and growing faster than any time in the past 20,000 years. The high concentration of CO_2 and other greenhouse gases (GHGs) generates significant radiative forcing that contributes to warming. To reduce radiative forcing and the human contribution to warming, GHG concentrations must fall. To reduce GHG concentrations, emissions must fall below the rate at which GHGs are removed from the atmosphere. Anthropogenic CO_2 emissions are now roughly double the removal rate, and the removal rate is projected to fall as natural carbon sinks saturate. Emissions must therefore fall by more than half even to stabilize CO_2 at present record levels. Such reductions greatly exceed the Kyoto targets, while the Bush administration’s Clear Skies Initiative calls for continued emissions growth. Does the public understand these physical facts? We report experiments assessing people’s intuitive understanding of climate change. We presented highly educated graduate students with descriptions of greenhouse warming drawn from the IPCC’s nontechnical reports. Subjects were then asked to identify the likely response to various scenarios for CO_2 emissions or concentrations. The tasks require no mathematics, only an understanding of stocks and flows and basic facts about climate change. Overall performance was poor. Subjects often select trajectories that violate conservation of matter. Many believe temperature responds immediately to changes in CO_2 emissions or concentrations. Still more believe that stabilizing emissions near current rates would stabilize the climate, when in fact emissions would continue to exceed removal, increasing GHG concentrations and radiative forcing. Such beliefs support ‘‘wait and see’’ policies, but violate basic laws of physics. We discuss implications for education and public policy.

Syst. Dyn. Rev. **18**, 207–240, (2002)

Climate change presents a classic system problem. The spatial scale is global; the time scale dwarfs normal human concerns. The dynamics of the climate are exquisitely complex and imperfectly understood. ‘‘Climate’’ is an aggregate average of local and noisy weather over long time frames and broad geographical extent and can only be perceived with a long delay; many effects of climate change will be hard to attribute properly. As a global resource shared by all, the climate is vulnerable to the tragedy of the commons since individuals, firms, and nations benefit in the short-run from high greenhouse gas (GHG) emissions, while the costs are borne by all. Worse, the costs and benefits are distributed inequitably, both between rich and poor and between future

---

* John D. Sterman, MIT Sloan School of Management, 30 Wadsworth Street, Room E53-351, Cambridge, MA 02142, U.S.A; E-mail: jsterman@mit.edu

System Dynamics Review Vol. 18, No. 2, (Summer 2002): 207–240
DOI: 10.1002/sdr.242

Received March 2002
Accepted May 2002

generations and ourselves. Small wonder that the challenge of implementing policies to reduce the impact of warming remains unanswered.

While surveys show the vast majority of Americans believe global warming is real, the majority also believe, as a recent survey found, that ‘‘its effects will be gradual, so we can deal with the problem gradually’’ or that ‘‘until we are sure that global warming is really a problem, we should not take any steps that would have economic costs.’’[1] The lack of government support for the Kyoto Protocol, carbon taxes, and other policies designed to reduce GHG emissions can be traced to public indifference. Many people believe that there is plenty of time to act. They reason that there is little evidence of warming so far, and less that it is harmful. Why take costly action today for uncertain benefits tomorrow? Instead, they say, it is better to wait and see—if warming turns out to be greater and more harmful than expected, policies to mitigate it can then be implemented.

Wait and see policies, however, often do not work in systems with long time delays, significant accumulations, multiple feedback processes, nonlinearities, and other elements of dynamic complexity (Sterman 1994). When there are long time delays between actions and their effects, as in the climate system, wait and see can be disastrous. Actions to halt warming must be taken decades before we can know what the consequences of warming will be, and before scientific certainty about the dynamics of the global climate can be gained.

We hypothesize that much of the complacency about climate change arises from poor systems thinking skills. If people do not recognize the existence of important feedbacks, time delays, and stock and flow structures in the climate system, or if they are unable to relate these structures to the dynamics of the climate, they are likely to draw erroneous inferences about the response of the climate to human activity. Without such understanding people are likely to rely on the intuitive ‘‘wait and see’’ strategy that works well in a range of everyday tasks.

In this article we report experiments assessing people’s understanding of the structure and dynamics of the global climate. We present highly educated adults enrolled in university graduate programs with descriptions of the climate and data on past emissions, atmospheric CO_2 concentrations, and global mean temperature, then ask them to identify the likely response of the system to various scenarios for CO_2 emissions or concentrations. People perform poorly on these simple tasks. They often select trajectories for emissions or CO_2 concentrations that violate conservation of matter. They significantly underestimate the inertia in the system. Many believe global temperatures respond immediately to changes in CO_2 emissions or concentrations, so that a peak and decline in emissions causes an immediate peak and decline in temperature. Still more believe that stabilizing emissions near current rates, as called for in the Kyoto Protocol, would stabilize the climate. Such beliefs support a ‘‘wait and see’’ policy, but violate the basic laws of physics. Stabilizing emissions at Kyoto rates ensures continued growth in GHG concentrations and continued warming (IPCC 2001a, b). To reduce radiative forcing, concentrations of greenhouse gases (GHGs) must fall. To reduce GHG concentrations, emissions must fall below the rate at which GHGs are removed from the atmosphere. Anthropogenic CO_2 emissions now exceed removal by about a factor of two, and the removal rate is projected to fall as natural carbon sinks saturate (IPCC 2001b; Sarmiento *et al.* 1995). Emissions must therefore fall by more than half even to stabilize CO_2 at present record levels. Reducing CO_2 concentrations and radiative forcing require even deeper cuts. Such reductions greatly exceed the Kyoto targets, while the Bush administration's Clear Skies Initiative calls for continued emissions growth.

Our work extends the research reported in Booth Sweeney and Sterman (2000), which reported experiments assessing people's intuitive systems thinking skills, including stocks and flows, time delays, and simple feedbacks. The results were not encouraging. For example, subjects were shown a picture of a bathtub and given a graph showing the inflow and outflow over time. They were asked to sketch on another graph the quantity of water in the tub. The tasks were simple and required no knowledge of higher mathematics and little calculation. Yet highly educated subjects with extensive training in mathematics and science performed poorly. Subjects frequently violated fundamental relationships between stocks and flows, including conservation of matter. They had poor understanding of the relationship between the net flow into a stock and the rate of change of the stock. Many subjects appeared to rely on a heuristic in which they matched the shape of the output of the system to the shape of the input, a heuristic that leads to gross error in systems with stock and flow structures and time delays (Moxnes 2000 also found some support for pattern matching in an experimental study of renewable resource management).

It is not surprising, as research suggests, that school children do not understand the physical processes that govern climate change (Meadows and Wiesenmayer 1999; Groves and Pugh 1999). Our results are more disturbing. The errors and misconceptions exhibited by highly educated adults constitute a serious challenge to informed debate over climate change policy. We discuss possible sources for these errors, their consequences, and ways they might be overcome.

Structure and dynamics of the climate

Despite the complexity of the climate, the essentials are simple and can be easily understood with basic knowledge of stocks and flows. The temperature at the earth's surface—the land, lower atmosphere, and surface layer of the ocean (the top 50 to 100 meters, where most sea life exists)—is determined primarily by the balance of the incoming solar radiation (insolation) and the energy the earth radiates back to space. The earth is a warm mass surrounded by the cold of space and like all such masses emits black body radiation whose frequency distribution and intensity depends on its surface temperature. The warmer the earth, the greater the flow of energy radiated back into space. The result is a negative feedback process: Incoming solar energy heats the earth until it is just warm enough for the energy radiated back to space to balance the solar energy input.

The amount of energy radiated back into space depends on the composition of the atmosphere. Greenhouse gases such as carbon dioxide absorb some of the energy radiated by the earth, instead of allowing it to escape into space. Thus an increase in GHGs causes the earth to warm. The earth heats up until the energy escaping through the atmosphere rises enough to again balance the incoming solar energy. Naturally occurring greenhouse gases—including water vapor—reduce the emissivity of the atmosphere enough to warm the surface of the earth (including the oceans) to a life-sustaining average of about 15 °C (59 °F). Without GHGs in the atmosphere, the mean global temperature would be about −17 °C (0 °F) and a blanket of ice would perpetually cover the earth.

Natural biogeochemical processes have caused the concentration of carbon dioxide in the atmosphere to fluctuate over geological time, and surface temperatures have fluctuated with it. Human activity has now reached a scale where it affects these processes significantly. Anthropogenic GHG emissions have been growing exponentially since the beginning of the industrial age. Consequently, atmospheric concentrations of CO_2 and other GHGs including nitrous oxide (N_2O), methane (CH_4), chlorofluorocarbons (CFCs), hydrofluorocarbons (HFCs), perfluorinated carbons (PFCs), and others have been growing exponentially. The preindustrial atmospheric CO_2 concentration was about 280 ppm; today it is about 370 ppm and rising. Concentrations of other important GHGs are also rising; for example, N_2O, and CH_4 concentrations are up by 17 percent and 151 percent, respectively, since 1750. The UN-sponsored Intergovernmental Panel on Climate Change (IPCC) notes that:

> The present CO_2 concentration has not been exceeded in the last 420,000 years and likely not during the past 20 million years. The current rate of increase is unprecedented during at least the past 20,000 years. (IPCC 2001a: 7)

Current greenhouse gas concentrations contribute about 2.4 watts per square meter of net radiative forcing, that is, incoming solar radiation exceeds outgoing radiation by 2.4 w/m^2. Consequently, mean global surface temperature has been rising. The IPCC (2001a, b) reports that mean global temperatures rose in the twentieth century by $0.6 \pm 0.2\,^{\circ}C$. The warming has been accompanied by glacier retreat and a decline in winter snow cover, a 40 percent decline in summer sea-ice thickness in the Arctic, an increase in average precipitation and in extreme weather events, and a rise of 0.1–0.2 meters in sea level, among other effects.

Debate continues about the dynamics of the global climate system, its response to forcing by human activity, and the consequences of warming. The public discussion has been polarized by well-financed campaigns to discount the science. Nevertheless, consensus is emerging. The IPCC concludes in its most recent report that global warming is real and that ‘‘most of the warming observed over the last 50 years is attributable to human activities’’ (IPCC 2001a: 10).

Simulation models of various types are the primary research tools used to explore the likely response of the climate to forcing by human activity. Detailed general circulation models (GCMs) simulate climate at finely spaced intervals covering the entire surface of the earth, but take GHG emissions as exogenous inputs. At the other extreme, so-called integrated climate–economy models close some of the feedbacks among the human economy, carbon emissions, and global climate but treat the carbon cycle and climate as global aggregates with a small number of stocks. Despite the differences among the models, all show the climate to possess enormous inertia. Changes in GHG emissions only slowly affect temperature and climate, and their impact persists for decades.

Fiddaman (1997) analyzes many of the most widely used climate–economy models, identifying a number of problems and inconsistencies in them, including many of the errors in stock-flow representation discussed in this article. For example, in his widely cited DICE model, Nordhaus (1992a, b) violates the law of conservation of mass by assuming that a significant fraction of carbon emissions simply disappear (Nordhaus assumed these emissions flow into a limitless sink outside the model boundary). Fiddaman (2002, in this issue) developed a model that corrects these and other defects in common climate–economy models and linked it to a model of the economy and energy system. The model is carefully calibrated to the available data; we use it as the basis for one of the tasks in our experiments.

Figure 1 shows a simplified representation of the stock and flow structure of the global carbon cycle and heat balance. The global carbon stock is

Fig. 1. Simplified representation of the carbon cycle and global temperature

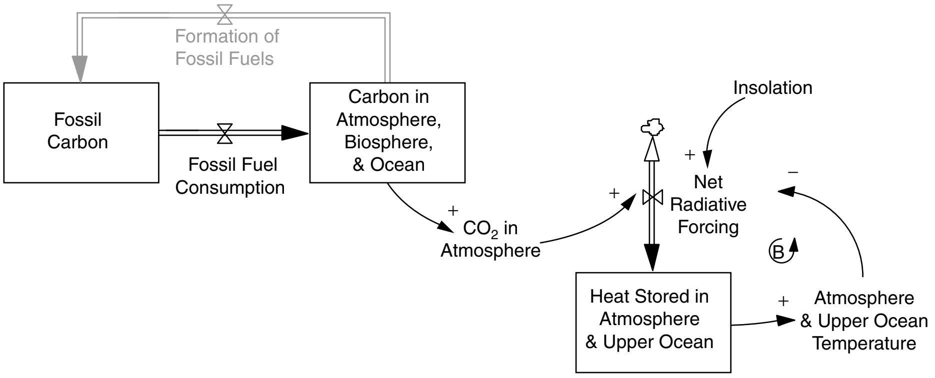

divided into two compartments: (1) carbon sequestered in, e.g., fossil fuels and therefore not interacting with the climate at present; and (2) carbon in the atmosphere, in biomass, and dissolved in the ocean. Over geological time carbon in terrestrial biomass and in the ocean gradually accumulated in stocks of fossil fuels, reducing the quantity of carbon circulating among the atmosphere, terrestrial biomass, and the ocean.[2] Since the beginning of the industrial revolution, however, fossil-fuel consumption has grown exponentially, injecting previously sequestered carbon into the atmosphere. While some of this carbon is taken up by biomass or dissolved in the ocean, the equilibrium concentration of CO_2 in the atmosphere rises. The rise in atmospheric CO_2 increases net radiative forcing, causing the stock of heat at the surface and in the surface layer of the ocean to rise until the surface is warm enough for radiation of energy back to space to balance the incoming solar energy. The stock and flow structure in Figure 1, simplified though it is, shows that an injection of fossil carbon to the atmosphere leads to a rise in average surface temperatures, though with a long lag.

Figure 2 shows a more detailed representation of the carbon cycle and heat balance, disaggregating non-fossil carbon into stocks of atmospheric CO_2, carbon in terrestrial biomass, carbon in soils, and carbon in the ocean (both dissolved and in oceanic biomass). Burning fossil fuels adds CO_2 to the atmosphere. Higher atmospheric CO_2 then increases the rate at which CO_2 is consumed by aquatic life or dissolves into the mixed layer of the ocean. Eventually, carbon taken up by the surface layer is transported to deeper waters, through ocean currents, diffusion, and precipitation of detritus from aquatic life. The transfer of carbon to the depths is slow, and mixing between the surface and abyssal waters is weak, so many carbon-cycle models disaggregate the water column into a number of distinct states and model the transfer of carbon between adjacent layers explicitly. Fiddaman's model utilizes ten layers, enough to capture the slow adjustment of abyssal carbon concentrations to changes in carbon concentrations in the mixed layer.

Increased atmospheric CO_2 concentration also stimulates uptake of carbon by terrestrial plants (the flux of carbon to biomass). Most of the carbon in biomass is released back into the atmosphere through respiration and metabolic activity of animal and bacterial life, and by fire (natural and human-caused). A relatively small fraction of the carbon in biomass is transferred to soils. The carbon in soils can be taken up directly into biomass as plants grow or can be released into the atmosphere through decay.

Note that Figure 2 does not show the conversion of carbon in soils or the abyssal layer of the ocean into new stocks of fossil fuels. Although the dynamics of global warming will play out over the next several centuries, this interval is so short relative to the millions of years required to form oil, gas, and coal that these carbon flows are negligible over the relevant time horizon.

Figure 2 also shows the transfer of heat from the surface layer to the deep ocean. The heat capacity of the ocean is large and mixing is slow, slowing the

Fig. 2. More detailed representation of the carbon cycle and global temperature balance

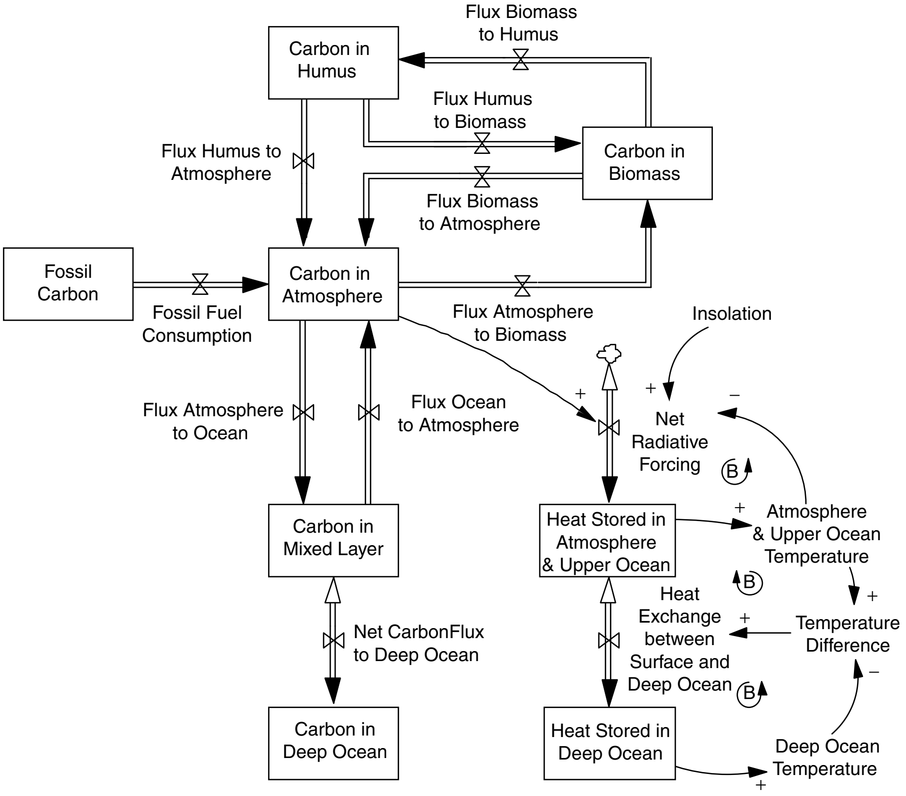

adjustment of global temperature to the equilibrium level determined by the concentration of greenhouse gases.

Method, subjects, and procedure

We created two tasks to explore people's intuitive understanding of global warming and the dynamics of the climate, the Zero Emissions (ZE) task and the Stable CO_2 Concentration (SC) task. Each consists of a short description of the climate system, graphs showing the history of human CO_2 emissions, the concentration of CO_2 in the atmosphere, and global mean temperature. Subjects were presented with scenarios for the evolution of emissions or atmospheric CO_2 and asked to specify the likely response of the climate. Both tasks can be done without use of mathematics beyond simple arithmetic.

We administered the tasks to students from MIT and Harvard enrolled in the introductory system dynamics course at the MIT Sloan School of Management, and to students at the Graduate School of Business at the University of Chicago.

The Zero Emissions task was administered in the fall of 1999. These students received the task early in the semester, before the concept of stocks and flows was introduced. Students were given approximately 10 minutes. They were told that the purpose of the questions was to illustrate important systems thinking concepts they were about to study and to develop a tool to assess systems thinking skills. Students were not paid or graded. The University of Chicago students were MBAs taking a course in negotiation in the spring of 2001. They received, in a single session, one of the bathtub tasks described in Booth Sweeney and Sterman (2000) and the ZE task. They had about 10 minutes to do both tasks; many finished before the end of the allotted period. Again, they were told that the purpose of the questions was to illustrate important systems thinking concepts and were neither paid nor graded. The Stable CO_2 Concentration task was administered in the spring term of 2002 to the introductory system dynamics course at the MIT Sloan School, on the first day of the term.

The MIT subjects were primarily MBA students but also included students in other master's degree programs, Ph.D. students, undergraduates, and students cross-registered from graduate programs at other local universities, primarily Harvard (Table 1). Average age was about 29; the majority had at least several years of business experience. The students were highly international, with several dozen countries represented; English was a first language for about half. About three-quarters were male. Two-thirds had undergraduate degrees in engineering, computer science, mathematics, or the sciences; most of the rest studied business or a social science (primarily economics). Fewer than 5 percent had degrees in the humanities. More than a third already had a graduate degree in another field, including 5 percent with doctorates. Most of these advanced degrees were in technical fields. About three-quarters of the Chicago students were regular MBA students (85 percent second year, 15 percent first year); the remainder were enrolled in the evening program.

Initial coding criteria were developed and tested on a subsample. The coding criteria were revised to resolve ambiguities. A detailed coding guide, covering the global warming tasks and the tasks reported in Booth Sweeney and Sterman (2000), is available from the authors.

The Zero Emissions task

The Zero Emissions task (Figure 3) describes a thought experiment in which subjects are asked to imagine that human CO_2 emissions fall instantly to zero in the year 2000. The task provides a brief description of the global warming issue. Subjects received one of two experimental treatment conditions. In the ‘‘CO_2 Graph’’ condition subjects were provided with graphs showing data from 1950 for CO_2 emissions, CO_2 in the atmosphere, and global mean temperature. They were asked to sketch the likely path of atmospheric CO_2 and mean temperature

Table 1. Subject demographics

All entries are %		**MIT Group 1 ZE Task**	**MIT Group 2 SC Task**
Age	19–24	10	14
	25–30	59	53
	31–35	20	22
	36 and up	11	11
Gender	M	76	78
	F	24	22
Student Status			
	1st/2nd yr MBA	37	39
	Executive MBA[a]	16	11
	LFM[b]	16	5
	Other[c]	32	46
Prior Field of Study			
	Business/management	27	11
	Engineering	45	52
	Social sciences	4	18
	Science	10	7
	Computer science	5	8
	Math	6	2
	Humanities	4	3
Highest Prior Degree			
	BA	17	16
	BS	43	40
	MA/MS	28	35
	Ph.D.	5	5
	High school	2	1
	BE, JD, BBA, MD	1	2
	BA & BS	3	1
Region of Origin			
	North America (+ Australia/NZ)	50	39
	Europe	16	18
	Asia and Middle East	22	33
	Latin America	10	9
	Africa	2	1
English	First language	54	50
	Not first language	46	50

[a] The executive MBA students were enrolled in various one-year degree programs at MIT and are typically mid-career executives with extensive business experience.

[b] LFM = Leaders for Manufacturing, a dual degree program awarding both an MBA and MS in engineering.

[c] "Other" includes Ph.D. and graduate students from other MIT departments (averaging about 10–15 percent), graduate students from other universities (primarily Harvard), averaging about 10–15 percent, and MIT undergraduates (about 5 percent).

Totals may not equal 100% due to rounding.

Consider the problem of global warming. Carbon dioxide (CO_2) is a greenhouse gas that traps heat and contributes to warming. CO_2 emissions from combustion of fossil fuels like oil, gas, and coal have been increasing since the start of the industrial revolution. The curve labelled "Anthropogenic CO_2 Emissions" in Figure 1 shows the worldwide emission rate of CO_2 from fossil fuel combustion since 1950. Figure 2 shows the stock of carbon dioxide in the atmosphere, along with a trend line generated by a global climate simulation model. Figure 3 shows data on average global temperatures since 1950, along with a trend line.

In 1995, a UN scientific panel concluded that these emissions were contributing to global warming, stating that "The balance of evidence suggests a discernible human influence on climate." In 1997 the industrialized nations agreed to stabilize their CO_2 emissions near mid 1990 rates. Implementation, however, remains elusive.

Now let's do a mental exercise we call an extreme conditions test. What do you think would happen to the average global temperature if anthropogenic CO_2 emissions suddenly stopped completely, so that annual emissions were zero? This imaginary scenario is shown in Figure 1. In the year 2000 anthropogenic CO_2 emissions drop instantaneously to zero and remain there forever.

Assume anthropogenic CO_2 emissions follow this scenario. Sketch the likely path (the continuation of the simulation) for atmospheric CO_2 for the next 50 years using the space provided in the right half of Figure 2. Then sketch the likely path (the continuation of the simulation) for the average global temperature for the next 50 years using the space provided in the right half of Figure 3.

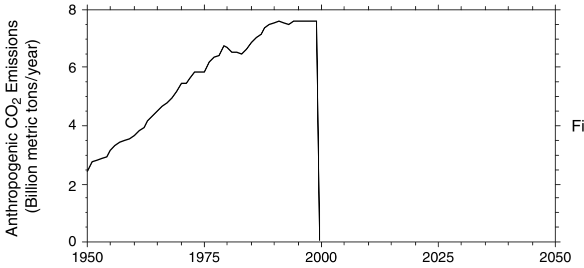

Figure 1: Anthropogenic CO_2 emissions

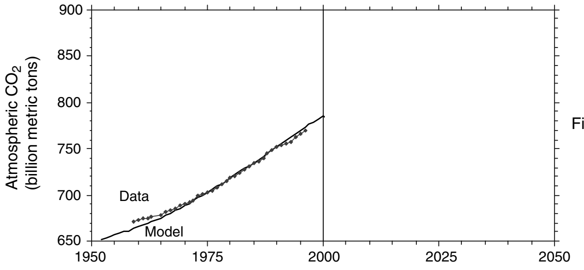

Figure 2: Carbon dioxide in the atmosphere

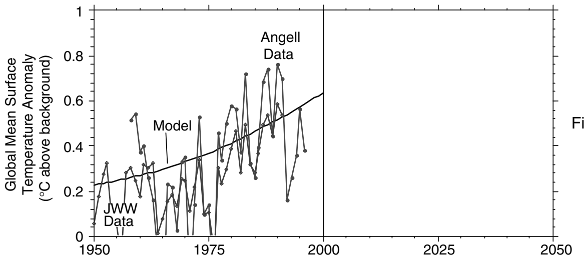

Figure 3: Global mean surface temperature

Fig. 3. The Zero Emissions task (showing the CO_2 graph treatment)

from 2000 to 2050, assuming the hypothesized drop to zero in human CO_2 emissions. In the ‘‘No CO_2 Graph’’ condition, the graph of atmospheric CO_2 was omitted, the text adjusted accordingly, and subjects were asked only to sketch the global mean temperature. We hypothesized that performance in the CO_2 Graph condition would be better than in the no graph condition because explicitly showing the stock of atmospheric CO_2, which determines net radiative forcing and thus the change in global mean temperature, should help subjects think about the stock and flow relationships better than when they are not prompted to consider the stock of atmospheric CO_2.

Of course, no one knows exactly how the carbon cycle and climate would respond to such a shock. Uncertainty about temperature is high; the likely behavior of carbon concentrations is fairly well known. Despite the uncertainties the stock/flow structure of the global climate system and basic laws of physics sharply constrain the possible trajectories. Figure 4 shows the response of Fiddaman’s (1997) model to the Zero Emissions scenario. Though the rate of CO_2 emissions falls to zero in the year 2000, mean global temperature continues to rise for about three more decades. It then falls very slowly.[3]

Simple stock and flow considerations explain how it is possible for the global temperature to rise even after human CO_2 emissions fall to zero. When emissions fall to zero the inflows to the stock of atmospheric carbon fall below the outflows (absorption of carbon by biomass and by the oceans). Therefore the stock of CO_2 in the atmosphere peaks and begins to fall. The concentration of CO_2 in the atmosphere falls only slowly, however. First, the uptake of carbon by biomass falls as the concentration of CO_2 in the atmosphere declines, while CO_2 continues to flow into the air as stocks of carbon in terrestrial biomass and soils, both swollen from absorption of extra carbon from the CO_2 enriched atmosphere, are oxidized by microbial action and fire. Similarly, as atmospheric CO_2 falls, the flux of carbon from the air to the surface layer of the ocean falls, while the flux of carbon from the ocean to the air remains high as some carbon previously absorbed by the ocean and oceanic biomass returns to the atmosphere. These compensatory responses slow the decline of atmospheric CO_2 so that 50 years after human emissions stop completely, the concentration of CO_2 in the model atmosphere has fallen back only to its 1990 level.

Similar stock-flow reasoning explains the dynamics of global temperature. Global mean temperature rises as long as incoming solar radiation exceeds the heat radiated back to space or transferred to the deep ocean. Though falling after the year 2000, elevated atmospheric CO_2 concentrations still cause net radiative forcing. Global mean temperature continues to grow, but at a diminishing rate as the net forcing slowly declines. By about 2030 the surface has warmed enough and the concentration of CO_2 in the atmosphere has fallen enough for insolation to be balanced again by radiation of heat to space and the rate of heat transfer to the deep ocean. Temperature then peaks and thereafter

Fig. 4. Simulation of the Zero Emissions Task. Atmospheric CO_2 peaks when emissions fall to zero in 2000, but mean surface temperature continues to rise for several decades because net radiative forcing, though falling, is still positive. *Source*: Sterman (2000); based on the model in Fiddaman (1997)

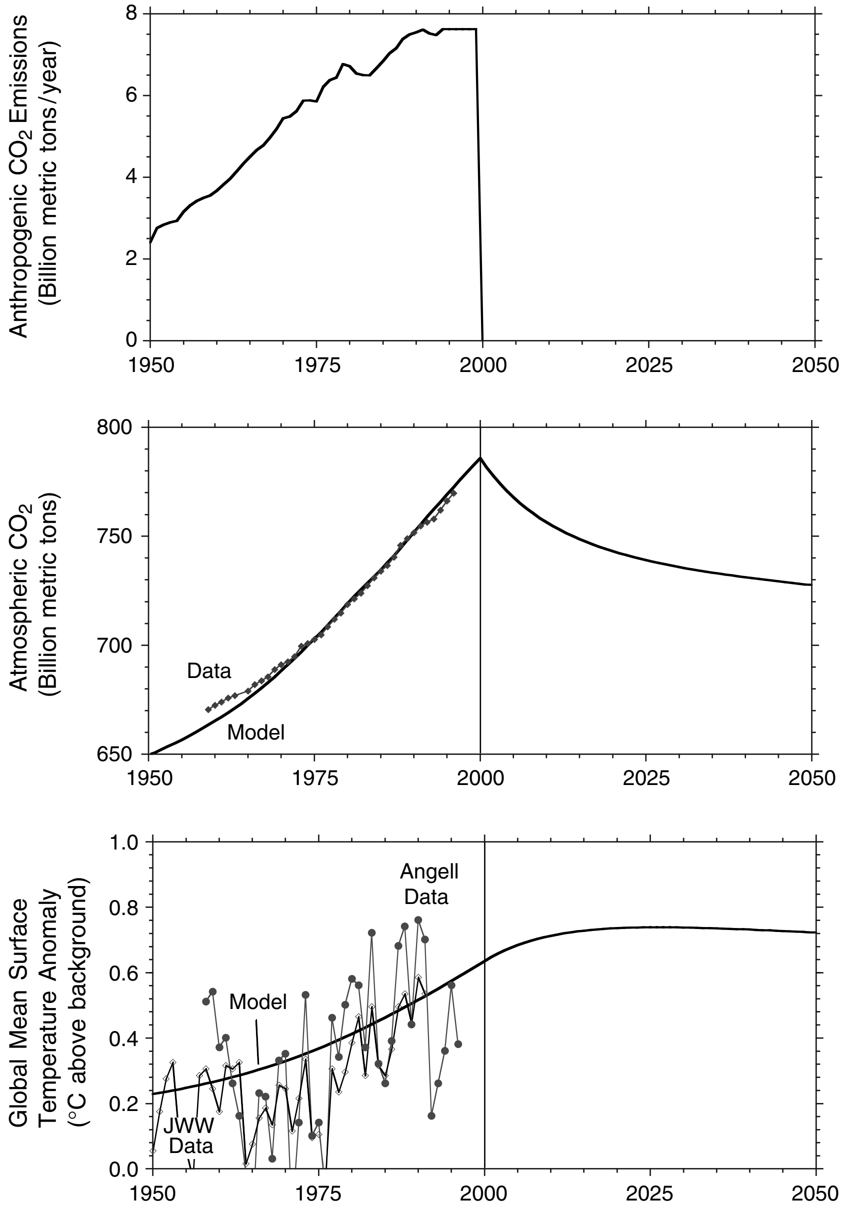

declines. The decline is slow. First, the slow decline of GHG concentrations after 2000 slows the decline in net radiative forcing. Second, as the atmosphere cools, the massive quantity of heat previously stored in the ocean begins to flow back to the atmosphere, slowing surface cooling.[4]

The dynamics may seem complex, but the basic stock/flow structure is straightforward, and indeed, familiar to all. For example, in the temperate latitudes of the northern hemisphere, maximum insolation occurs at the summer solstice, while the hottest surface temperatures occur more than a month later, in late July or early August, when the days are already significantly shorter. That is, from mid June through August the energy input from the sun is falling, but the average surface temperature keeps rising. Though the flux of solar energy heating the surface declines after the solstice as the hours of daylight diminish, it still exceeds the average energy radiated back to space by the earth. Energy inflow exceeds outflow, so average temperature rises even as the input of solar energy falls.

We coded each response into categories describing different patterns, such as continued growth, immediate peak and decline, etc. (Table 2). Because the numerical values in the simulation are uncertain, responses were considered correct if they qualitatively approximated the patterns shown in Figure 4. Specifically, the CO_2 trajectory was considered correct if it peaked at or shortly after the year 2000, then declined. The temperature trajectory was considered correct if it continued to rise, peaked sometime after the year 2000, then declined. In coding responses we were very generous in the timing and magnitudes, marking an answer as incorrect only when a sketch clearly violated one of the basic physical criteria dictated by the stock and flow structure (for example, if a subject showed temperature peaking in the year 2000 and declining immediately).

Even with our generous coding, only 22 percent of the CO_2 trajectories drawn by the MIT students were correct (Table 2). The most common error was to show the stock of atmospheric CO_2 stabilizing in 2000 and remaining constant thereafter. Such trajectories would be correct if the flux of CO_2 into the atmosphere from natural sources was exactly balanced by the absorption of CO_2 out of the atmosphere, or if, as seems more likely, subjects ignored the natural flows and assumed that human emissions are the only flow affecting the stock.

Across both treatments (with and without the CO_2 graph), only 36 percent of the subjects approximated the correct trajectory for global mean temperature (Table 2). The single most common error was to show temperature peaking in the year 2000, at the same moment emissions fall, and then dropping, as illustrated in Figure 5. We had hypothesized that the treatment in which subjects were asked to sketch the trajectory for the stock of atmospheric CO_2 would improve performance on the trajectory for global temperature, but just the opposite is observed: 46 percent of those asked to sketch temperature only were correct, while only 28 percent of the temperature trajectories were correct for those asked to sketch both; the difference is highly significant ($p = 0.014$ by the Fisher Exact Test). It is possible that performance is worse when both graphs are requested because subjects have more to do in the time available. More likely, subjects with a weak grasp of stock-flow concepts may have

Table 2. Performance of MIT students on the Zero Emissions task

	CO_2 trajectory	
1	**Correct:**	
	CO_2 peaks at or very shortly after the year 2000, then declines at a diminishing rate.	22%
	Incorrect:	78%
2	CO_2 stabilizes in or after 2000 and never drops	31%
3	CO_2 keeps rising forever	8%
4	CO_2 immediately drops and continues to go down (shows a sudden, discontinuous jump down at or very shortly after 2000)	4%
5	CO_2 stabilizes, then decreases	7%
6	CO_2 increases, then decreases	16%
7	The CO_2 trajectory is discontinuous (has a sudden jump up or down at some other time than at or very shortly after 2000).	1%
8	The CO_2 trajectory follows some other path.	10%
	$N =$	97

	Global mean temperature trajectory	**Total**	**CO_2 graph**	**No CO_2 graph**
1	**Correct:**			
	Temperature continuing to rise for about 20–30 years, then falls (slowly).	36%	**28%**	**46%**
	Incorrect:	64%	**72%**	**54%**
2	Immediate peak and drop in temperature in or very shortly after the year 2000	22%	23%	20%
3	Temperature rising forever	11%	13%	9%
4	Temperature stabilizing in or after 2000 and never declining	18%	**25%**	**9%**
5	A fluctuation in temperature	1%	0%	3%
6	A discontinuous path (temperature has a sudden jump up or down)	3%	4%	1%
7	Temperature stabilizing then decreasing	2%	**3%**	**0%**
8	Temperature decreasing, then increasing.	2%	2%	3%
9	Some other path for temperature	6%	3%	10%
	$N =$	186	106	80

Notes: Bold figures indicate significant differences between the CO_2 and No CO_2 graph conditions at $p < 0.05$ by the Fisher Exact Test. Totals may not add to 1.00 due to rounding. Number of responses for the CO_2 graph in the CO_2 graph condition (96) was less than the 109 who provided usable temperature graphs in that condition.

assumed temperature should follow the same pattern as CO_2, consistent with the belief that the inputs and outputs in a system should be highly correlated, as observed in Booth Sweeney and Sterman (2000). In fact, 76 percent of all subjects drew CO_2 and temperature trajectories with the same pattern.

The University of Chicago MBA students did about the same on the Zero Emissions task (Table 3). About 64 percent of those in the CO_2 graph condition

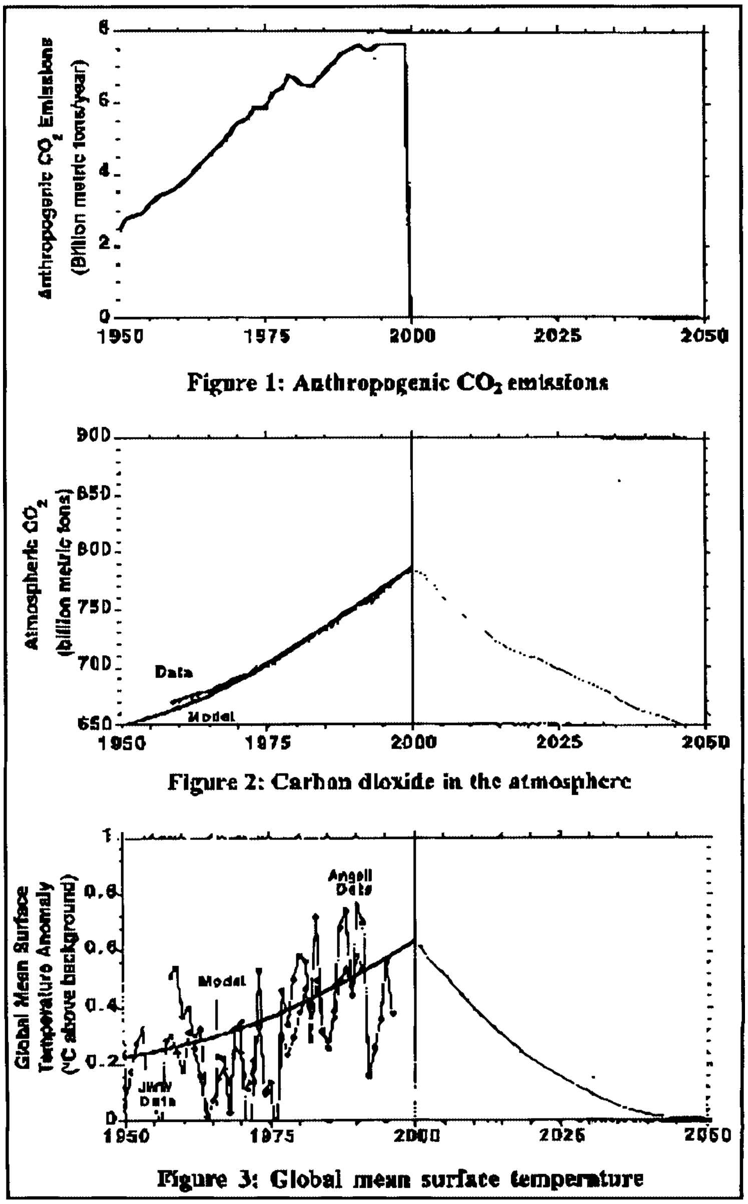

Fig. 5. Typical erroneous response to the global warming task. The subject's atmospheric CO_2 trajectory is qualitatively correct, but the global temperature is shown as following the same path, including a discontinuity in slope in the year 2000. Given the subject's path for CO_2, temperature could not peak until much later, after the CO_2 concentration falls enough to bring net radiative forcing to zero

drew incorrect trajectories (compared to 78 percent of the MIT students, but the difference is not significant ($p = 0.10$ by the Fisher Exact Test)). About 80 percent of the Chicago students receiving the CO_2 graph condition drew erroneous trajectories for temperature, compared to 64 percent of the MIT

Table 3. Performance of University of Chicago MBA students on the Zero Emissions task

	CO_2 trajectory	
1	**Correct:**	
	CO_2 peaks at or very shortly after the year 2000, then declines at a diminishing rate.	36%
	Incorrect:	64%
2	CO_2 stabilizes in or after 2000 and never drops	27%
3	CO_2 keeps rising forever	9%
4	CO_2 immediately drops and continues to go down (shows a sudden, discontinuous jump down at or very shortly after 2000)	2%
5	CO_2 stabilizes, then decreases	7%
6	CO_2 increases, then decreases	9%
7	The CO_2 trajectory is discontinuous (has a sudden jump up or down at some other time than at or very shortly after 2000).	2%
8	The CO_2 trajectory follows some other path.	7%
	$N =$	44

	Global mean temperature trajectory	Total	CO_2 graph	No CO_2 graph
1	**Correct:**			
	Temperature continuing to rise for about 20–30 years, then falls (slowly).	21%	20%	21%
	Incorrect:	79%	80%	79%
2	Immediate peak and drop in temperature in or very shortly after the year 2000	29%	36%	25%
3	Temperature rising forever	20%	22%	18%
4	Temperature stabilizing in or after 2000 and never declining	19%	18%	19%
5	A fluctuation in temperature	0%	0%	0%
6	A discontinuous path (temperature has a sudden jump up or down)	4%	0%	7%
7	Temperature stabilizing then decreasing	6%	4%	7%
8	Temperature decreasing, then increasing.	1%	0%	2%
9	Some other path for temperature	1%	0%	2%
	$N =$	102	45	57

Notes: Totals may not add to 1.00 due to rounding. None of the differences between the CO_2 and No CO_2 Graph conditions were significant at $p < 0.05$.

students; again the difference is not significant ($p = 0.32$). In the No CO_2 condition, 79 percent of the Chicago students drew erroneous temperature trajectories, compared to 46 percent of the MIT students, a significant difference ($p = 0.004$). There were no significant differences in temperature patterns between the CO_2 and No CO_2 treatments among the Chicago MBAs. Like the MIT group, about three-quarters of the Chicago MBAs drew temperature trajectories correlated with their CO_2 trajectories—if they showed CO_2 peaking and then falling, they were also likely to show temperature peaking and falling—providing more evidence that people believe temperature and CO_2 concentrations should follow the same pattern.

The stable CO_2 concentration task

It may be argued that the zero emissions scenario is so unrealistic that people should not be expected to do well. The stable CO_2 Concentration (SC) task presents subjects with more realistic scenarios (Figure 6). The first page presents a short description of climate dynamics, along with data for CO_2 emissions, atmospheric CO_2 concentration, and global mean temperature. The explanatory text was extracted and condensed from the non-technical summary of the IPCC's third assessment report (IPCC 2001a). The description explicitly discusses the rate of removal of CO_2 from the atmosphere by natural processes, stating that:

> Natural processes gradually remove CO_2 from the atmosphere (for example, as it is used by plant life and dissolves in the ocean). Currently, the net removal of atmospheric CO_2 by natural processes is about half of the anthropogenic CO_2 emissions.

Thus the text provides explicit cues encouraging subjects to notice the stock and flow relationship between CO_2 in the atmosphere, emissions, and natural uptake, along with quantitative information about the magnitude of the removal flow.

We defined two scenarios, shown in the second part of Figure 6. In the first, the concentration of CO_2 in the atmosphere gradually rises from current levels of about 370 ppm to 400 ppm by 2100, a rise of about 8 percent. In the second, CO_2 concentration gradually falls by about 8 percent, to 340 ppm. A graph of each CO_2 trajectory is shown. Subjects are then asked to select one of seven possible responses for the pattern anthropogenic emissions would have to take to achieve the specified CO_2 concentration. These ranged from continued growth in emissions through stabilization at current rates to an immediate drop below current rates. They are then asked to select one of seven possibilities describing the most likely path for global mean temperature given the CO_2 scenario they faced. The temperature trajectory options ranged from continued growth through 2100 to stabilization at today's levels to an immediate drop and decline through 2100.[5]

As with the zero emissions task, the stock/flow structure and basic physics constrain possible trajectories for emissions and temperature. According to the IPCC and the description provided to the subjects, the rate at which CO_2 is currently removed from the atmosphere is about half the human emission rate. To stabilize atmospheric CO_2, the emission rate must equal the removal rate. Whether that occurs through a drop in emissions or an

Consider the issue of global warming. In 2001, the Intergovernmental Panel on Climate Change (IPCC), a scientific panel organized by the United Nations, concluded that carbon dioxide (CO_2) and other greenhouse gas emissions were contributing to global warming. The panel stated that ‘‘most of the warming observed over the last 50 years is attributable to human activities.’’

The amount of CO_2 in the atmosphere is affected by natural processes and by human activity. Anthropogenic CO_2 emissions (emissions resulting from human activity, including combustion of fossil fuels and changes in land use, especially deforestation), have been growing since the start of the industrial revolution (Figure 1). Natural processes gradually remove CO_2 from the atmosphere (for example, as it is used by plant life and dissolves in the ocean). Currently, the net removal of atmospheric CO_2 by natural processes is about half of the anthropogenic CO_2 emissions. As a result, concentrations of CO_2 in the atmosphere have increased, from preindustrial levels of about 280 parts per million (ppm) to about 370 ppm today (Figure 2). Increases in the concentrations of greenhouse gases reduce the efficiency with which the Earth's surface radiates energy to space. This results in a positive radiative forcing that tends to warm the lower atmosphere and surface. As shown in Figure 3, global average surface temperatures have increased since the start of the industrial revolution.

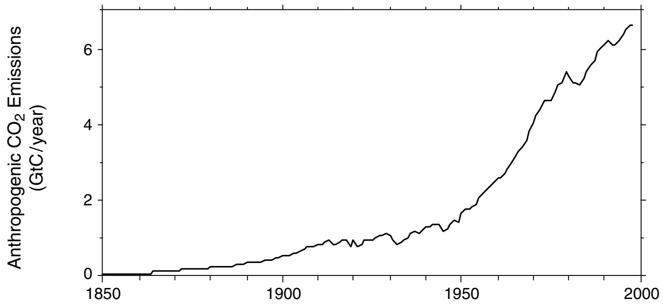

Figure 1. Global CO_2 emissions resulting from human activity (billion tons of carbon per year)

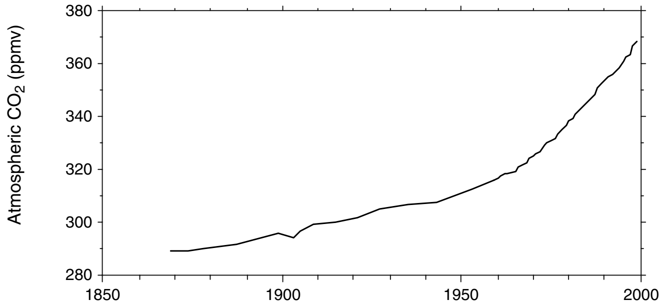

Figure 2. Atmospheric CO_2 concentrations, parts per million

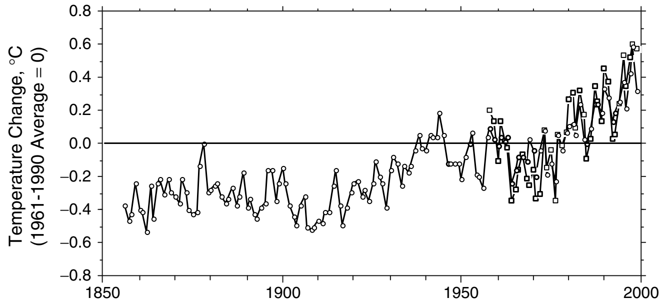

Figure 3. Average global surface temperatures, °C. The zero line is set to the average for the period 1961-1990

Fig. 6. The CO_2 Stabilization task (part 1: description)

Now consider a scenario in which the concentration of CO_2 in the atmosphere gradually rises [falls] to 400 ppm [340 ppm], about 8% higher [lower] than the level today, then stabilizes by the year 2100, as shown here:

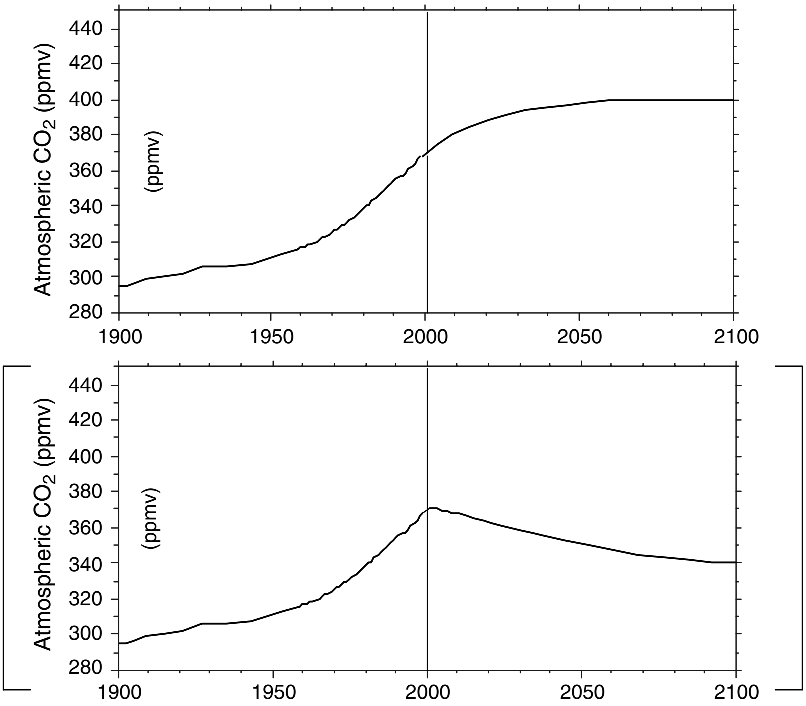

1. For this to occur, CO_2 emissions resulting from human activity would have to:
   - ☐ Continue to rise through the year 2100.
   - ☐ Gradually rise about 8% and then stabilize by the year 2100.
   - ☐ Gradually rise less than 8% and then stabilize by the year 2100.
   - ☐ Stabilize now at current rates.
   - ☐ Gradually fall about 8% and then stabilize by the year 2100.
   - ☐ Gradually fall more than 8% and then stabilize by the year 2100.
   - ☐ Immediately drop more than 8% and then stabilize by the year 2100.
2. Assuming CO_2 concentrations follow the scenario above, the average global temperature would most likely:
   - ☐ Continue to rise through the year 2100.
   - ☐ Continue to rise, then stabilize by the year 2100.
   - ☐ Rise for a few more years, then peak, gradually fall and stabilize above current levels.
   - ☐ Stabilize now at current levels.
   - ☐ Rise for a few more years, then peak, gradually fall and stabilize below current levels.
   - ☐ Rise for a few more years, then peak and continue to fall through the year 2100.
   - ☐ Immediately drop, then stabilize by the year 2100 below current levels.

3. Why? Explain your choices (*briefly*):

Fig. 6. (*Continued*). The CO_2 Stabilization task (part 2: response)

increase in removal depends on auxiliary assumptions subjects must make. There are several possibilities. First, subjects may assume (either because they believe it likely or as a simple approximation that reduces cognitive effort) that removal will remain at current rates. If so, the emission rate must fall by half to yield equilibrium. Alternatively, subjects may assume removal grows as atmospheric CO_2 rises, reasoning that higher CO_2 concentrations stimulate uptake by biomass and the ocean. If so, the most natural assumption is an increase or decrease in removal of about 8 percent, depending on whether they receive the 400 or 340 ppm scenario. Alternatively, sophisticated subjects may reason that future removal will drop as the sinks currently absorbing carbon saturate (as the IPCC projects). In all these cases, CO_2 removal remains far below current emissions. To achieve equilibrium by 2100, emissions must fall by roughly half from current rates—or more.[6]

In fact, climate models project removal will decline below current rates. First, carbon removal from the atmosphere is nonlinear in CO_2 concentrations: As atmospheric CO_2 rises, carbon uptake by aquatic and terrestrial life is increasingly limited by other nutrients or available sunlight so net primary production does not rise in proportion to CO_2. Second, terrestrial and oceanic sinks have absorbed much of the fossil carbon humanity has injected to the atmosphere since the industrial revolution. That carbon eventually cycles back to the atmosphere; full equilibration will take several centuries. Third, warming increases microbial respiration, returning previously sequestered carbon to the atmosphere at faster rates, for example, as tundra thaws (Goulden *et al.* 1998). To stabilize CO_2, emissions must therefore fall below the current removal rate and then continue to drop (IPCC 2001b, Sarmiento *et al.* 1995).

In the 400 ppm scenario, the CO_2 concentration continues to increase after the year 2000, but at a diminishing rather than increasing rate. Emissions must peak in 2000 (the inflection point in CO_2 concentration), then gradually fall to the removal rate. Therefore the correct response for the emissions trajectory is "Gradually fall more than 8% and then stabilize by the year 2100." Considering any decline greater than 8% to be correct constitutes an *a fortiori* coding that favors high subject performance: Many trajectories in which emissions fall by more than 8% would still leave emissions greater than removal so atmospheric CO_2 would continue to rise.

In the 340 ppm scenario, the CO_2 concentration peaks in the year 2000, then falls at a diminishing rate. Therefore emissions must immediately fall below the removal rate so that the net change in atmospheric CO_2 is negative. Emissions must then gradually rise until they equal removals so that the net removal rate gradually approaches zero from below, at which point atmospheric CO_2 stabilizes. Therefore the correct response for the emissions trajectory is "Immediately drop more than 8% and then stabilize by the year 2100."[7]

While the quantitative path for emissions cannot be inferred from the information provided, the qualitative dynamics are not in doubt. The stock and flow structure of emissions, removal, and CO_2 concentrations dictates that emissions must fall from current levels. We hypothesize, however, that subjects with weak understanding of stocks and flows will, as seen in prior work (Booth Sweeney and Sterman 2000), use a pattern-matching heuristic in which system inputs and outputs are assumed to be correlated, and conclude that the path of emissions should look like the path of atmospheric CO_2. Such subjects are likely to select continued growth in emissions in the 400 ppm scenario, and a gradual drop in emissions in the 340 ppm scenario.

The exact trajectory for global temperature in the 400 and 340 ppm scenarios is of course unknown, but again the basic stock and flow structure of the global heat balance sharply constrains the possibilities. Currently, as noted in the description of the SC task, CO_2 concentrations have risen enough to cause a ‘‘positive radiative forcing that tends to warm the lower atmosphere and surface.’’ In the 400 ppm scenario, rising GHG concentrations increase radiative forcing above current rates, causing warming to continue until the temperature of the surface has warmed enough for outgoing energy to once again balance insolation. The task does not provide the quantitative data required to determine whether the temperature would equilibrate by 2100 or later, so correct responses for the 400 ppm scenario are ‘‘Continue to rise through the year 2100’’ and ‘‘Continue to rise, then stabilize by the year 2100.’’ Immediate stabilization, rise and decline, and drops below current levels are all ruled out by the basic physics of the global heat balance.[8]

Temperature change in the 340 ppm scenario is harder for subjects to assess. However, net radiative forcing must become negative for temperature to peak and fall. In the 340 ppm scenario CO_2 begins at current levels of 370, which yield positive net forcing, then gradually declines to 340 ppm. It always remains well above the preindustrial level of 280 ppm, when net forcing was approximately zero and there was no significant warming trend (as seen in the graphs shown in the SC task). Hence the human contribution to net forcing would decline from its current rate, but remain positive. Warming would continue, though perhaps at a diminishing rate. The last four choices in the SC task can be ruled out. Temperature would most likely ‘‘continue to rise through 2100;’’ it might ‘‘continue to rise, then stabilize by the year 2100.’’ The third response, ‘‘Rise for a few more years, then peak, gradually fall and stabilize above current levels’’ is plausible only if subjects interpret the phrase ‘‘a few more years’’ to mean ‘‘at least several decades’’ and/or believe net forcing will drop in the future faster than it rose up to now.

The SC task was given to students enrolled in the introductory system dynamics class at the MIT Sloan School on the first day of the spring term 2002 (Table 4). Overall performance reveals significant deficiencies in student understanding of the stock and flow structure relating emissions and concentrations. Only 18 percent of those receiving the 400 ppm scenario correctly recognized that emissions must peak and fall to equilibrium at the removal rate. Nearly half (44 percent) of those receiving the 340 ppm scenario correctly recognized that emissions must immediately fall more than 8 percent,

Table 4. Performance of MIT students on the Stable CO_2 Concentration task

	CO_2 emissions would have to...	**CO_2 scenario:**	
		400	**340**
1	Continue to rise through the year 2100	6%	1%
2	Gradually rise about 8% and then stabilize by the year 2100	**21%**	**0%**
3	Gradually rise less than 8% and then stabilize by the year 2100	**20%**	**2%**
4	Stabilize now at current rates	14%	7%
5	Gradually fall about 8% and then stabilize by the year 2100	**7%**	**23%**
6	Gradually fall more than 8% and then stabilize by the year 2100	18%	22%
7	Immediately drop more than 8% and then stabilize by the year 2100	**14%**	**44%**
	N	85	86

	Temperature would most likely...	**CO_2 Scenario:**	
		400	**340**
1	Continue to rise through the year 2100	**31%**	**16%**
2	Continue to rise, then stabilize by the year 2100	**38%**	**7%**
3	Rise for a few more years, then peak, gradually fall and stabilize above current levels	26%	15%
4	Stabilize now at current levels	1%	0%
5	Rise for a few more years, then peak, gradually fall and stabilize below current levels	**2%**	**41%**
6	Rise for a few more years, then peak and continue to fall through the year 2100	**1%**	**9%**
7	Immediately drop, then stabilize by the year 2100 below current levels	**0%**	**11%**
	N	84	85

	Emissions trajectory		**Temperature trajectory**	
Scenario	400 ppm	340 ppm	400 ppm	340 ppm
Correct Response	Item 6	Item 7	Items 1 or 2	Items 1 or 2
Fraction correct	**18%**	**44%**	**69%**	**23%**

Note: **Bold** indicates the differences are significant at $p < 0.05$ by the Fisher Exact Test.

then stabilize, a highly significant difference ($p = 0.00024$ by the Fisher Exact Test).

In the 400 ppm scenario, nearly half of the subjects (47 percent) thought emissions would continue to rise. Only 39 percent thought emissions would fall by any amount. In the 340 ppm scenario, only 3 percent thought emissions would continue to rise. These results are again consistent with the hypothesis that people believe the output of a system should follow the same pattern as the input. In particular, many conclude that rising CO_2 concentrations imply rising emissions. The correct response, that for concentration to rise from 370 to equilibrium at 400 ppm requires emissions to fall, appears to be counterintuitive. The higher performance of subjects in the 340 ppm scenario is again consistent with our hypothesis: 90 percent selected an option in which emissions fall, matching the pattern of atmospheric CO_2. Nearly a quarter (23 percent) erroneously conclude emissions gradually fall by about 8 percent, the same fractional drop as CO_2 concentration, and another 22 percent selected a gradual drop by more than 8 percent. These errors are not trivial: a gradual drop by about 8 percent means emissions would continue to exceed removal, so CO_2 concentration would keep rising. A gradual drop by more than 8 percent means emissions remain above removal for at least some period of time, so CO_2 would continue to rise. CO_2 would peak and decline later only if the drop in emissions was large enough for emissions to fall below removal (i.e., much larger than 8 percent).

The subjects' temperature trajectories show little understanding of stocks and flows and are again consistent with the pattern-matching heuristic. In the 400 ppm scenario 38 percent indicate that temperature would most likely ''continue to rise, then stabilize by the year 2100''—the same as the pattern of CO_2 concentration they faced. Another 31 percent indicated continued warming. Using our generous coding criteria, both these responses are considered correct, though the simulations reported by the IPCC suggest continued warming through 2100 is the most likely outcome. In contrast, only 23 percent of the subjects in the 340 ppm scenario chose the correct responses (items 1 or 2: continued warming through 2100 or continued warming to equilibrium above current levels); the difference is highly significant ($p < 3 \times 10^{-9}$). Fully 76 percent select a pattern in which temperature falls (though perhaps after a few more years of warming), and 61 percent select an option in which temperature falls below current levels by 2100.

The written explanations varied considerably, from ''Just wild guesses'' and ''I intuitively came up with these answers, I really cannot explain why'' to sophisticated comments showing clear understanding of the stock-flow structure, as illustrated by this subject in the 340 ppm condition who correctly writes:

> Since CO_2 removal is currently half of emissions, even a substantial drop will result in continued global warming. Thus, for the level of CO_2 to eventually stabilize—this implies that emissions and removal are in balance, which implies an immediate drop in emissions.

A number of written explanations indicate explicit use of the pattern-matching heuristic. A subject in the 400 ppm condition notes ''All of the graphs seem to have the same shape[;] for this reason the reaction (temp[erature]) should follow the action (emission).'' Similarly, a subject in the 340 ppm condition indicated (incorrectly) that emissions would gradually fall by 8 percent and then stabilize, explaining that emissions, CO_2 concentration, and temperature would show the ''same tendency, but a little time gap among the three factors.'' Another said the ''Concentration of CO_2 is positively correlated to emissions but not perfectly.''

Interestingly, many subjects offer detailed explanations showing some awareness of the stock and flow structure, but then reach erroneous conclusions. For example, a subject in the 340 ppm condition concludes incorrectly that emissions would rise less than 8 percent and then stabilize, writing:

> If CO_2 emissions from human activity rose a bit more and then stabilized, the surrounding environment (plant life, algae in the ocean, etc.) would adapt to this equilibrium, allowing atmosphere CO_2 to fall a bit and also stabilize. Temperatures would rise a bit as the earth's ecologic system ''catch[es] up,'' and would then stabilize.

Despite use of words such as ''stabilized'' and ''equilibrium'' the subject asserts atmospheric CO_2 would fall after emissions rose ''a bit more,'' something that is possible only if natural removal more than doubled.

Another in the 400 ppm condition incorrectly asserts emissions would rise about 8 percent and then stabilize (the same pattern as CO_2 concentration), writing:

> 1) Since CO_2 in the atmosphere will gradually rise, while CO_2 net removal will stay approximately at the same level, emission[s] will gradually rise til 2100 then stabilize.
> 2) Global temp. rises as CO_2 concentration rises.

Overall, the written explanations provide further evidence of weak understanding of basic stock and flow structures, and the prevalence of the belief that emissions, concentration, and temperature should be correlated.

Discussion

We presented highly educated subjects with two different tasks testing their ability to understand the most basic stock and flow structures governing the climate. Yet subjects do quite poorly, and their responses exhibit systematic errors. The concentration of CO_2 in the atmosphere must continue to rise as long as emissions exceed uptake; surface temperature must continue to rise as long as incoming solar energy exceeds outgoing radiation and heat transfer to the deep ocean. Subjects, however, regularly violate these basic laws of physics. They consistently underestimate the delay in the response of temperature to changes in CO_2 concentration, selecting trajectories in which temperature responds far too much and too fast. The vast majority believe that temperature should follow the same pattern as CO_2 concentration, rising when CO_2 is rising and falling when CO_2 is falling. In fact, temperature can rise even as CO_2 drops. Similarly, most believe that stabilizing CO_2 concentrations can be accomplished by stabilizing emissions near current rates, when emissions must fall significantly, to the removal rate, for concentrations to stabilize.

Before discussing the implications, we consider alternative explanations for the results. One possibility is that the subjects did not apply much effort because there was insufficient incentive. The role of incentives in judgment and decision making is subtle. In an extensive review, Camerer and Hogarth (1999) found incentives sometimes improved performance, sometimes had no impact, and sometimes worsened performance. It is possible that additional incentive in the form of grades or monetary payment would improve the results. The resolution of this issue awaits future research. It is also possible that the subjects were given insufficient time. This question also must be left for future research. We also note that time pressure is a pervasive feature in many organizations where decisions involving complex dynamics are made. Limited time is the realistic condition in which most people are asked to evaluate information on global warming.

It may also be argued that performance would improve if subjects were given more extensive data and background on climate dynamics. We did not examine variations in information provided. We note, however, that the SC task includes explicit information on the rate at which natural processes remove CO_2 from the atmosphere. Subjects' written explanations show that a number used this information, some effectively. Yet many others explicitly point to the past correlations among emissions, concentration, and temperature as evidence justifying pattern matching, ignoring the cues in the task description increasing the salience of the stock and flow structure. Many more use the information incorrectly, violating conservation laws and reaching erroneous conclusions. Most discussions of warming in the media do not present information on the removal rate by natural processes, and do not draw attention to the stock–flow structure as in our tasks. We hypothesize that performance without this information would be substantially worse than observed here. We also note that in our prior work (Booth Sweeney and Sterman 2000) similar subjects did poorly and exhibited the same systematic violations of conservation laws in much simpler tasks such as filling a bathtub. We certainly hope people would do better if they were first able to take courses in systems thinking and climatology. At present, few do so.

Advocates of the naturalistic decision making movement argue that many of the apparent errors documented in decision making research arise not because people have poor reasoning skills but as artifacts of unfamiliar and unrealistic laboratory tasks. While emphasizing the bounded rationality of human decision making, they argue that people can often perform well in complex decision making settings because we have evolved ‘‘fast and frugal'' heuristics that "are successful to the degree they are ecologically rational, that is, adapted to the structure of the information in the environment in which they are used..." (Gigerenzer *et al.* 1999; vii). Perhaps people understand stocks, flows, delays, and feedback well, but do poorly here because of the unfamiliar and unrealistic presentation of the information. We agree that people sometimes perform well in familiar, naturalistic settings yet poorly on the same type of task in an unfamiliar setting. Our decision making capabilities evolved to function in particular environments; to the extent the heuristics we use in these environments are context-specific, performance will not necessarily transfer to other situations even if their logical structure is the same.

The evolutionary perspective suggests that the errors people exhibit in our tasks should be expected. It is not necessary to understand the relationship between flows and stocks to fill a bathtub—nature accumulates the water "automatically." It is far more efficient to simply monitor the level of water in a tub and shut off the tap when the water reaches the desired level—a simple, effectively first-order negative feedback process, which, experiments such as Diehl and Sterman (1995) show, people can do well. As Laplace remarked, "Nature laughs at the difficulties of integration."[9] That is, stocks in nature always properly accumulate their flows even when mathematicians cannot solve the equations of motion for the system. Thus, for a wide range of everyday tasks, decision makers have no need to infer how the flows relate to the stocks—it is better to simply wait and see how the state of the system changes, and then take corrective action. In such settings intuitive understanding of stocks and flows offers no survival value and is unlikely to evolve.

Unfortunately, the wait-and-see strategy can fail spectacularly in systems with long time delays, multiple feedback processes, and other elements of dynamic complexity (Diehl and Sterman 1995; Sterman 1989a, b). More and more of the pressing problems facing us as managers and citizens alike involve long delays. The long time scale means there is little opportunity for learning through outcome feedback. Instead, we must rely on models of various types to help us project the likely dynamics of the system. These models typically present information in the form of spreadsheets, graphs, or text—the same type of data presentation in our experiments. Managers are called on to evaluate spreadsheets and graphs projecting revenue and expenditure, bookings and shipments, hiring and attrition. These modes of data presentation are not unique to business. Epidemiologists must understand the relationship between the incidence and prevalence of disease, urban planners need to know how migration and population are related, and everyone, not only climatologists, needs to understand how emissions of greenhouse gases alter global temperatures. For global warming, and many of the most pressing issues in business and public policy, the mode of data presentation in our tasks *is* the naturalistic context.

The results suggest that highly educated people have extremely poor understanding of global warming. There are several lessons. Most people drew trajectories in which CO_2 and temperature followed the same pattern—they intuitively feel CO_2 and temperature should be correlated. But the human contribution to global warming cannot be proven or disproven by simply correlating emissions and temperature: The stock/flow structure means climate dynamics are fundamentally incompatible with such naive ''common sense'' approaches. For example, the full impact of past emissions has not yet been observed. Since the industrial revolution the oceans and terrestrial carbon stocks have been absorbing carbon out of the atmosphere at higher rates, suppressing the rise in atmospheric CO_2 concentrations. As these stocks increase, their absorption capacity diminishes. The impact of future emissions on atmospheric CO_2 is likely to be larger than that observed in the past.[10] The inertia of the system means further warming and climate change are already under way. Actions to halt warming must be taken decades before we can know what the consequences of warming will be, and before scientific certainty about the dynamics of the global climate can be gained. Yet many people drew trajectories in which global temperature responds immediately to changes in emissions of greenhouse gases, significantly underestimating the time delays and inertia of the system.

The IPCC considered 40 scenarios for future GHG emissions (IPCC 2001b: 62). In many of these greenhouse gas emissions eventually stabilize or even fall. Yet in every scenario, CO_2 concentrations continue to rise through 2100, to levels ranging from 540 to 970 ppm. In every scenario, net radiative forcing increases, from a low of about 4 W/m^2 in 2100 to more than 9 W/m^2. In every scenario temperatures continue to climb, with warming over the next century of between 1.4 and 5.8 °C (2.5 to 10.4 °F). For comparison, the mean global temperature during the last ice age, when sheets of ice thousands of feet thick covered much of the northern hemisphere, was only about 5 °C colder than today.

At the 1997 Kyoto conference, 38 industrialized nations agreed to reduce emissions to about 95 percent of 1990 rates by 2012. While the agreement is better than business as usual, rapidly developing nations like China are not signatories, and their emissions continue to grow. The policy debate has become a fight over whether to stabilize the emission *rate*, not the *stocks* of greenhouse gases that drive the climate. Even if Kyoto were fully implemented, emissions would continue to exceed removal and GHG concentrations would continue to rise. The fight over implementation of the Kyoto Protocol, therefore, has become a debate about how much more GHG concentrations in the atmosphere will rise, and how much faster the global climate will warm. Halting warming, much less reversing it, is not even on the table.

The situation is, of course, worse. The G.W. Bush administration formally repudiated the Kyoto treaty in 2001. Instead, the administration explicitly adopted the wait and see approach. In February 2002, President Bush proposed the so-called ‘‘Clear Skies Initiative,’’ stating:

> My administration is committed to cutting our nation’s greenhouse gas intensity—how much we emit per unit of economic activity—by 18 percent over the next 10 years. This will set America on a path to slow the growth of our greenhouse gas emissions and, as science justifies, to stop and then reverse the growth of emissions. This is the common sense way to measure progress.... If, however, by 2012, our progress is not sufficient and sound science justifies further action, the United States will respond with additional measures....[11]

Even if achieved, the nonbinding target would not ‘‘stop or reverse’’ warming. It would not lower GHG concentrations. It would not even reduce emissions. It proposes only to slow the *growth* of emissions. Emissions will continue to grow. Greenhouse gas concentrations will therefore continue to rise, and warming will accelerate. The language used to describe Clear Skies plays upon the public’s poor understanding of stocks and flows. By focusing on the GHG intensity of economic output, the proposal can be presented with words like ‘‘cutting’’, ‘‘slow’’ and ‘‘reverse’’. In fact, at the rate of real economic growth assumed in the administration’s 2002 budget, emissions in 2012 would rise by 17 percent.[12] The target 18 percent cut in emissions intensity over 10 years is designed to sound like a large amount, but it is about the same as the 17.4 percent drop in emissions intensity that occurred from 1990 to 2000 with no policy intervention. Clear Skies is the continuation of business as usual. A more accurate statement of the policy would be ‘‘My administration is committed to further growth in our nation’s greenhouse gas emissions, ensuring higher greenhouse gas concentrations and continued global warming. If it turns out that sound science justifies further action, it will be too late.’’

Here lies both a challenge and an opportunity. The challenge: The majority of Americans believe warming is real and should be addressed. But, as we have seen, people’s intuitive understanding of even the simplest dynamic systems is poor. As long as people’s common sense tells them that stabilizing emissions is sufficient there can be little political will or public pressure for policies that could stabilize climate and prevent further warming. As long as people believe the delays in the response of the system are short, they will conclude it is best to ‘‘wait and see’’ if warming will occur and how much more harmful it will be before taking action. Such heuristics often work well in everyday tasks with low dynamic complexity, where delays are short, outcome feedback is unambiguous and timely, the opportunities for corrective action frequent, and the costs of error are small. But none of these conditions hold in systems with high dynamic complexity, where delays between actions and impacts are long, outcome feedback is ambiguous and delayed, many actions have irreversible consequences, and the costs of error are often immense. The same decision making heuristics that serve us well in simple systems may lead to disaster in complex dynamic systems such as the climate.

Initiatives like Clear Skies present two possibilities: Perhaps cynical policymakers and special interests are taking advantage of the public's poor understanding of stocks and flows to make do-nothing policies sound like major initiatives. Alternatively, policymakers themselves may not fully appreciate the dynamics of the climate and sincerely believe a wait-and-see policy is prudent. In free societies, the media are supposed to act as a check on such misinformation. Yet, judging by the poor performance of the highly educated subjects in our experiments, many honest policymakers, members of the media and citizens do not understand the most fundamental principles of dynamics, much less the immensely complicated models developed by climate scientists. People judge the plausibility of model-based projections such as those of the IPCC by whether the projections ''make sense'' relative to their intuitive understanding of the system, and their intuitive ability to relate flows and stocks, understand time delays, and account for simple feedbacks. The results of such models are presented, even in nontechnical reports such as those of the IPCC (2001a), in the form of text, charts and graphs—the same mode of data presentation used in our tasks. Supposedly nontechnical reports such as the IPCC's (2001a) *Summary for Policymakers* are far too technical for the average person—and elected official—to understand. Such presentations are unlikely to overcome the power of ''common sense'' and ''wait and see.'' In such an environment, the corrective feedback of a free press may be ineffective.

The opportunity

Dana Meadows (1989), explaining why she wrote a weekly newspaper column to communicate systems concepts to the public, noted that ''Even the simplest systems concepts help. The level of public discussion is so simpleminded that it doesn't take much to raise its quality. The most fundamental tenets of system dynamics can clear up significant muddles in public thinking.'' While the complexities of the global climate are daunting, the essence of the problem is as simple as filling a bathtub. Learning about stocks and flows may go a long way in overcoming people's poor intuitive understanding of dynamics. To illustrate, consider the following letter to the editor of *The Chemical Engineer*, by A. Lodge (1999), in response to an article about the complexities of global climate policy:

> The article… gives a fascinating insight into the way international politics struggles with complex technical issues. I was inspired to set up an experiment to test some of the ideas, and hit upon the analogy of using my bath instead of the Earth and taking the water as carbon dioxide. I jammed the plug firmly, and turned one tap to full. I observed that the bath was filling with water. I turned the flow down to 80%—a

> massive 20% reduction—only to discover that it was still filling but slightly more slowly. At this point I was joined by my neighbour, an American. He pointed out that reducing the flow by 20% was out of the question; we haggled for a bit before agreeing on a reduction to 94.8%. We thought the 5.2% reduction had a nice ring to it. Oddly, the bath was still filling up with water at almost the same rate that it had been initially. My friend then gave me a five pound note to turn the tap down by another 20%. I did so. He then turned on the other tap to exactly counter the 20% saving. I complained, only to be told that he had ‘‘bought my credits,’’ whatever that means. He then rushed out, returning with a bucket which he put under the second tap. I was so impressed that I did not notice for a moment that the bath was still filling up and that the bucket would soon overflow. We decided we had experimented enough for one day and went off to the pub. We were on our third pint when we remembered that the experiment was still running.

The letter is a fine example of the power of basic stock and flow logic. The higher the concentration of greenhouse gases in the atmosphere, the higher the global temperature will eventually become. The stock and flow structure of the global climate means stabilizing emissions near current rates will not stabilize the climate, but ensures continued growth in GHG concentrations, and continued warming. The climate is complex and our ignorance great. Yet at another level it is as simple as filling a tub: Humanity is injecting CO_2 into the atmosphere at about twice the rate it is drained out. Stabilizing the concentration of CO_2 requires substantial cuts in emissions. Kyoto, while better than business as usual, will not stabilize GHG concentrations at even the record levels they have now attained, much less end the inadvertent experiment in global climate change humanity is now conducting. The sooner people understand these dynamics the sooner they will call for leaders who reject do-nothing wait and see policies and turn down the tap—before the tub overflows.

Acknowledgements

Financial support for this project was provided by the MIT Sloan School of Management Organizational Learning Fund. We thank Jim Doyle, Tom Fiddaman, Jonathan Foley, Jeff Greenblatt, Michael Radzicki, Nelson Repenning, Jorge Sarmiento, Anjali Sastry, and seminar participants at the University of Chicago and MIT for helpful suggestions. We thank Yuval Rottenstreich for administering the experiments with his students at the University of Chicago, and Florian Kapmeier for careful coding and review of the coding criteria. Claudia Forero assisted with data entry.

Notes

1. A Center on Policy Attitudes poll in November 2000 showed that while most Americans believe global warming is real, only 39 percent agreed that ‘‘global warming is a serious and pressing problem [and]

---

we should begin taking steps now even if this involves significant costs.'' 19 percent agreed that ''Until we are sure that global warming is really a problem, we should not take any steps that would have economic costs'' while another 39 percent agreed that ''the problem of global warming should be addressed, but its effects will be gradual, so we can deal with the problem gradually by taking steps that are low in cost.'' (http://www.pipa.org/OnlineReports/GlobalWarming/buenos_aires.html). A Harris poll conducted in August 2001(http://www.harrisinteractive.com/harris_poll/index.asp?PID=256) found that 88 percent of Americans have heard about global warming and most (75 percent) believe it is real (compared to 19 percent who don't). Most who have heard about the Kyoto and Bonn accords approve (by 70 percent to 22 percent), but Republicans much less likely to believe warming is real or support Kyoto than Democrats. See also Immerwahr (1999).

2. This description omits consideration of carbon involved in the formation and weathering of rocks, processes that, like the formation of fossil fuels, operate over a much longer time frame than the injection of fossil carbon into the atmosphere by human activity.
3. The climate change setting necessarily introduces some uncertainty into the experiments because there is no single, unambiguously correct answer to the tasks. The ambiguity is the unavoidable cost of tasks with high ecological validity. Providing additional information on, e.g., CO_2 removal or the stock and flow structure of the climate might reduce ambiguity but would also reduce realism: most information about global climate change available to the general public (and also some of the professional literature) does not provide data on removal or the stock–flow structure. Note also that Booth Sweeney and Sterman (2000) report tasks in which subjects are provided with full information about all relevant stocks and flows; these tasks have unambiguously correct responses. Subjects in those experiments exhibit the same errors and weak understanding of stocks and flows seen here.
4. Temperature change also depends on the concentrations of other GHGs besides CO_2, on particulate aerosols, and on a wide range of feedbacks, both positive and negative, among GHG concentrations, the hydrological cycle, weather patterns, cloud formation and precipitation, surface albedo, ocean chemistry, and many other biogeochemical processes. Consistent with many other climate economy models, Fiddaman implicitly assumes other GHGs such as methane and particulates are proportional to CO_2 emissions. The model also does not include a detailed treatment of feedbacks between, e.g., albedo and temperature. The tasks do not mention these, provide data about them, nor specify their future paths. One might therefore argue that the temperature trajectory could follow any path, depending on what the subjects assumed about these factors. This is implausible for several reasons. First, subjects unaware that there are other GHGs, and unaware

---

of the debate over cloud feedbacks, or the role of black carbon and other aerosols cannot take these factors into account. Subjects who are aware of these factors are also likely to know that emissions and concentrations of other GHGs have grown exponentially since the industrial revolution, following the same qualitative pattern as CO_2. Third, while some other GHGs have greater warming potential per mole, CO_2 is by far the dominant contributor to radiative forcing, accounting for more than half of current net forcing and projected by the IPCC to rise to about three-quarters (IPCC 2001b: 66).

5. We also tested a graphical response mode in which subjects are asked to sketch a graph of emissions rather than select from the list of trajectories. This task, in which subjects were asked only to sketch or select a trajectory for CO_2 emissions, but not temperature, was administered to a large group of subjects in the fall of 2001 under conditions similar to those described here. We coded the sketches into the same categories provided in the multiple choice condition (plus an ''other'' category for those trajectories that differed from any of the options in the multiple choice format). While there are some differences in the response frequencies, the overall distribution of sketch vs. multiple choice responses is similar.
6. Equilibrium could also be achieved by raising the removal rate through technological interventions such as seeding oceans with iron to stimulate plankton growth or injecting CO_2 into the deep ocean. Such interventions are unlikely to work on the scale needed to avoid large emissions cuts, and many argue they would cause their own harmful side effects to ecosystems (e.g., Chisholm *et al.* 2001; Siebel and Walsh 2001).
7. While the IPCC Third Assessment Report does not include simulations of these scenarios, it does report scenarios in which CO_2 concentrations stabilize at values from 450 to 1000 ppm (IPCC 2001b: 75-77). In the 450 ppm scenario, emissions drop shortly after the year 2000 and stabilize by 2300 between half and one-third current rates. Achieving the lower concentrations of 400 or 340 ppm would require much larger and earlier drops in emissions.
8. IPCC (2001b) simulations of the 450 ppm scenario show temperature rising about 1.8 °C by 2100, with continued growth to almost 2.2 °C above current levels by 2350. The 500 ppm scenario yields temperature gains of about 2.1 °C by 2100, and more than 2.8 °C by 2350. Assuming linearity of response (approximately exhibited by the IPCC simulations from 450 to 1000 ppm), the 400 ppm scenario would yield temperature increases of 1.65 °C by 2100 and 1.9 °C by 2350.
9. Quoted in Krutch (1959), p. 510.
10. The IPCC (2001b: 51) concludes ''The fraction of emitted CO_2 that can be taken up by the oceans and land is expected to decline with increasing CO_2 concentrations.... [Despite uncertainties] current models consistently indicate that when the effects of climate change are considered, CO_2 uptake

---

by oceans and land becomes smaller.'' The limits on sinks include a decline in the uptake of CO_2 by forests, including North America, where a great deal of land cleared for farming through the mid nineteenth century has been abandoned and allowed to return to forest over the past century; that process is saturating and may reverse due to development.

11. See http://www.whitehouse.gov/news/releases/2002/02/20020214-5.html.
12. The administration's proposed 2002 budget projects real GDP will grow from an estimated \$9,382 billion in 2002 to \$12,973 billion in 2012, a 38.3 percent increase (a compound rate of 3.24 percent per year); see <http://w3.access.gpo.gov/usbudget/>. U.S. GHG Emissions would therefore grow to $1.383/1.18 = 1.17$ times the projected 2002 rate.

References

Booth Sweeney L, Sterman JD. 2000. Bathtub dynamics: initial results of a systems thinking inventory. *System Dynamics Review* **16**(4): 249–294.

Camerer C, Hogarth R. 1999. The effects of financial incentives in experiments: a review and capital-labor-production framework. *Journal of Risk and Uncertainty* **19**(1–3): 7–42.

Chisholm S, Falkowski P, Cullen J. 2001. Dis-crediting ocean fertilization. *Science* **294**(12 October): 309–310.

Diehl E, Sterman J. 1995. Effects of feedback complexity on dynamic decision making. *Organizational Behavior and Human Decision Processes* **62**(2): 198–215.

Fiddaman T. 1997. *Feedback Complexity in Integrated Climate–Economy Models*. PhD thesis, MIT Sloan School of Management, Cambridge, MA 02142. Available at http://home.earthlink.net/~tomfid/.

Fiddaman T. 2002. Exploring policy options with a behavioral climate–economy model. *System Dynamics Review* **18**(2): 243–267.

Gigerenzer G, Todd P *et al*. 1999. *Simple Heuristics that Make Us Smart*. Oxford University Press: New York.

Goulden M *et al*. 1998. Sensitivity of boreal forest carbon balance to soil thaw. *Science* **279**(9 January): 214–217.

Groves F, Pugh A 1999. Elementary pre-service teacher perceptions of the greenhouse effect. *Journal of Science Education and Technology* **8**(1): 75–81.

Immerwahr J. 1999. Waiting for a signal: public attitudes toward global warming, the environment and geophysical research. American Geophysical Union. Available at http://www.agu.org/sci_soc/attitude_study.html.

IPCC 2001a. *Climate Change 2001: The Scientific Basis. Summary for Policymakers. A Report of Working Group 1 of the Intergovernmental Panel on Climate Change*. IPCC Third Assessment Report. Available at http://www.ipcc.ch.

IPCC 2001b. *Climate Change 2001: The Scientific Basis. Technical Summary. A Report of Working Group 1 of the Intergovernmental Panel on Climate Change*. IPCC Third Assessment Report. Available at http://www.ipcc.ch.

Krutch JW. 1959. The colloid and the crystal. In *The Armchair Science Reader*, Gordan I, Sorkin S (eds). Simon and Schuster: New York.

Lodge A. 1999. Letter to the editor. *The Chemical Engineer* 25 November, p. 25.

Meadows DH. 1989. System dynamics meets the press. *System Dynamics Review* **5**(1): 68–80.

Meadows G, Wiesenmayer R. 1999. Identifying and addressing students' alternative conceptions of the causes of global warming: the need for cognitive conflict. *Journal of Science Education and Technology* **8**(3): 235–239.

Moxnes E. 2000. Not only the tragedy of the commons: misperceptions of feedback and policies for sustainable development. *System Dynamics Review* **16**(4): 325–348.

Nordhaus W. 1992a. *The "DICE" Model: Background and Structure of a Dynamic Integrated Climate–Economy Model of the Economics of Global Warming*. Cowles Foundation Discussion Paper No. 1009. Cowles Foundation for Research in Economics: New Haven, CT.

Nordhaus W. 1992b. An optimal transition path for controlling greenhouse gases. *Science* **258**(20 November): 1315–1319.

Sarmiento JL, Le Quéré C, Pacala SW. 1995. Limiting future atmospheric carbon dioxide. *Global Biogeochemical Cycles* **9**(1): 121–137.

Seibel B, Walsh P. 2001. Potential impacts of CO_2 injection on deep-sea biota. *Science* **294**(12 October): 319–320.

Sterman J. 1989a. Misperceptions of feedback in dynamic decision making. *Organizational Behavior and Human Decision Processes* **43**(3): 301–335.

Sterman J. 1989b. Modeling managerial behavior: misperceptions of feedback in a dynamic decision making experiment. *Management Science* **35**(3): 321–339.

Sterman J. 1994. Learning in and about complex systems. *System Dynamics Review* **10**(2–3): 291–330.

Sterman J. 2000. *Business Dynamics: Systems Thinking and Modeling for a Complex World*. Irwin/McGraw-Hill: New York.