14 An Intelligent Tutoring System for Exploring Principles of Economics

Valerie Shute
Air Force Human Resources Laboratory

Robert Glaser
Learning Research and Development Center

INTRODUCTION

> *I am impressed by the possibility that some experience in discovering principles in a field of knowledge will radically alter the relation the learner perceives between himself and the knowledge, and his way of behaving when he forgets a solution or encounters an unprecedented problem.*
>
> —(Cronbach, 1966, p. 86)

The intelligent tutoring system we discuss in this chapter provides an environment for discovering principles in a field of knowledge. It was designed to achieve this by enhancing students' inductive inquiry skills using the specific subject-matter knowledge of elementary economics for exploring the laws of supply and demand. *Inductive inquiry skills* in this context refers to the students' effectiveness in collecting, organizing, and understanding data, concepts, and relationships in a new domain. This environment uses both discovery learning and more directive approaches: when appropriate, the students are free to explore the domain under study, extracting facts, and organizing principles as they work. When needed, the system can take charge and direct the student in the activities that are most likely to explicate the topic under study.

The implementation of this approach has been developed on a Xerox 1108/1186 Lisp machine, a powerful stand-alone computer that allows self-paced, individualized, and interactive instruction in a rich data source. The plea for such a system can be witnessed as far back as 28 years ago when Suchman (1961) wrote,

> The need for improvement is great. Current educational practice tends to make children less autonomous and less empirical in their search for understanding. . . . The schools must have a new pedagogy with a new set of goals which subordinate retention to thinking. It is clear that such a program should offer large amounts of practice in exploring, manipulating and searching. The children should be given a maximum of opportunity to experience autonomous discovery. (cited in Wittrock, 1966, pp. 37–38).

We hypothesize that discovery learning can contribute to a rich understanding of domain information by increasing students' ability to access and organize information. Knowledge can be viewed as being structured in a network composed of units of "nodes" (concepts) and "links" (relations) (Anderson & Bower, 1973; Collins & Quillian, 1969; Norman & Rumelhart, 1975). Various cognitive processes can create or operate on these structures and activation may spread through the declarative network arousing associated concepts (e.g., Collins & Loftus, 1975). Active exploration should lead to more interconnected links being established between 'nodes' during acquisition, as compared to more passive knowledge acquisition about the same concepts.

We also believe that effective interrogative skills are instructable if the particular skills involved can be articulated and practiced under circumstances which require them to be used. Cronbach (1966) called for assessing the relative effectiveness of learning by discovery versus more didactic approaches by focusing on a narrow problem under limited circumstances. Today, such a circumscribed learning arena is typically called a *microworld environment*. This type of instructional setting can be designed so that students can engage in discovery through determining what data are gathered, making observations, formulating explanatory generalizations, and making experimental predictions. Moreover, a system could be designed so that students could learn from and reflect upon their own knowledge-acquiring activities.

More recently, Sleeman and Brown have commented on the benefits of learning-by-doing where factual knowledge is transformed into experiential knowledge. Tutorial intelligence (i.e., knowing *what* to say, *how* best to say it, and *when* to interrupt the student's problem-solving activity), in conjunction with a microworld environment, can potentially "transform a student's conceptual flounderings and misconceptions into profound and efficient learning experiences, ones rooted in an individual's own actions and hypotheses" (Sleeman & Brown, 1982, p. 2). In experiential learning, as students interact with new subject-matter situations, they compare their observations with their current beliefs and theories. Consequently, these beliefs may be temporarily rejected, accepted, modified, or replaced (see Glaser, 1984). In the course of this developing knowledge, students ask questions, make predictions, make inferences, and generate hypotheses about why certain events occur with systematic regularity.

Such knowledge interrogation is made possible in the present system by having an underlying architecture of knowledge hierarchially organized according to a pedagogical sequence of prerequisite knowledge. This knowledge base can be interactively accessed in various ways by the student. Students can explore the microworld, *Smithtown*, by way of menu-driven options that allow them to: Change the population, change the weather conditions, see sales market information, and so on. Additionally, there are online tools available to help students organize and systematize their information. These tools include: a Notebook to collect data, A Table to sort data, a Graphing package to plot data, as well as a Hypotheses menu to state different variable relationships and conditions under which they hold. There is also an ongoing record maintained of the students' actions (i.e., their menu selections and experiments) which is accessible to both the system and the student. Market simulations that students run comprise "experiments" that allow them to generate and test hypotheses, as well as form generalizations about economic phenomena.

The remainder of this chapter, which is an initial report of this work, is organized as follows.

1. Purpose of this research
2. Parts of the system
    - The Simulation: *Smithtown*
    - Tools for interrogating the system
    - The Knowledge base
    - The Diagnostician
    - The Student Model
    - The Coach
3. Individual differences in interrogative strategies
    - Protocol analysis: Effective interrogative behavior
    - Protocol analysis: Less effective interrogative behavior
4. Evaluation issues relating to the effectiveness of the tutor
5. Summary and future directions

PURPOSE

Students are provided with a microworld environment that they can actively explore. The main purposes of this system are: (a) to investigate individual differences in inductive reasoning and hypothesis generation of the kind involved in scientific inquiry; (b) to study how students' general inquiry skills might be enhanced; and (c) to teach elementary concepts and relationships in microeconomics. The domain of economics is used as an exemplary vehicle to this end. The unique feature of our system, called *Smithtown*,[1] is that in addition to providing a microworld environment for exploring and discovering, it has been designed to be an intelligent, interrogatable microworld where the intelligence is incorporated into the "Diagnostician" and the "Coach" components. These components serve to assist students in becoming more effective interrogators rather than explicitly instructing the domain knowledge.

Smithtown contains a range of levels-of-guidance that can be gradually increased or decreased, depending on the characteristics of the learner at any point in time. The lower end of the range is the more guided environment and this explicitly teaches the interrogative skills that we have determined to be most effective in extracting information from a body of knowledge.[2] The less structured end of the range, or discovery environment, allows the learner to exercise those skills without tutorial intervention.

Students start out in the discovery mode of the tutor; however, it is believed that individuals will show different patterns of exploratory behavior and will differentially benefit from such an unstructured environment. Those individuals who do not do well in the discovery mode are provided with supplemental tutoring in effective interrogation skills. That is, students are placed in the discovery environment until they are unequivocally unsuccessful, at which point the system automatically places them into progressively more structured modes of inquiry.

PARTS OF THE SYSTEM

The system is composed of four main components: (1) the Knowledge Base, (2) the Diagnostician, (3) the Student Model, and (4) the Coach. The knowledge base includes the targeted elements to be learned, such as the *law of demand*, which is an example of economics domain knowledge and *generalization of a concept*, which is an example of inquiry skill knowledge. The Diagnostician is a set of software "critics" that monitor the student's success in applying the inquiry behaviors and learning the domain concepts. The student model is the updated representation of the student's evolving knowledge base, both in terms of economics knowledge and inquiry skills. Finally, the Coach instructs the student based on information provided to it by the student model regarding unsystematic or ineffectual skills. Before discussing each of these components, we provide an

[1]The system is named for Adam Smith who observed in *The Wealth of Nations*, 1776, that the forces of the market are guided as if by an invisible hand.

[2]These interrogative skills were derived from protocol analyses of individuals interacting with the system and are summarized in the section outlining the Diagnostician.

overview of the simulation, and how it appears to the student. This is followed by a description of the inquiry tools.

The Simulation: *Smithtown*

The microworld is a simulated town, *Smithtown*, where the student can participate in various markets and manipulate economic conditions. Student interrogation of the microworld is through certain allowable actions in a scenario which take place in some market where a particular good is bought and sold. Students can select any good or service from a menu and run a market simulation on it; for example, adjusting the market price and observing the repercussions in the market (e.g., changes in quantity supplied or demanded, and in surplus or shortage, etc.) or allow the computer to make a price adjustment. In addition, town or global conditions may be altered, such as per capita income, or the town's population, as well as specific conditions per good, such as the number of current suppliers in town.

To begin an experiment or query, the following flow of menus occurs (see Fig. 14.1): First, a student selects an item from the "Goods and Services" menu. Second, a "Planning menu" pops up containing all possible variables that either may be changed by the student, or that change as a function of one of the independent variables. The student must state the variables they are interested in investigating; that is she must make a statement of intention. This assists the system in understanding and classifying student activities. Third, a menu with relevant economic indicators for *Smithtown* appears. These variables include average income, population, weather, consumer preference index, number of suppliers, and labor costs. Each of these variables has a system supplied default value (e.g., population = 10,000), and the current value for each of these indicators is shown on the screen (see Fig. 14.2: the gauges on the left side of the figure).

After the student examines and/or modifies the indicators, she may see the "Prediction" menu. The student is asked if she would like to make a prediction regarding the outcome of the current investigation. Making a prediction is optional. After observing events and effects in the microworld for some time, the student should be ready to predict outcomes. When the student does make a prediction, the system classifies the student as conducting an experiment rather than just exploring the world. For example, suppose a student was interested in looking at how changing the price of donuts affected the market for donuts in *Smithtown*. The student could collect data on what happened in the market after increasing and decreasing the price of donuts. This would be classified as 'exploratory behavior' by the system. Over time, the student would be in a position to correctly predict that, when the price of donuts increased: (a) the quantity demanded would decrease while (b) the quantity supplied would increase.

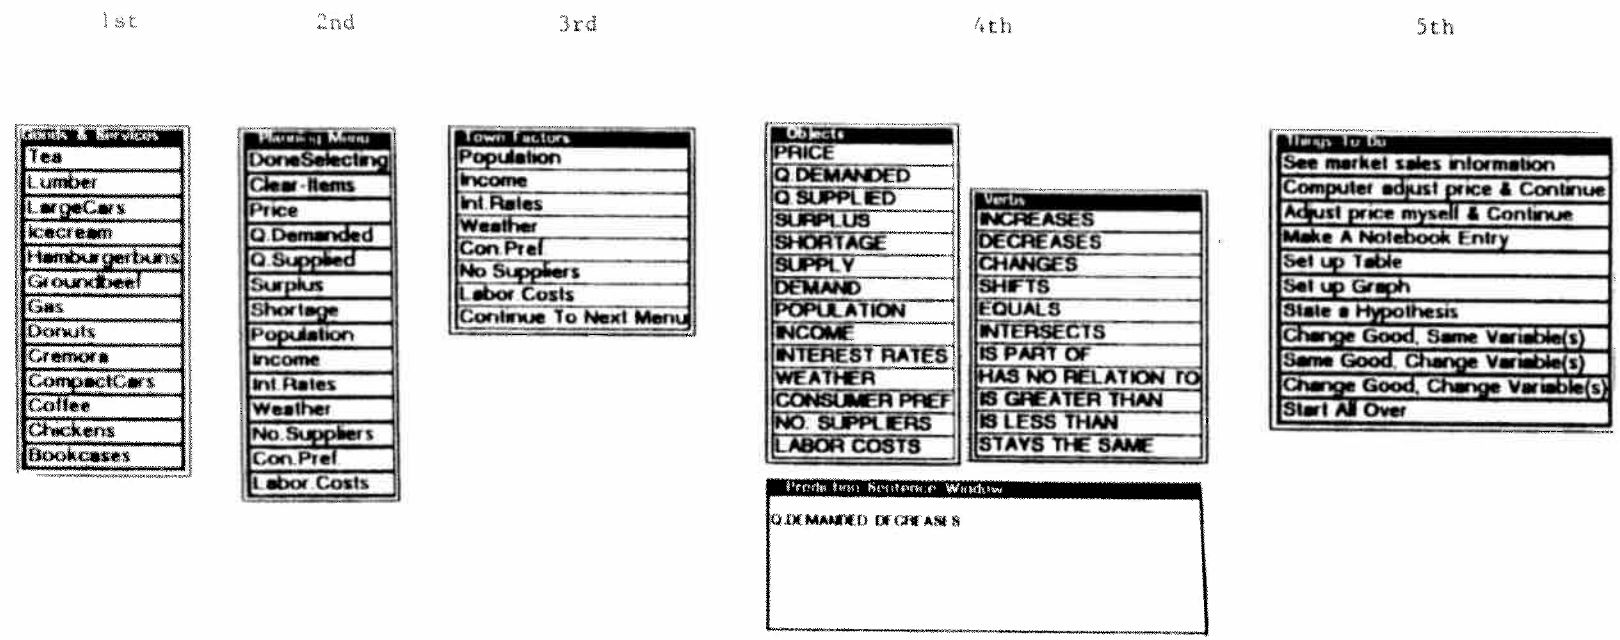

FIG. 14.1. Flow of menus in *Smithtown.*

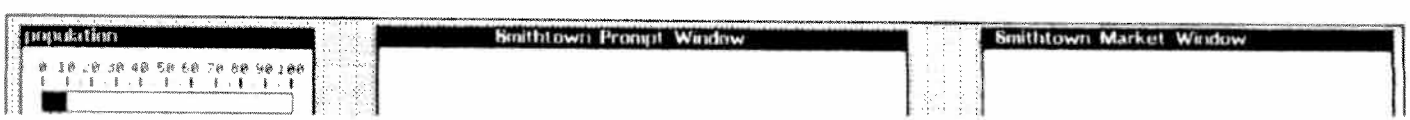

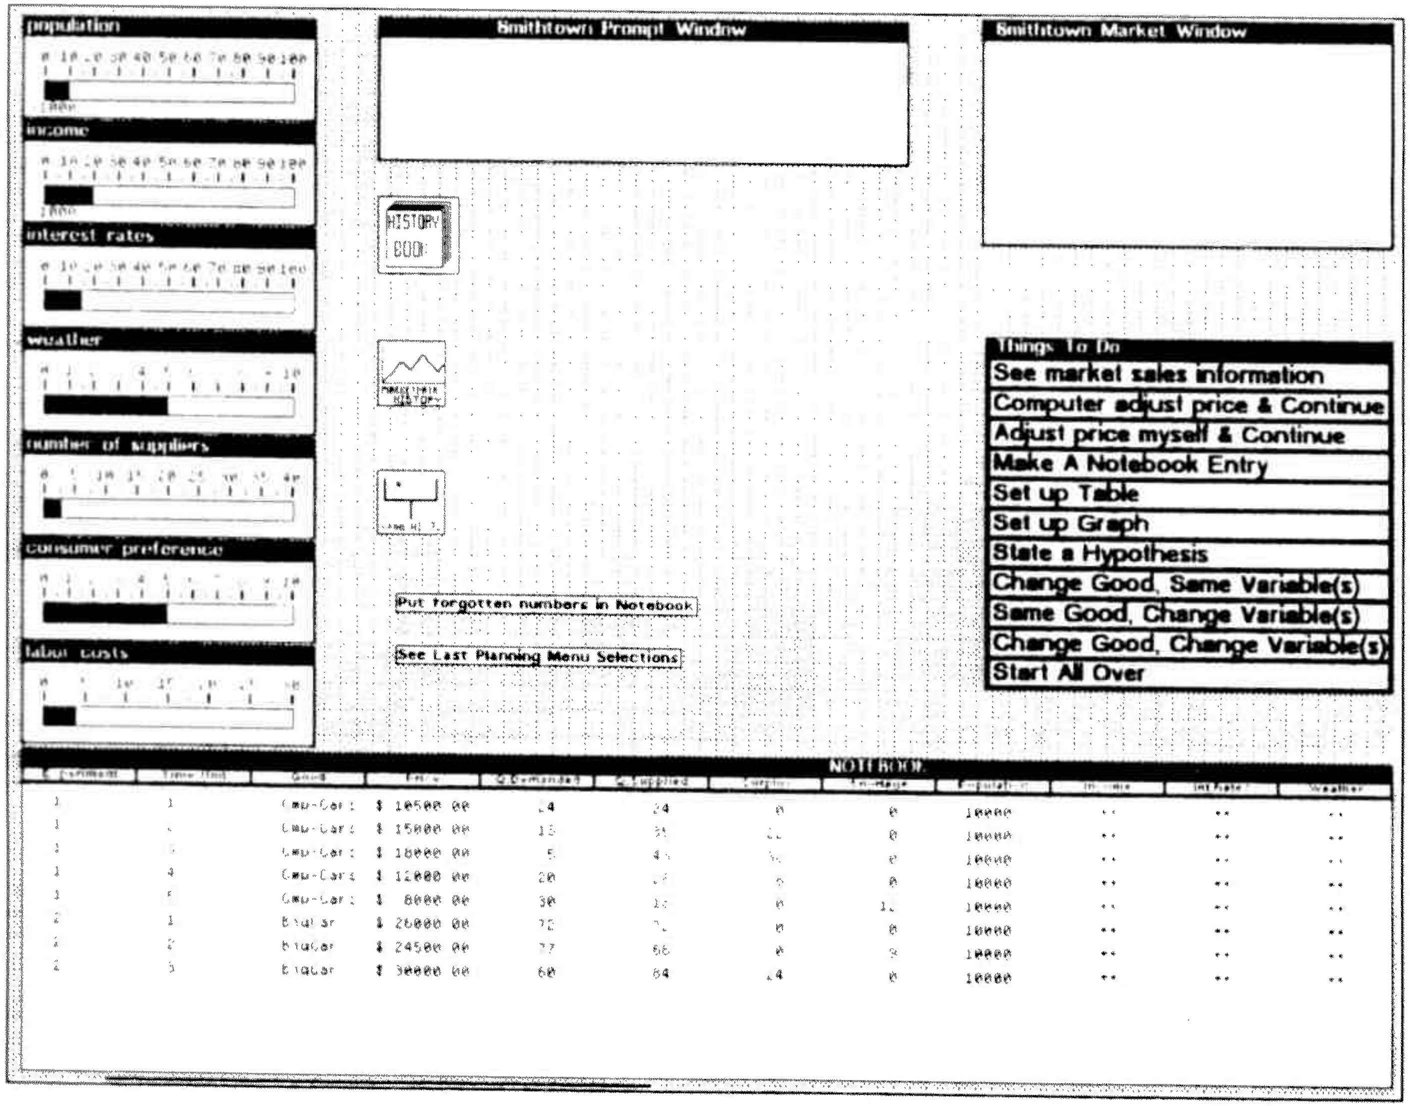

FIG. 14.2. Screen interface in *Smithtown.*

Finally, the student is presented with a "Things to Do" menu where she can see the effect of market manipulations on price, quantity demanded, quantity supplied, surplus, or shortage. At this point, the student may do one of three global actions: (1) Adjust the market price or have the computer make a price adjustment, (2) Use the inquiry tools to assist in the investigation (e.g., graph two variables), or (3) Select an experimental framework. The experimental frameworks let the student manipulate the market in various systematic ways and observe the effects. The three global actions increase in terms of complexity, and we expect to see individual differences in applying them in the microworld, starting off more tentatively changing the price of goods around and gradually learning to employ the various tools and the experimental frameworks for more formal and sophisticated experiments.

An experiment is defined as a series of student actions or tests carried out to see how variables relate to each other in *Smithtown,* or to find out what happens as a result of some specific parameter change. More sophisticated actions are made possible via three experimental frames available to the student:

1. Change Good, Same Independent Variable(s)
2. Same Good, Change Independent Variable(s)
3. Change Good, Change Independent Variable(s)

The first frame can either be a continuing experiment or a new experiment, depending on the intentions of the student. That is, if the student elects to change a good (e.g., coffee) to either a substitute (e.g., tea) or complement (Cremora) while holding the independent variables constant then this is classified as a continuing experiment; otherwise, changing the good to something else, (e.g., large cars), would constitute a new experiment. Changing the good while holding the independent variables constant is an experiment in the generalization of a concept as it holds across various goods. The second frame defines a new experiment where the effects of different independent variables are investigated within a common market (e.g., donuts). Finally, the third frame is a new experiment with different independent variable(s) and different goods. For any experiment, in this sense, the dependent variables in the marketplace that may change are the: market price, quantity demanded, demand, quantity supplied, supply, surplus and shortage.

The demarcation into either *continuing* or *new* experiments presumes that the student attempts to explore the *Smithtown* microworld in a systematic manner. However, often in initial explorations, individuals have no such systematic plan in mind; rather, they randomly change variables and observe the effects. The system recognizes several types of systematic investigations. These include *explorations:* observing and obtaining information from the microworld in order to generate and refine hypotheses about the microeconomic concepts and laws;

experiments: a series of student actions conducted to confirm or differentiate hypotheses; and *exercises:* tests on a previously confirmed hypothesis, perhaps to see the extent or limitations of its application (see Shrager, 1985, for a similar demarcation). Experiments carry a specific prediction from the "Prediction" menu while explorations do not.

Given this framework for collecting student actions, we are interested in monitoring changes in these experiments over time. To do so, we employ the Diagnostician component, described in a later section. Now we discuss the inquiry tools that the system makes available to the student to assist in his or her investigations.

Tools for Interrogating the System

As indicated, we have included several "online tools" for the students to use in their explorations of this microworld. These include: a Notebook for collecting data and observations, a Table to organize data from the notebook, a Graph Utility to plot data, a Hypothesis Menu to state relationships among variables, and three History Windows that allow the students to see a chronological listing of behaviors, data, and concepts learned so far. Each of these tools is now discussed in turn.

The students can keep a Notebook of data about their explorations of *Smithtown*. For the user, this is optional, but the computer always keeps a history of all data that have been collected. An example of the online Notebook is seen at the bottom of Fig. 14.2. The students can select which variables they want recorded, and the values are automatically put into the notebook.

Once students have collected data in the Notebook, they can elect to isolate some of the data and put them into a special Table where various sorting tools for reordering the entries are available (see Fig. 14.3). When students want to separate out some variables for more intensive study, they choose the option of setting up a Table from the "Things to Do" menu. The screen is cleared, and they can specify which variables they would like to put together (e.g., market price and quantity supplied). This is an important tool for reducing and making sense of raw data. For instance, after collecting and recording data on the good, compact cars, a student may begin to sense the relation between the price of compact cars and the quantity supplied (i.e., as the market price increases, the quantity supplied increases). To test this developing idea, he or she can have the system isolate these two variables in the Table and sort one of the variables by ascending (or descending) order. Once this is done, the function relating the two variables should be clearer. This table reveals suppliers' behavior and allows for a possible hypothesis of an economic principle to emerge; that is, the law of supply.

The Graph utility allows the student to plot data, either from the notebook or from the table that has been constructed. This is included as an alternative way of

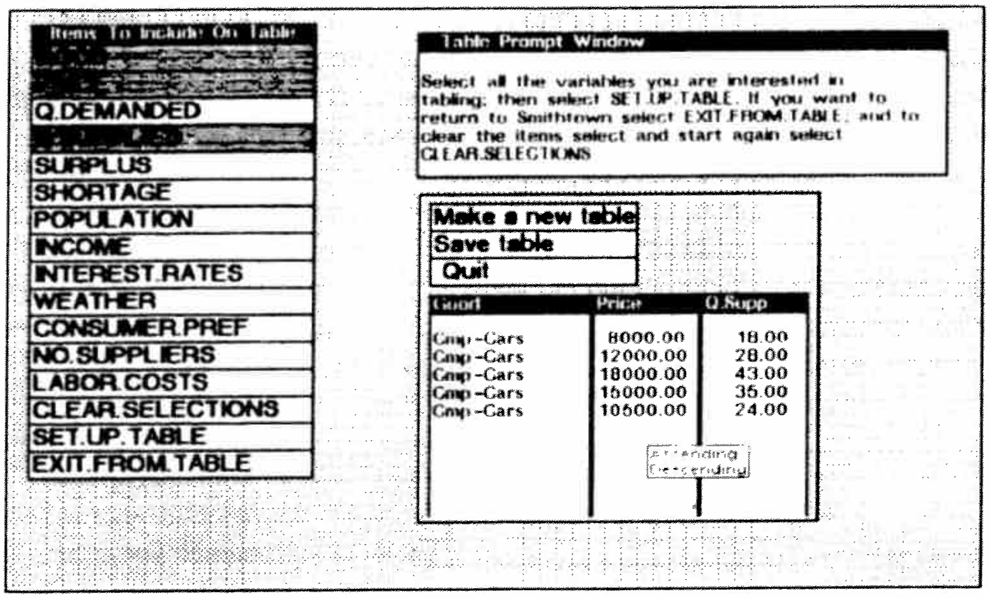

FIG. 14.3. Table package.

viewing relations between variables. Some individuals can better understand information when it is depicted in a graph than in tabular form. To apply this utility, students must select the menu option: "Set up a Graph" from the "Things to Do" menu, informing the system of the variables they wish to plot. An example of a graph with two variables plotted together can be seen in Fig. 14.4. The system has two other options in the Graph utility: (1) Save a Graph, and (2) Superimpose Graphs. The first option saves a particular graph by shrinking it down to a small window, preserving the graph functions on the outside of the window. The shrunken window can be moved and placed anywhere on the screen by the student, and enlarged again by just buttoning within its region.[3] The second option lets the student plot two or more graphs together, providing they share a common independent variable on one of the axes. This allows students to see curve shifts, interactions, and so on.

A framework is provided for the students to make inductions or generalizations of relationships from the data in the form of a Hypothesis menu (see Fig. 14.5). When a student believes she is ready to state a hypothesis, she can choose this option from the "Things to Do" menu. As with the Table and Graph options, the screen clears and the instructions to the student appear, prompting on how to select words from the following four interconnected menus to construct a statement about some variable relationships. One menu consists of *connectors* such

[3] 'Buttoning' refers to the action taken with the 'mouse', the device which lets students make menu selections by pressing one button on its top.

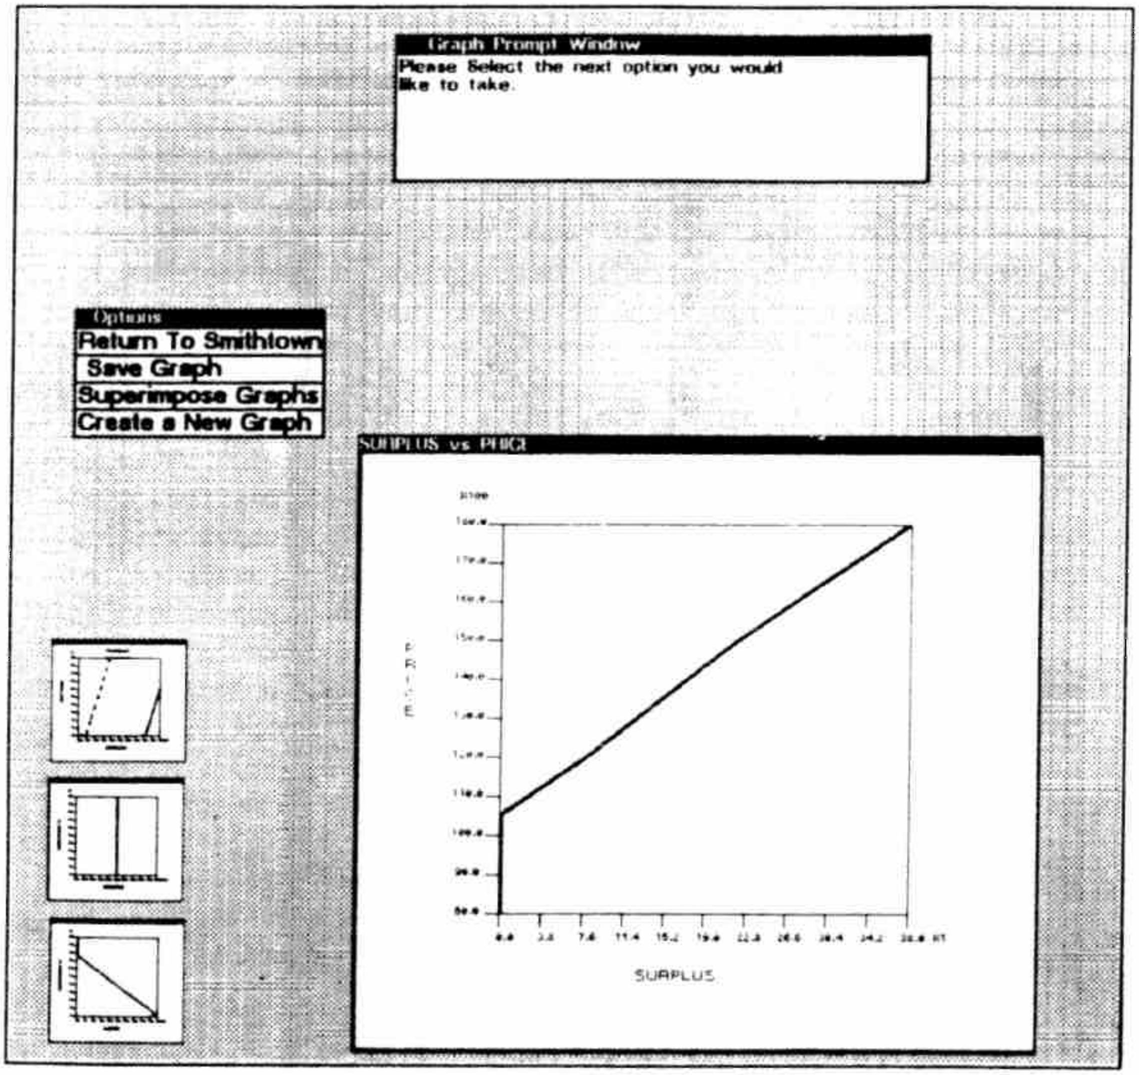

FIG. 14.4. Graph utility.

as: if, then, as, when, and, the. Another menu consists of *market variables* such as: income, population, quantity demanded, demand, quantity supplied, supply, market price, surplus, shortage, and so on. A third menu consists of *direct objects* including: over time, down/right, up/right, down/left, up/left, along the demand curve, along the supply curve, zero, left, right, price changes, changes other than price, and changes to. Finally, the last menu includes *verbs* describing the change: decreases, increases, equals, intersects, is part of, has no relation to, is greater than, is less than, slopes, moves, shifts as a result of, changes as a result of. As students choose words from these menus, the emerging statement appears in the Hypothesis Statement Window. For the example given above in the market for compact cars, a student could state, "As price increases, quantity

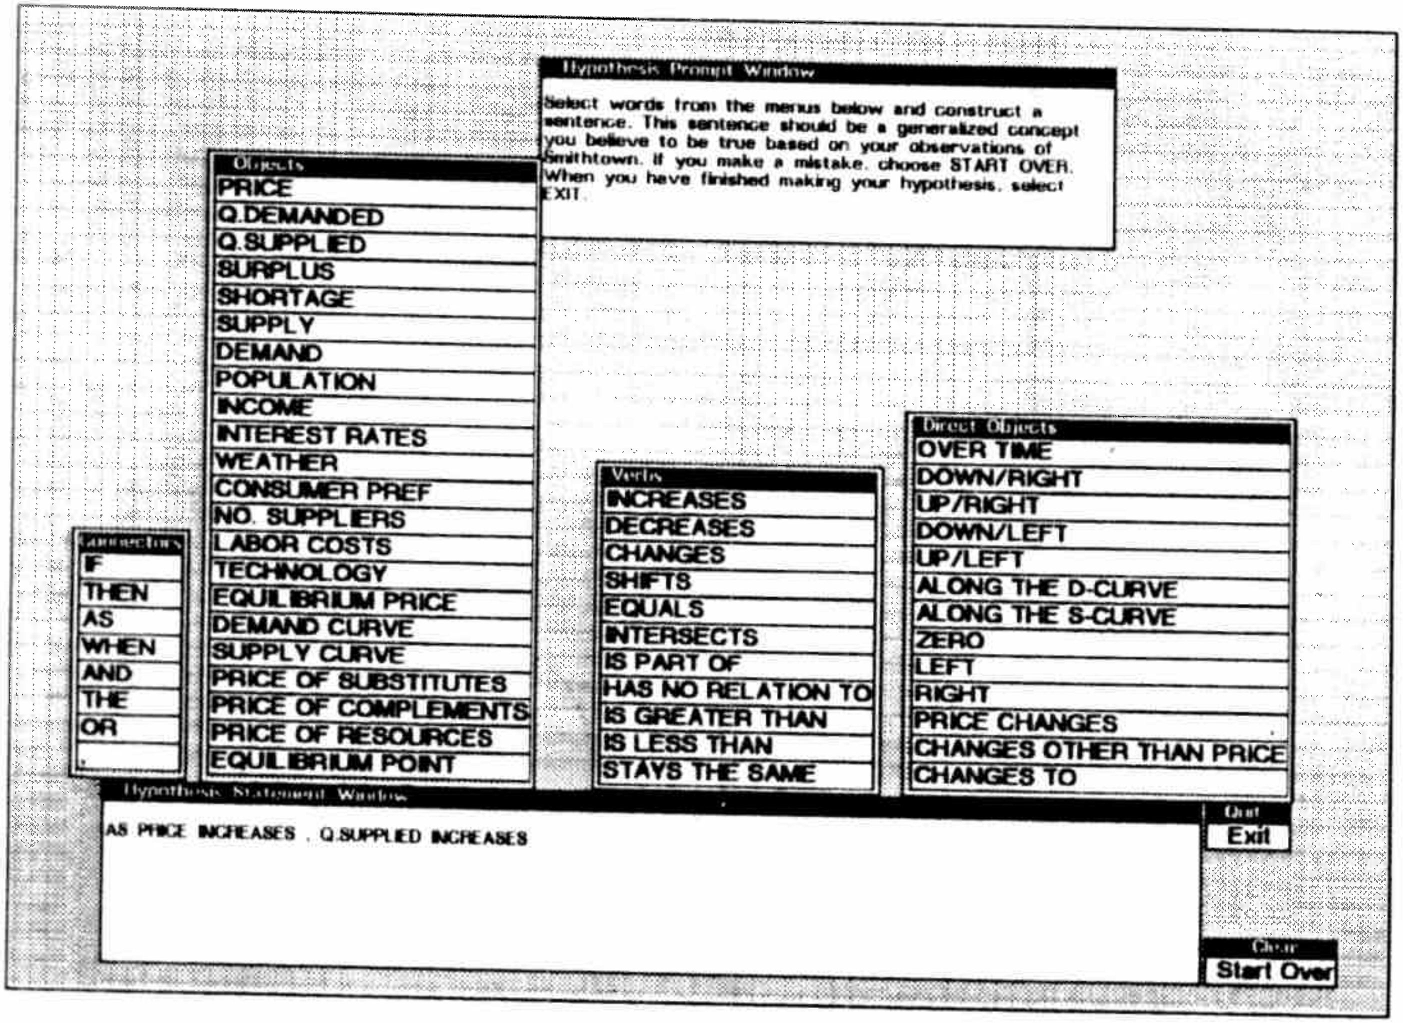

FIG. 14.5. Hypothesis menu.

supplied increases".[4] A pattern matcher analyzes key words from the input and checks to see if this matches any stored relationships for each targeted concept. If so, the system flags that concept as having been conditionally learned. Otherwise, the student is informed that the statement is not understood.

Three history windows are included in the system. As students continue to interact with the microworld, histories accumulate summarizing the various student actions resulting from different explorations and experiments. This summary is maintained in the *Student History Window* and is accessible to both the student and the system. The grain size for this summary is at a small enough level to give a detailed chronological accounting of all menu items chosen, predictions and hypotheses made, and how the tools were employed (e.g., STUDENT SET UP A GRAPH PLOTTING QUANTITY DEMANDED ON THE X AXIS AND PRICE ON THE Y AXIS). The *Market Data Window* keeps a record of all variables and associated values that have been manipulated. If a student forgot to enter some values in the notebook, she may go to this window and retrieve the necessary values for any given experiment or time unit within an experiment. Finally, there is the *Goal History Window*. This provides a representation of what the student has successfully learned in terms of concepts targeted by the system. Moreover, there is a listing of the concepts not yet learned. The student can see this list at any time by enlarging this window. The concepts already learned are shaded while those concepts remaining to be learned are left untouched. This provides a means for the student to gauge progress in the acquisition of relevant concepts.

Given these tools, the system can judge how effectively the student applies them in the interrogation of the microworld. Since coaching assistance is embedded in the domain knowledge, we now present the organization of the domain knowledge.

The Knowledge Base

The tutor has a well-defined instructional domain which is broken down into key concepts that are organized in a bottom-up manner (i.e., from simpler to more complex ideas). An understanding of these concepts should result from the students' experiments in the microworld.

The hierarchy of domain knowledge was developed by first reviewing six introductory microeconomics textbooks and determining the presentation order of information, and second, discussing the optimal ordering of these concepts for student learning in the classroom with a college instructor of economics. Although a student is not required to learn the concepts in any prescribed order, the hierarchy shown in Figure 6 provides the system with information about where

[4]This is equivalent to expressing the relation as, "Quantity supplied increases as the market price increases".

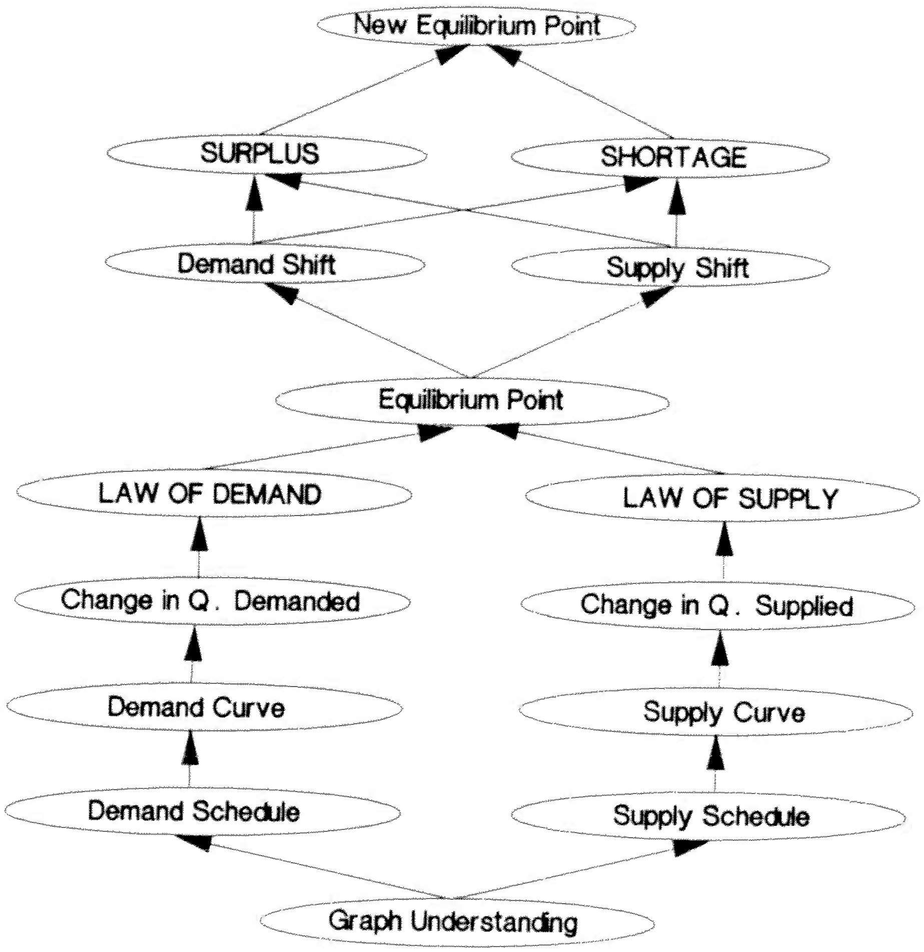

FIG. 14.6. Hierarchy of microeconomic concepts.

the student is likely to be with regard to his or her knowledge acquisition. That is, the concept of "equilibrium point" can be more readily understood after the individual concepts of supply and demand have been learned, both of which have their own prerequisite knowledge. Appendix I contains a summary of the content, concepts, and model underlying *Smithtown*. The knowledge base for the inquiry skills is delineated in the following section.

The Diagnostician

Smithtown was designed to provide a range of guidance that can be increased or decreased, depending on the characteristics of the learner's performance. At one end of the range is the discovery environment. As long as the student is progressing, the microworld will remain in a discovery mode. "Progress" is defined as (a) demonstrating appropriate investigative behaviors, and (b) learning the domain concepts in a reasonable amount of time. At the other end of the range is the directive environment that explicitly assists the student in using an interrogative skill that is deemed problematic for him/her. For example, the system might instruct a student to enter data in the notebook from the market before altering any variables.

The Diagnostician evaluates how a student interrogates the system, ascertaining whether he or she is proceeding in a systematic, efficient manner. The Diagnostician works in conjunction with the Student Model and Coach. The Coach intervenes when necessary. We characterized students' effectiveness in interrogative strategies over the course of their experiments by analyzing the protocols of effective and less effective behaviors (see section on Protocol Analysis). To do this, an explication of effective inquiry behaviors is required, as well as a definition of a test, which specifies the students' experiments within the microworld and determines their systematicity or buggy version of a test. We call these tests "critics." With this information, the Diagnostician can monitor how the student performs in the microworld. Using the inquiry tools provided, there are various dimensions on which student investigative performance can be evaluated. Following is a listing of inquiry behaviors:

Exploratory Data Collection

Baseline Data Collection: For an initial exploration, were data collected from the market in equilibrium, before any variables were altered?

Baseline Data Entry: For an initial exploration, were data entered into the notebook from the market equilibrium, before any variables were altered?

Monitoring the Market—Computer Price Change: Did the student allow the computer to make adjustments to the market price of a good (to see ensuing repercussions)?

Planning for an Experiment

Entry of Planning Menu Items: Did the student enter data into the notebook for those variables she expressed an interest in investigating from the planning menu?

Systematic Notebook Entries

Independent Variable Data Entry: Did the student make a notebook entry every time she changed a variable or variables?

Dependent Variable Data Entry: Did the student enter those variables that changed as a result of a manipulation?

Data Organization

Table Usage: Did the student use the table package to organize data that were entered in the notebook?

Isolating Variables in the Table: Did the student put only a few variables into the table to reduce and make sense of the data?

Sorting on Relevant Variables: Was the sorting option used on relevant variables in the table? For example, if the price was systematically varied, then it should have been sorted instead of something not varied, like income.

Graph Package Usage: Was the graph package used by the student to plot data?

Plotting Variables: Did the student select variables to plot that had been manipulated, or that changed as a result of a manipulation?

Saving Graphs: Did the student save a graph after plotting variables that had been changed?

Superimposing Graphs: Was there an attempt made to superimpose two graphs to see relationships between functions, like supply and demand, or two curves in parallel?

Systematic Experimental Behaviors

Manipulating the Market Price/User Price Change: Did the student make adjustments to the market price of a good to see the market changes? For experiments, manipulating the price oneself is probably more effective than observing the computer adjust the price.

Sufficiently Large Change to Variables: Were initial changes made to variable(s) sufficiently large enough to detect any market effects?

Sufficiently Small Change to Variables: Were later changes made to variable(s) sufficiently small enough to discriminate and refine patterns in the data?

Number of Variables Changed: Did the student change only one variable at a time for comparison and/or recording? This is related to what a student already knows; over time, a student can progressively handle more variables.

Inductive/Generalization Strategies

Replicating an Experiment: Did the student attempt to redo an experiment upon getting confirming or negating results for a given prediction?

Generalize a Concept Across All Goods: Did the student try to generalize an economics principle as it holds for all goods? For example, having learned the law of demand operating in the donut market, did the student attempt to generalize the principle for *all* goods?

Generalize a Concept Across Related Goods: Did the student try to generalize an economics principle across specifically related goods, like between tea, coffee, and Cremora, or between lumber and wooden bookcases?

Hypothesis Specification: Were any hypotheses stated as a result of observed systematicities in the data from the Notebook, Table, or Graph?

Complexity of Hypotheses: Was there an increase over time in the chaining of variables when the student generated hypotheses? As knowledge increases, the number of variables strung together should go from two or more, in progressively more complex relationships.

Strategies Used With Negating Evidence

Failure-Driven Behaviors: If an experiment was conducted testing some specified belief or prediction, and the results of the experiment were not supported, did the student: (a) Redo the experiment with parameter changes, attempting to get the data to fit the existing theory, or (b) Generate a new hypothesis that fit the observed data?

Each of these specific actions has been translated into a critic which is a test of sequences of student actions. In addition, there are points in each critic where the student might go astray constituting unsystematic behaviors, or buggy versions of the critics (see asterisk in example, below). For instance, the critic for the first inquiry behavior would look like the following (paraphrased):

BASELINE DATA COLLECTION: CRITIC 1

If: It is an exploration or experiment, AND
It is the first action in this investigation, AND
A good or service was selected, AND
From the Planning menu there is a list of variables interested in, AND
* From the Economic indicators menu there are no variables changed

Then: BASELINE DATA COLLECTION: CRITIC 1 has been demonstrated.

Otherwise: Alert the Coach for possibly helping the student, depending on the value of the Critic-counter which tallies the number of times the student erred on this skill (applied a buggy version of it) or failed to apply it when it was relevant. The Coach then may help the student according to the appropriate level of explicitness required.

The Diagnostician evaluates an experiment in terms of the values assigned to each specific critic and indicates to the Coach whether or not a given sequence of student actions has reached some predetermined level of adequate performance.

Because the system knows when each concept has been learned (i.e., when the student specifies a correct economic principle), and there is a fixed set of concepts, the system can ascertain relationships between specific inquiry behaviors and the accuracy and quantity of concepts being learned. That is, some skills may be more powerful and necessary than others in the extraction of domain information. So, the Diagnostician characterizes students in terms of their declarative and procedural knowledge at any point in time. For *declarative knowledge*, the system knows when concepts the student currently understands, as evidenced by a successful statement about variable relations and conditions under which they hold. For *procedural knowledge*, the system knows which and how the student applies the various tools to extract information, and how she creates tests for conducting different experiments on the data. Specifically, the Diagnostician determines how a student is: collecting data and observations from experiments (Notebook entries); isolating out specific variables for study of relationships (Table or Graph); testing the emerging hypotheses by way of different experiments; and specifying variable relationships (Hypothesis menu).

The Student Model

The hierarchy of concept elements provides the basis for a model of student knowledge. Each concept is associated with a rule or relationship among economic variables. For example, the law of demand relates price and quantity demanded in an inverse relationship. When a student uses the hypothesis menu and generates a valid statement about the underlying variable relationships for a particular concept, the system flags that concept or relationship as having been learned. This appears in the Goal History Window, checked off from the curriculum list of domain elements to be learned.

Similarly, a student model is constructed for effective usage of the inquiry tools. Actual student performance is compared to optimal performance in an overlay sense (Carr & Goldstein, 1977). The effectiveness of performance is determined by two statistics maintained for each student as he or she interacts with the microworld: (a) Demonstrating an inquiry behavior when it was appropriate; or (b) Applying a buggy version of that behavior. A "batting average" is computed for each behavior consisting of the number of times the skill was used, divided by the number of times it *should* have been used. If that number is less than some threshold (e.g., <50%), the Coach would be prompted to intervene. Also, if several inefficient behaviors were demonstrated, the Coach can address each one in turn based on a hierarchy of coaching where, for instance, the behavior for *Baseline Data Collection* takes precedence over *Save a Graph,* and so on. The way the Coach actually operates is described next.

The Coach

The purpose of the Coach is to take the information provided by the Student Model, and if there is evidence of student floundering, to act appropriately. An example of less effective interrogative behavior would be if the student consistently made only slight changes to variables being investigated during preliminary experiments (i.e., not employing the critic: *Sufficiently Large Change to Variables*). Such manipulations, as changing the population from 12,000 to 12,500, or the number of suppliers of a particular good from 10 to 11, are typically too small to result in any real changes in the marketplace. The Coach's intervention would consist of a suggestion to try to make the changes larger, and hence, more observable.

For each student experiment, the Diagnostician assesses the specific inquiry actions, then updates information in the Student Model. Those actions failing to reach the predetermined level of success or criterion performance would thereby become candidates for possible assistance. Because our Coach is relatively unobtrusive, it requires several corroborations of failure (i.e., ineffective behaviors across successive experiments) before actively intervening. In other words, it maintains an ongoing list of student problems and, when the evidence for floundering is unequivocal, it intervenes. In the event that there are several ineffective behaviors, the Coach makes a decision as to which one needs addressing first. Again, it addresses the students' interrogative behaviors and not the domain knowledge, *per se*. Assistance is provided in the context of the student's current investigation.

In order to encourage an independence of thinking, tutorial guidance progresses in terms of explicitness. For the initial intervention, the Coach provides a subtle hint. The next time it is called for the same behavior, the intervention is clearer, providing analogies to the student to map onto the current situation. Finally, the third intervention provides explicit instructions for conducting a particular experiment. For example, if the student continued to alter variables before recording the results of the market prior to the changes, the Coach would intervene with the following, varying levels of explicitness:

> *Explanation 1 (vague):* You changed ⟨variable = X⟩. If you are interested in seeing how the change affects the market, what do you need to do before you begin making changes?
>
> *Explanation 2 (analog mapping):* You changed ⟨variable = X⟩. If you wanted to measure the affects of some drug on performance, first you would have to find out what the performance level of the subject was *before* the drug was given. Apply the same logic here.
>
> *Explanation 3 (explicit):* You changed ⟨variable = X⟩ without recording its value before the change. You should choose the 'Continue To Next Menu' option from the Town Factors menu and then make a notebook entry of the current market information before changing the variable's value. This gives you baseline data to make comparisons against later on.

When a student has conducted several experiments and discovered an economic principle, she may use the hypothesis menu to generate a statement about the underlying variable relationships for that concept. If the student is successful, she is provided with a congratulatory statement by the Coach, including the proper label for the concept. For instance, the system may respond to a student having just specified the law of demand (i.e., As price increases, quantity demanded decreases), with: "*Congratulations! You have just discovered what economists refer to as the Law of Demand. Please note that the converse is also true. As price decreases, quantity demanded increases.* Can you conduct an experiment that illustrates the Law of Demand?" The request for an experiment guarantees that the student possesses not only declarative knowledge about the concept in question, but also procedural knowledge about how to construct an exemplary experiment. When the student has successfully stated the concept in the hypothesis menu and has generated an experiment typifying this regularity or law, the concept is checked off as "learned" in the Goal History Window.

In summary, the Coach attends to three preliminary principles for interrupting a student's activity:

1. The student's particular problematic behavior must be apparent. An "apparent" level of ineffective behavior was initially an arbitrary threshold ratio, later refined after experimentation.
2. Feedback should not be always critical, but also appreciative of any good aspect of a student's behavior.
3. Any instruction or guidance should be spaced. At least a gap of two events should occur between interruptions.

INDIVIDUAL DIFFERENCES: INTERROGATIVE STRATEGIES

Over the course of system development, we collected protocol data on effective and less effective interrogative strategies from individuals interacting with *Smithtown*. The pilot data we collected have been used in building the Diagnostician, discussed earlier. This information focuses the Diagnostician on particular issues in the developing student model. We anticipated that the better students would attempt to derive and test generalizations about their experiences whereas the less successful students would work longer at the more concrete level of manipulating variables, changing parameters, and observing specific effects. In addition to knowledge-level differences, we also expected that the histories of students' interrogations would give evidence of differences in self-regulatory behavior. We anticipated that successful learners would show more activities summarizing and organizing information, more frequent testing of their current knowledge, and systematic predictions to check the limitations of the hypothesis being considered.

There were, as expected, large individual differences in the strategies used to interrogate a new domain where the purpose was to extract the underlying principles and systematicities. None of the subjects had any formal economics training and all were volunteers.

Protocol Analysis: Effective Interrogative Behavior

To illustrate a contrast in interrogative behaviors, two extractions from protocols follow. The first is an effective methodology and the second is a less effective one. To maintain subject confidentiality, we use the generic pronoun "he" to describe our two subjects.

The first subject clearly shows how someone can optimally extract information from an unfamiliar domain. This individual had no formal economics training. He spent one and a half hours interacting with the system and verbalized his actions, the justifications for them, current hypotheses of variable relationships, and predictions of different events. He completed a total of six distinct experiments within the system and showed a marked learning of the domain knowledge from the beginning to the end of the session.

Our subject used the same good, ice cream, for the six experiments and changed only one variable per trial (e.g., increasing overall population from 10,000 to 20,000). This subject was aware of his systematicity. When he was asked about his strategy, he replied,

> I'm trying to hold everything the same, then change one variable at a time, one of the ones I've identified as being independent for some reason, and keep running the simulation through the time units until there's 0 shortage and 0 surplus. The market would be stable.

Initially he used the option of letting the computer adjust the market price of a selected good. Our subject only began adjusting the price himself in the 4th of the total six experiments that he conducted. This was in accord with our belief that a more efficient interrogator would initially observe a phenomenon before becoming an active experimenter within the system. For instance, he observed how the computer made changes in the market price and the resulting changes that occurred to such variables as quantity demanded, quantity supplied, and so on. This served as a model for future explorations by the subject.

Our subject also demonstrated increasing selectivity in the variables he chose to record. Initially, he recorded all of the variables in the notebook.

> So, now I'm sort of regretting having used all these fields since I'm only experimenting with the suppliers. So I'm going to throw out the rest, other than the basic variables: quantity demanded, quantity supplied, surplus, shortage, market price and suppliers. I don't want to fill out the rest since I've decided that these are more critical, like they are independent variables and I'm only looking for variation within suppliers.

As his domain knowledge increased, this was clearly reflected in the content of hypotheses generated at the beginning of each experiment. For instance, after checking out and recording the market sales information for ice cream, he wanted to change the population of the town to see the effect on quantity demanded, quantity supplied, and other variables. When asked for specific hypotheses, he had none. In experiment #3, when he increased the number of suppliers, his hypothesis was a little more substantive,

> If there are more suppliers, the going price will be lower than it was because competition will be more intense.

Contrast that with a hypothesis generated during the final experiment. He doubled the population of Smithtown, had previously cut everyone's income by half, and had 10 ice cream shops in town. He hypothesized,

> Since the population is up, and even at their newly poor state, quantity demanded ought to go up. In fact, it ought to double. So, suppliers can raise their prices since there are going to be shortages. The new equilibrium price is going to be higher, and the route to this end is that there's going to be big shortages, and the price will move up until it reaches equilibrium.

Another strategy that changed was the *type* of experiments he created. His experiments became increasingly more controlled and precise. He began his explorations with a vague experiment, saying, "I'd sort of like to see what's gonna happen, like if I changed the population to see what happens to quantity demanded, and stuff." By contrast, at the end of his interaction with *Smithtown* he stated,

> You know what I'd like to do now? A whole series of experiments using 10, 11, 12, 13, and 14 suppliers, and look at what the market price is for each one of those and actually get curves . . . jump up a level of detail and look just at the equilibrium price and ignore the other variables.

The effective behavior of questioning the meaning of variable relationships is illustrated when our subject reflected on the repercussions of his cutting the average per capita income by half.

> So, what's the basic mechanism? Why is it that when income goes down, and people have less money to spend . . . they're less likely to buy ice cream? So they don't have money to blow on frivolties like ice cream. Hmmm, they're just gonna stay away from ice cream stores in droves. So, in order to encourage people to go to their stores, the shopkeepers lower their prices. As the price goes down, the people are more likely to go into the stores. If the shopkeepers drop their price way too far, people will clean them out and the shopkeepers will think, "We can do better than that", so they will up the price to readjust itself. Hey, they makes sense!

Finally, our subject's initial experimentations dealt with factors he was familiar with. When asked why he chose the number of suppliers as one of the first variables he manipulated, he responded that it was something he knew about and that he also knew, theoretically, the effects it would have in the marketplace.

Protocol Analysis: Less Effective Interrogative Behavior

The second subject illustrates how someone can extract minimal domain knowledge with ineffective inquiry skills. This individual also had no formal economics training and spent 1½ hours interacting with *Smithtown*, verbalizing all actions, predictions, and hypotheses. He completed a total of three experiments and gave evidence of learning only one elementary principle of supply and demand.[5]

For experiment #1, the market selected was compact cars, and the variable manipulated was income, which was increased from $8,000.00 to $10,000.00 in order to make it more "reasonable." This strategy of selecting values for variables that were more "normal" persisted throughout the entire session. There was no attempt to conduct scientific investigations of causal relationships. Instead, all attempts and changes centered around "making the town more normal." At the end of experiment #2, when the subject had decreased the price of compact cars and was asked why, he responded, "Because I wanted people to be able to afford to buy a car, really. I was putting myself in their place." In addition, the changes that *were* made were often inadequate. For instance, in one experiment, the price of chicken was changed by the subject from $1.33/lb. to $1.29/lb. This action reduced the price by only 4 cents to make it conform to what the subject believed was a "reasonable" and observed sales price. The action was not carried out as an experiment on the potential affects of changing price on the other variables.

Notebook entries continued to be incomplete, and contained unnecessary information. In his final experiment, the subject selected the market for chickens and made no changes to any variables. After seeing "Sales Market Information," he made a notebook entry of the variables: market price, quantity supplied, surplus, suppliers and consumer preference. The other essential variables, good, quantity demanded, and shortage, were missing. The inclusion of suppliers and consumer preference was unnecessary, as they had nothing to do with the subsequent manipulation, which dealt with price changes in the market.

Hypotheses regarding relationships between variables were based on surface structure characteristics of the simulations. When the subject was asked at the end of experiment #1 (market for compact cars, income changed) about any hypotheses for resulting variable relations, he responded,

> It just seems to me that there doesn't really seem much to do in this town for these people to need cars, you know what I mean? Not a whole lot, except to get back

[5]The law of demand was understood, illustrated by the statement: *"If I decreased the market price to something like $6,500.00, then more people are going to want to come out and buy the cars."*

> and forth to work. Most of them just can't afford a car. You make $10,000.00 and you're gonna take $7,000.00 and buy a car?!

Similarly, the subject understood causal relations among variables from an ethical perspective; if suppliers reduced the price of chickens, for instance, then they "should have" supplied a sufficient quantity of chickens for all who wanted them. The basic mechanism underlying the supply and demand model eluded the subject, as seen in the following statement, when the subject was astounded to find out that when the price of chickens was reduced, the amount of chickens supplied actually *decreased.*

> Why did they go down on the supply then? I mean, if they [the suppliers] knew it was going to go down, that they were going to have this sale for $1.29/lb., they should have . . . [put more chickens out]. Why did they do that? The shortage of chickens equals 65 pounds, I want to know why they did that. Now, wouldn't *you* provide more chickens?

When the price of chickens was increased from $1.29/lb. to $1.43/lb., the response was again incredulous,

> They gave *more* chickens! When they knew it was going to go up [the price], and they knew people wouldn't buy it at that rate . . . that *really* doesn't make sense! If this is really true, somebody should put a stop to it!!

Finally, when asked at the end of the experiment if he had applied any strategies to the explorations, he responded, "No, not really. I was just going by what I know from day to day experience."

The contrast between these two interrogative styles is clear. In terms of the inquiry behaviors listed earlier, our more effective subject demonstrated a systematicity in the varying of variables, a concern with making changes to variables sufficiently large, a selective recording of notebook entries, a progression from 'computer' to 'self' price adjustments,[6] as well as a tendency to start with more familiar knowledge, and increase the complexity of his hypotheses as his domain knowledge increased. The relationship between his behaviors and his growing comprehension of the domain knowledge was also apparent. The less effective subject's behavior showed an absence of most of these skills. In addition, he attended more to the surface structure characteristics of events at the expense of the more principled relationships describing the laws of supply and demand and their interaction.

---

[6]Alternatively, this strategy could be characterized as progressing from *observing* the microworld to more actively *experimenting* within the microworld.

EVALUATION

In order to determine what the experience with the microworld actually taught, we investigated student knowledge and skill acquisition in a controlled design (see Shute, Glaser, & Raghavan, 1989 for a complete discussion of these results). Pretest and posttest batteries covering economic principles dealing with the laws of supply and demand were developed in conjunction with an economics expert working on the project. There were three tests per battery: alternative/matched forms of a multiple choice test, a short answer test, and complex scenarios test requiring subjects to solve "What would happen if . . ." type of problems. The complex scenarios test was designed to externalize an individual's problem representation by having the subject identify the relevant variables in the problem, chart which variables were related, and indicate the nature of the relationships.

These test batteries were administered to 30 subjects divided equally into three groups: a control group receiving no intervention, a group from an introductory microeconomics classroom, and a group interacting with *Smithtown*. Relative performance on these tests showed that *Smithtown* was effective in imparting domain knowledge. Subjects in the economics classroom and those working only with *Smithtown* demonstrated very similar gain scores (from pretests to posttests) even though the classroom group spent more than twice as much time on the same curriculum material than the *Smithtown* group. Moreover, *Smithtown* did not directly instruct economics; rather, the instruction was specific for inquiry skills.

Two data sources were used in the analysis of experimental subjects interacting with the microworld. First, there were verbal protocols consisting of the verbalizations from the subject. Second, there were the computer-recorded student histories of all actions in the microworld. A rational categorization of behaviors was made based on the verbal protocols, dividing the learning behaviors into three global categories: (a) gross activities, (b) data management behaviors, and (c) scientific actions. These were further subdivided into more refined categories, each containing individual learning indicators or critics (see Shute & Glaser, 1990, for more information on the validity of these categories).

A large-scale study (N = 530) was conducted using *Smithtown* to investigate individual differences in learning behaviors within this type of environment (see Shute & Glaser, 1990). This study showed that a factor indicating the degree to which a person engaged in hypothesis-generating and testing was the most predictive of successful learning in *Smithtown*, accounting for considerably more of the variance in our learning criterion (i.e., total number of concepts acquired) than a measure of general intelligence. The individual learning indicators comprising this factor consisted of: total number of hypotheses made, number of times findings from one experiment were generalized across related and unrelated markets, and having sufficient data collected prior to making a hypothesis.

SUMMARY AND FUTURE DIRECTIONS

The economics microworld discussed in this chapter was designed as an instructional system using a laboratory environment for studying inquiry and inferencing behaviors. If a student is not progressing in that environment, there is a mechanism to provide tutorial assistance. Such discovery microworlds provide excellent opportunities for observing and reacting to individual differences in learning styles, creating a more active and adaptive curriculum than is currently possible in many traditional classroom or computer settings. The system can react to individual differences by exemplifying effective interrogative strategies that can be emulated by individuals, or by providing explicit instruction on interrogation strategies to uncover consistencies and regularities in a body of knowledge.

In summary, *Smithtown* is based on a set of general instructional design principles:

- The instructional situation provides the learner with opportunities to interrogate the microworld and to formulate and test hypotheses through questions and simulated experiments.
- The instructional environment permits actions at different levels. The student is able to conduct experiments, record the results, and select and analyze subsets of that database. For example, the student can tabulate variables and see the results, make a statement of a hypothesized relationship, and obtain further information about its adequacy or limitations.
- The computer screen serves as a memory aid for students, allowing them to see the results of explorations to date. The information is displayed in various forms (e.g., graphical or tabular) or in terms of verbal statements of relationships.
- There is hierarchy of domain knowledge available in the microworld, and the system has the information about what subset of it has been learned by the student and what subset of it is still available for learning. This information is accessible to the student.
- A history of the student's actions, in selecting information and formulating and testing hypotheses, are sequenced and stored in a way that permits student performance to be displayed and evaluated in a structured manner.
- As the student explores *Smithtown,* the Diagnostician accesses prerequisite skills necessary for activity in the microworld, knowledge components that the student has discovered and acquired, and strategic issues like the moves of the student with respect to interrogation and hypothesis formation.
- The Diagnostician passes a list of problem behaviors on the Student Model which uses this information, when necessary, to recommend situations to the

Coach for possible student assistance. This guidance should enable the student to demonstrate more effective interrogative procedures.

Because an individual's learning style (e.g., passive vs. active) is believed to interact with the instructional environment (e.g., discovery vs. didactic), as well as the domain (i.e., microeconomics), this raises several questions regarding *Smithtown* that have not yet been tested: (1) How much flexibility of behaviors should be allowed before the Coach intervenes? Because there is not just one way to use the tools optimally to conduct experiments to induce regularities, what should the criteria be for saying that a given student is floundering or behaving unsystematically? In a primarily discovery environment such as *Smithtown*, how do we disambiguate the passive learners who are not active because of cognitive deficits rather than because of an acquired learning style? (2) Because the Diagnostician monitors students' experiments in terms of the systematicity of their actions, how should it identify those interrogative skills that have not been demonstrated? Some effective interrogative behaviors are not demonstrated because they are not understood, or they have yet to be applied but are understood, or they are simply not preferred (e.g., the student chose not to graph data). (3) What are the trends of the learning indicators over time? How does increasing domain knowledge relate to inquiry behaviors? For example, as an individual learns more about economics, can he or she manipulate progressively more independent variables simultaneously? (4) How generalizable are these findings regarding the delineated "effective" and "less effective" inquiry skills? Does training on these skills transfer to successful performance in a different domain? These issues should be sorted out as more data accumulate.

To test the effectiveness of the acquisition of inquiry skills, we have planned a series of transfer studies where students interacting with one system will subsequently learn a new body of knowledge from another inquiry system. There are several microworlds now developed that are similar in system architecture to *Smithtown*, but differ in the domain knowledge contained: (a) light refraction and reflection, (b) orbital mechanics, and (c) basic principles of electricity.

In conclusion, we have been obtaining rich information on how students collect and organize information, how they use evidence to generate a test hypotheses, and form generalizations in the course of scientific inquiry with computer-based laboratories.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the many persons whose contributions are invaluable to this project: Jeff Blais, Jeff Bonar, Alan Lesgold, Kalyani Raghavan, Peter Reimann, and Paul Resnick. Last but hardly least, we offer a special thanks to Jamie Schultz for his contribution of ideas and code, making this project an online reality.

The research reported herein was supported by the Learning Research and Development Center, supported in part as a research center by funds from the Office of Educational Research and Improvement (OERI) of the U.S. Department of Education. The opinions expressed do not necessarily reflect the position or policy of OERI and no official endorsement should be inferred.

REFERENCES

Anderson, J. R. (1983). *The architecture of cognition.* Cambridge, MA: Harvard University Press.

Anderson, J. R., & Bower, G. H. (1973). *Human associative memory.* Washington: Winston and Sons.

Carr, B., & Goldstein, I. P. (1977). Overlays: A theory of model for computer-aided instruction (*AI Memo 406*), MIT A.I. Lab.

Collins, A. M., & Loftus, E. F. (1975). A spreading activation theory of semantic processing. *Psychological Review, 82,* 407–428.

Collins, A. M., & Quillian, M. R. (1969). Retrieval time from semantic memory. *Journal of Verbal Learning and Verbal Behavior, 8,* 240–247.

Cronbach, L. J. (1966). The logic of experiments on discovery. In L. S. Shulman & E. R. Keisler (Eds.), *Learning by discovery.* Chicago: Rand McNally.

Glaser, R. (1984). Education and thinking: The role of knowledge. *American Psychologist, 39,* 93–104.

Norman, D. A., & Rumelhart, D. E. (1975). *Explorations in cognition.* San Francisco: Freeman.

Pitt, R. B. (1983). Development of a general problem-solving schema in adolescence and early adulthood. *Journal of Experimental Psychology: General, 112*(4), 547–584.

Shrager, J. C. (1985). *Instructionless learning: Discovery of the mental model of a complex device.* Unpublished doctoral dissertation, Carnegie-Mellon University.

Shute, V. J., Glaser, R., & Raghavan, K. (1989). Inference and discovery in an exploratory laboratory. In P. L. Ackerman, R. J. Sternberg, & R. Glaser (Eds.), *Learning and individual differences.* New York: W. H. Freeman.

Shute, V. J., & Glaser, R. (1990). A large-scale evaluation of an intelligent discovery world: *Smithtown. Interactive Learning Environments,* Vol 1.

Sleeman, D., & Brown, J. S. (1982). *Intelligent tutoring systems.* New York: Academic Press.

Wittrock, M. C. (1966). The learning by discovery hypothesis. In L. S. Shulman & E. R. Keisler (Eds.), *Learning by discovery.* Chicago: Rand McNally.

APPENDIX I

Smithtown: The Content, Concepts and Model

This appendix overviews the content and concepts that are included in the system, representing the central economic principles that we want students to learn. We discuss how these concepts are implemented in our system, detailing the underlying model which drives the simulation. The principles, thus, will be discussed as they exist and function in *Smithtown.*

Microeconomic Principles

One of the central concepts in economics is the concept of a market. In the economic sense, a market is the interaction between buyers and sellers. Their interaction determines the quantities of various goods and services (products and resources) which will be bought and sold during some period of time, as well as the prices at which these exchanges take place. Thus, through markets, we make most of our basic decisions about how resources in our society will be used: What will be produced, how will it be produced, and who will get it once it is produced. In a market, buyers and sellers independently consider the price of the product in determining the quantities they want to buy and sell; however, price affects buyers and sellers differently.

Supply and Demand. The buyers' side of the market is called *demand*. The law of demand states that the quantity of a product which consumers would be willing and able to purchase during some period of time is negatively related to the price of the product. If the price of gasoline goes up, we will demand a smaller quantity of gasoline; if the price goes down, we will demand a larger quantity. The same is true for most products. If we draw a graph of the combinations of price and the resulting quantities demanded, we get what is called a *demand curve*.

The sellers' side of the market is called *supply*. The law of supply is that the quantity of a product which producers would be willing and able to produce and sell during some period of time is positively related to the price of the product. If the price of color television sets goes up, producers will be willing and able to offer more television sets for sale. If the price of color television sets goes down, producers will reduce the number of television sets they put on the market. If we draw a graph of the combinations of price and the resulting quantities supplied, we get what is called a *supply curve*.

The Equilibrium Point. Both buyers and sellers consider the market price in making decisions, but price affects buyers and sellers differently. If the price rises, sellers will make larger quantities available for sale, but buyers will demand smaller quantities. If the price falls, buyers will demand larger quantities, but sellers will make smaller quantities available for sale. It is because of these opposite reactions to price changes that buyers and sellers can reach an agreement. When a price is reached where the quantity that sellers want to sell is equal to the quantity that buyers want to buy, we say that the market is at a point of *equilibrium*.

Competitive markets always tend toward points of equilibrium. If the market is higher than the equilibrium price, buyers will demand smaller quantities than sellers are supplying. *Surpluses* of unsold goods will convince sellers to lower their price down toward the equilibrium level. If, for some reason, the market price is lower than the equilibrium price, buyers will demand larger quantities than sellers are supplying. *Shortages* will lead to price increases, and the price will rise toward the equilibrium level.

Changes in Demand and Supply. The equilibrium point in the market can change if either demand or supply change. Price determines the quantity demanded. Other factors also play a role in determining demand. Some of these other determinants are: 1) consumers' income—normally, if income increases, demand will increase; 2) consumers' tastes—if tastes shift in favor of a product, the demand will increase; 3) the price of substitute products—the demand for a product will increase if the price of a substitute increases; and 4) the price of complementary products—the demand for a product will decrease if the price of a complement increases. Other factors such as interest rates, the weather, population, and expectations about the future can also affect the demand for some products.

Graphically, a change in one of these other determinants results in a shift in the demand curve. This is called a *change in demand.* A shift to the right represents an increase in demand where a larger quantity will be demanded at each price. A shift to the left represents a decrease in demand where a smaller quantity will be demanded at each price.

In addition to price, other factors play a role in determining supply, including: 1) the cost of resources—if the cost of resources rises, the supply will decrease; 2) technology—if technology improves, supply will increase; 3) profits available in other lines of production—the supply of a product will decrease if other lines of production available to the seller become more profitable; and 4) the number of sellers—if the number of sellers increases, supply will increase.

Graphically, a change in one of these other determinants results in a shift to the supply curve. This is called a *change in supply.* A shift to the right represents an increase in supply—a larger quantity will be supplied at each price. A shift to the left represents a decrease in supply where smaller quantity will be supplied at each price.

Establishing a New Equilibrium. Competitive markets tend to converge toward equilibrium points. Equilibrium, once established, can be disturbed by changes in demand and/or supply. If demand and/or supply change, a surplus or shortage will result at the original price, and price will move toward a new equilibrium. A shortage at the original price will cause price to rise to the new level and cause changes in the quantities supplied and demanded. A new equilibrium will be established at the second price and the second quantity.

Economic Misconceptions. There are certain beliefs about the market that students typically have as they begin an introductory economics course. Three fairly ubiquitous misconceptions are outlined below.

1) Students often have difficulty distinguishing between a change in the quantity demanded (i.e., a movement along the curve) and a change in demand (i.e., a shift of the curve). They often believe that the curves shift in response to price changes, but in reality, curves shift as a function of factors *other* than price, like population changes or interest rate changes.

2) In order to understand how markets work, supply and demand must be seen as independent of each other—linked only by a common determinant (i.e., price). Students often think that changes in demand cause changes in supply instead of just changes in the quantity supplied, which is the true state.

3) Students often think of the demand curve as having a positive slope: "The price of a product will be high if people are demanding a large quantity of it." They also see the supply curve as having a negative slope: "Sellers lower the price when they want to supply a large quantity." These errors result from a misunderstanding of the demand and supply functions. It is price which determines the quantities demanded and supplied and not vice versa.

Student Actions in *Smithtown*

On their initial encounter with the system, students are instructed to experiment within the microworld, make changes to different variables, and see the results. Depending on the efficiency of their inquiry behaviors (e.g., systematicity of experiments), students can extract differing amounts of information regarding how supply and demand interact in a competitive market. For example, students starting out with no previous economic knowledge may design the following simple experiment: select a familiar good (e.g., donuts), not make any changes to the global variables (income, population, etc.), then collect information about the markets as it stands (baseline data).

To learn an elementary concept like the law of demand (i.e., the inverse relationship between price and quantity demanded), students can alter the price of donuts and see sales market information. Sales market information includes: quantity demanded, quantity supplied, surplus, and shortage. After collecting several instances of price changes and affects, the student can use some of the available tools included in the program to sort and order the data. One observation which should become apparent, then, is that as price goes up, quantity demanded goes down.

As discussed earlier, for both demand and supply, we have included in the model the factors that shift the respective curves. For each good represented in the system, coefficients are assigned to these variables indicating strength or relative importance. For instance, interest rates will have a larger coefficient for large and compact car markets, yet only minimal or no influence on the market for donuts. On the other hand, weather conditions will have a significant impact on the demand curve for ice cream, but less influence on the market for chicken.

Following are the variables the student may manipulate, affecting the demand and supply curves.

Demand Shifters

Population: For most goods, as population increases, the demand curve shifts to the right, the magnitude of shift determined by the strength of the coefficient.

Income: As income increases, the demand curve generally shifts to the right. Again, each good has its own coefficient indicating how far to shift the curve, the amount of shift dependent on the particular good. Some goods or services may have 0 as coefficients (e.g., water) or even negative coefficients (e.g., inferior goods).

Weather Conditions: For some goods, better weather may cause the demand curve to shift, either left or right. We have represented differential weather conditions in the system on a scale of 1 (cloudy, cold day) to 10 (beautiful, sunny day).

Interest Rates: As interest rates increase for some goods, the demand curve shifts to the left (e.g., cars). Again, the degree of shift is dependent on the underlying coefficient per market.

Consumer Preferences: A number, from 1 (low) to 10 (high) may be assigned to represent the consumer taste or preference for that good. This variable is unrelated to price, and can be affected by things like advertising, word of mouth, and so on.

Price of Substitute Goods: As this price goes up, there is a shift to the right of the good in question since more people will demand the alternative good (e.g., the price of butter increases, resulting in a shift to the right for the demand curve for margarine). Coefficients for this variable are positive.

Price of Complementary Goods: Complementary goods, those associated with the selected good, have negative coefficients, therefore as the price of complementary goods go up, the demand curve for the current good shifts left since they are typically purchased in conjunction with each other.

Preference Changes to Substitute and Complementary Goods: In addition to reactions through price, the demand for certain goods will react to tastes/preferences (scaled 1 to 10) for certain other goods.

Supply Shifters

Number of Suppliers: A change in the number of suppliers of a good will shift the supply curve for that good. Each good has a coefficient indicating how far to shift the supply curve. The coefficient is positive where, more suppliers of a good results in more of the good being supplied (i.e., a shift to the right of the supply curve). The converse is also true.

Cost of Resources: As this variable increases, there is an inverse effect on the supply curve for the good that uses the resource, thus it is a negative coefficient. For instance, if the cost of sugar increased, this would shift the supply curves for donuts and ice cream to the left.

Technology: If there are technological advances impacting a particular market, then this will have a positive effect on the supply curve for that market. To illustrate, suppose a new technology was developed for increasing the production of compact cars (robots in the factory). The supply curve for small cars would subsequently shift right reflecting the increased output.

Labor/Wage Costs: If the labor cost increases, this will have an inverse effect on the supply curve for the good that requires this service, so it has a negative coefficient.

Computer Calculations and Representations

Currently, we have a total of 13 goods in the system for students to manipulate. These are: coffee, tea, Cremora, donuts, ice cream, compact cars, large cars, gasoline, chicken, ground beef, hamburger buns, lumber, and wooden bookcases. For each one of these there is a list of variables containing default values: that is, the variables set in *Smithtown* when the student begins his or her investigations. Some of the more important ones include: equilibrium price, equilibrium quantity, intercept for supply curve, slope of supply curve, intercept for demand curve, slope of demand curve, list of substitute goods, list of complementary goods, current price, time frame,[7] as well as the coefficients for each of the shifter variables listed above.

Once a student has selected a good from the menu, the default values are initialized and a new menu appears. This menu contains a list of variables the student may choose to alter current default values (e.g., population is initially set at 10,000 Smithtonians). If an item from this menu is selected and changed, the action causes the demand or supply curve to shift from its default situation. Since both the intercept and the slope values are stored per good, and each good knows the relative influence of a particular variable on it, then the curve shifts by an amount indicated by the coefficients. To illustrate, suppose a student was investigating the effects of a changing population on the market for donuts, increasing the population from 10,000 to 20,000 residents. The demand curve will shift since population it is known to be a demand shifter, and the magnitude of the shift will be:

$$(0.1)(10{,}000) = 1000 \text{ scaled units}$$

where the 0.1 is the coefficient attached to population for donuts, and 10,000 it the amount of change (i.e., 20,000 − 10,000 = 10,000). The coefficients vary

[7]This represents the time period during which the market is monitored.

according to their relative influence on the shifters in a particular market.[8] The slope remains constant, so the shift is a uniform 1000 units to the right. Now, an updated equilibrium price can be calculated from the new point of intersection between the shifted demand curve and the unchanged supply curve.

In the demand and supply curves, the intercept and slope values for each good have been initially set at reasonable market values. For example, the equilibrium price for donuts is set at $0.50 per donut (i.e., the price at which quantity demanded equals quantity supplied), the quantity demanded intercept is set at 100 (i.e., the maximum number of donuts demanded in a given period of time such as one week), and the demand slope is set at −2, thus the initial demand curve is fully specified. Similarly, the supply curve for donuts has the same equilibrium price of $0.50, a supply intercept of −3, and a supply slope of 2.

[8]That is, the size of the shift is determined by the size of the coefficient. The coefficients for each good were based on observations about how much the different variables affected the different markets. For instance, interest rates strongly affect the market for large cars, but do not affect the market for donuts, thus, the interest rate coefficient for large cars is larger than for donuts.