Southern Economic Journal 2000, 66(3), 770–792

Targeting Teaching

Exploring Economic Models Using Excel

Miles Cahill* and George Kosicki†

This paper applies spreadsheet software to intermediate-level consumer theory concepts. Spreadsheets help make the concepts more accessible while allowing students to explore the ideas in more depth. Areas of application are utility functions, income and substitution effects, price indices, measures of welfare change, and the optimal saving rate. We chose the examples to stimulate awareness and discussion of the many classroom uses for four important Excel spreadsheet tools: three-dimensional (3-D) graphs, iteration, Goal Seek, and Solver.

1. Introduction

This paper uses spreadsheet software to explore a series of consumer theory examples appropriate to intermediate-level economics classes. The paper also illustrates how spreadsheets can make some advanced topics more accessible to students, thereby helping to bridge the gap between undergraduate and graduate education in economics.[1] The examples pertain to important topic areas in consumer theory: utility functions, substitution and income effects, price indices, measures of welfare change, and the optimal aggregate saving rate.

A web-based supplement accompanies this paper and includes additional examples pertaining to cost minimization, oligopoly models, IS-LM and aggregate supply-aggregate demand models, and the Solow growth model. This supplement is available on the Internet at http://sterling.holycross.edu/departments/economics/mcahill/sejpaper.html. The examples in the supplement draw on the same Excel tools that the paper presents.[2] The web supplement also contains student instruction

* Department of Economics, P.O. Box 99A, College of the Holy Cross, Worcester, MA 01610, USA; E-mail mcahill@holycross.edu; corresponding author.

† Department of Economics, P.O. Box 162A, College of the Holy Cross, Worcester, MA 01610, USA; E-mail gkosicki@holycross.edu.

We are grateful for useful comments from Jonathan Hamilton, John Carter, Mike Lovell, Mike Watts, and an anonymous referee.

[1] See the report by Kasper et al. (1991) on the preparation of economics students for graduate study. Three of the questions asked of representatives of highly selective liberal arts colleges were: "Is there much of a gap between the economic theory taught to your best majors, and what is taught during the first year of graduate work?" "Is that gap widening?" "Have you tried to reduce that gap?" All nine of the representatives indicated that there was a gap, six thought the gap was widening, and six indicated that their departments had tried to reduce the gap. One of the recommendations made in the report was that "the transition from undergraduate to graduate study would be eased if the undergraduate curricula offered advanced economic theory as regularly as it did 20 years ago" (p. 1106). For a critique of this recommendation, see Bateman (1992).

[2] Some of the Excel features that are utilized in this paper are relatively new to spreadsheets and have not been explored in earlier surveys such as Yohe (1989) and Judge (1990b). Judge (1990a), Welford (1990), Porter and Riley (1995), and Whigham (1998) contain collections of detailed spreadsheet assignments for economics. See Benninga (1997) for a number of financial modeling applications using Excel. Some of the examples in Porter and Riley (1995), Benninga (1997), and Whigham (1998) make use of the Excel tools highlighted in this paper.

sheets that give step-by-step instructions on completing the examples in this paper. In addition, annotated versions of the spreadsheets referred to in the paper are available at this address. It is possible to complete most of the applications discussed in the paper and the web supplement on all three of the leading spreadsheet packages (Microsoft Excel, Lotus 1-2-3, and Corel Quattro Pro), but only the Excel 97 commands are used.[3]

From a practical perspective, spreadsheet software such as Excel is a natural choice to use in exploring economic models because it is widely available on most campuses. This availability eliminates the task of seeking funding for the purchase and support of specialized software packages. In addition, spreadsheet software is relatively easy to use, and its flexibility makes it useful in many different courses at all levels of the traditional economics curriculum.[4] Most economics students almost certainly will use it after graduation in both career and personal settings. Most important, it minimizes black-box features that characterize much computer-assisted learning software.

Writing spreadsheet applications does take considerable time and effort. While it is often easy to show someone how to perform a certain operation, writing the sequence of commands can be difficult.[5] Finding and remedying the quirks of certain spreadsheet programs can be likewise burdensome. Yet, we believe these start-up costs are worth enduring because spreadsheets can do more than change the look of teaching, they can change the substance as well, in often surprising ways. To ease these start-up costs, the web supplement provides step-by-step student instructions for the applications. However, we provide these handouts with two notes of caution. First, reliance on these instructions may encourage students to follow the steps blindly without understanding the material. Second, the level of detail in the instructions may make the Excel commands seem more complicated and tedious than they really are.

The applications in this paper can complement many different approaches to teaching, from traditional lecture formats to more active-learning formats including discovery-oriented lab sessions.[6] It is possible to tailor many of the applications to suit the level, ability, and time constraints of students. Specifically, the instructor may choose to have students build the spreadsheets from scratch based on printed instructions, instruct students on how to build the spreadsheet in a lecture or lab setting, distribute completed or half-completed spreadsheets to students, or any combination of the above. Further, the instructor may choose to protect certain spreadsheet cells to prevent students from manipulating them.[7]

---

[3] The 1997 version of Lotus 1-2-3 cannot produce the three-dimensional (3-D) graphs discussed in section 2.

[4] While not focusing on spreadsheet software in particular, Blecha (1991) points out ways in which using computers in the classroom can add continuity to the students' learning experiences. Her basic premise is that "It is in the area of curricular integration that any revolutionary seeds sown by microcomputers are likely to take root" (p. 546). She presents examples to show how software can increase the degree of integration between topics in a course, between courses in a major, and between courses in different departments. Breece (1988) provides a thoughtful review of the different types of software available for use in the classroom. In the section concerning spreadsheets, he writes, "This technical innovation appears to have many applications in the teaching of economics; however, for some reason such programs are grossly underutilized" (p. 500).

[5] If this is all too apparent in the text below, our only suggestion is to try out the applications on a computer.

[6] See Bartlett and King (1990) and King and LaRoe (1991) for a description of a lab-based approach to teaching economics. The textbook by Yohe (1995) also takes a discovery-oriented approach to teaching. Judge (1996) evaluates a spreadsheet assignment in which students actively build their own economic models. We believe the spreadsheet applications in this paper could be used to facilitate the type of active-learning environment cited by Becker (1997) as one of the strategies for successful undergraduate teaching. Chizmar and Walbert (1999) also discuss using Excel spreadsheets in intermediate micro to facilitate a more active-learning environment.

[7] To protect all cells in a worksheet, choose Tools/Protection/Protect Sheet. To unprotect individual cells, unprotect the worksheet (Tools/Protection/Unprotect Sheet), highlight the cells you do not wish to protect, select Format/Cells, then deselect the Locked box under the Protection tab. Then when the spreadsheet is protected, the unlocked cells are not.

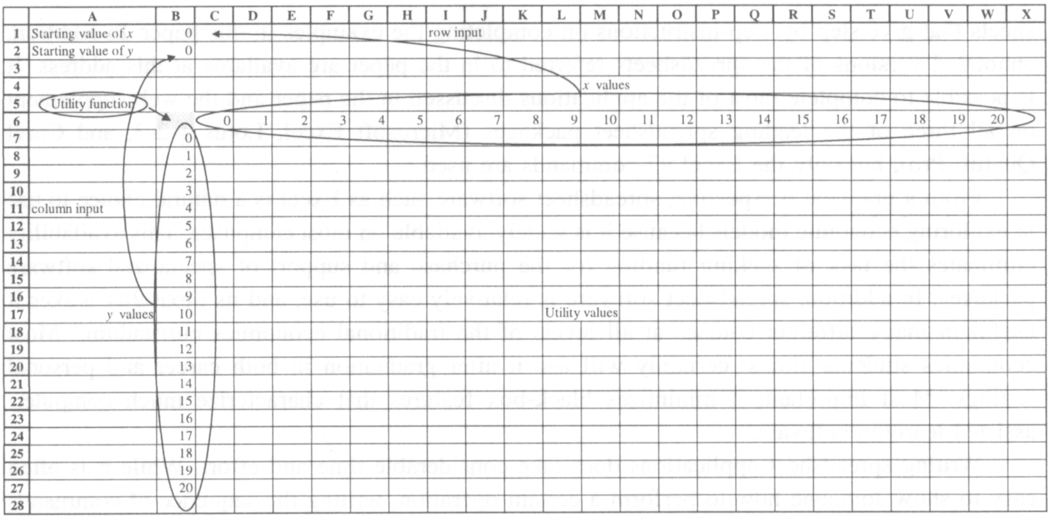

Figure 1. Tabular Setup of Utility Function Exercise

With regards to our own teaching, one of us has used spreadsheet applications primarily in a lab setting, and the other mainly in problem sets completed outside of class. We have found the use of spreadsheets enhances both teaching styles. When using either approach, the instructor is cautioned to carefully work out the exercises ahead of time, paying close attention to commands, options, and features that students may find confusing, perhaps distributing a handout to minimize confusion.[8] As students gain more experience, these problems naturally become fewer and fewer. Instructors can exploit this learning curve by, for example, giving detailed technical instructions for initial, simple assignments and giving fewer instructions for later more complicated exercises.

2. Utility Functions

A significant obstacle in microeconomics for many students is the connection between a utility function and the indifference map it generates. Excel can help to clarify these connections because of its ability to create three-dimensional (3-D) surface charts as well as the topographical or contour maps associated with 3-D functions.[9]

Excel's Data/Table command can generate the underlying data for the following exercise in plotting utility functions. Figure 1 shows the first step in the setup. To use the Data/Table command, you must form a table of cells in which the top row is the range of x values over which the utility function is to be evaluated, and the left-hand column represents the range of y values. The cells in the middle of the table will contain the utility assigned to each combination

[8] For example, Excel has an AutoComplete feature; if the first few characters you type in a cell match an existing entry in the same column, Excel fills in the remaining characters. This can be frustrating as the following example shows. Suppose you type in "x and y values" into cell A1 as a label. In cell A2, you intend to type in "x." However, Excel automatically places "x and y values" into cell A2 as soon as you type "x." To disable this feature, deselect the appropriate box under Tools/Options/Edit.

[9] Boyd (1998) explores the connection between a utility function and the indifference map using Maple, a symbolic computation program.

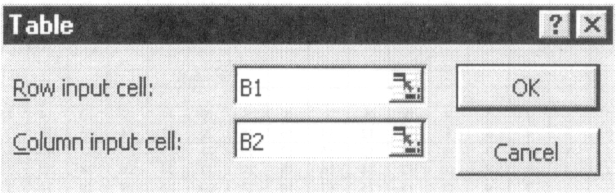

Figure 2. The Data/Table Command Window

of x and y. The spreadsheet in Figure 1 is set up to evaluate a utility function over the range of x and y values from 0 to 20 in increments of 1 unit.

The Data/Table command requires that somewhere on the worksheet there must be starting values for x and y. Using references to the cells containing the starting values, type the utility function into the corner of the table (as shown in Figure 1). For example, to generate data pertaining to a perfect substitute utility function $U = 2x + y$, use the formula =2*B1+B2. Then, the entire table of cells must be selected (cells B6 through W27). Implementing the Data/Table command spreads the utility function formula throughout the empty cells of the selected table by systematically replacing the starting value of x by each of the x values residing in the first row of the matrix. Similarly, Excel systematically replaces the starting y value with each of the y values residing in the first column of the table. The x values in this case are the "row input" and the y values are the "column input". To implement the Data/Table command, type into the Data/Table window the cell addresses associated with the starting values for the row and column inputs. Given the setup of Figure 1, Figure 2 shows the proper way to fill in the Data/Table window. Note that the response to the "Row input cell" prompt is cell B1 (the starting x value cell), while the proper response to the "Column input cell" prompt is cell B2 (the starting y value cell).

Implementing the Data/Table command generates the data set shown in the upper half of Figure 3. Unfortunately, Excel 97's Chart Wizard cannot plot this block of data unless the upper-left corner cell is deleted.[11] To make this deletion in a way that will not inhibit additional utility function simulations, copy the data and paste it to another location on the worksheet using the Edit/Paste Special/Paste Link sequence of commands. The Paste Link option appears in the Edit/Paste Special window shown in Figure 4. The Edit/Paste Special/Paste Link command sequence will create a new table in which each cell contains a formula that sets the cell equal to its counterpart cell in the original table. Delete the cell in the upper-left corner of this duplicate table (B31), and then select the entire table as shown in the lower half of Figure 3.

Once the data are selected, use the Surface Chart option of Excel's Chart Wizard (use the Insert/Chart command) to construct a plot as shown in Figure 5. Selecting the upper-left option in the Chart subtype matrix gives a multicolor, 3-D look at the utility function. Different colors

[10] In this paper, we denote a sequence of Excel commands by a sequence of words separated by slashes (/). For example, the reference "Data/Table" first directs the reader to choose the Data menu at the top of the screen, then the Table command within that menu.

[11] This was not the case in previous versions of Excel that contained the 3-D surface chart option. In those earlier versions, the corner cell could be part of the area selected provided that Excel was told to "Use First 1 Row for Category (X) Axis Label" and to "Use First 1 Column for the Series (Y) Axis Label." Excel 97 eliminated the dialog box containing these settings.

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W
1	Starting value of *x*	0																					
2	Starting value of *y*	0																					
3																							
4													*x* values										
5	$U = 2x + y$																						
6		0	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
7		0	0	2	4	6	8	10	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40
8		1	1	3	5	7	9	11	13	15	17	19	21	23	25	27	29	31	33	35	37	39	41
9		2	2	4	6	8	10	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40	42
10		3	3	5	7	9	11	13	15	17	19	21	23	25	27	29	31	33	35	37	39	41	43
11		4	4	6	8	10	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40	42	44
12		5	5	7	9	11	13	15	17	19	21	23	25	27	29	31	33	35	37	39	41	43	45
13		6	6	8	10	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40	42	44	46
14		7	7	9	11	13	15	17	19	21	23	25	27	29	31	33	35	37	39	41	43	45	47
15		8	8	10	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40	42	44	46	48
16		9	9	11	13	15	17	19	21	23	25	27	29	31	33	35	37	39	41	43	45	47	49
17	*y* values	10	10	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40	42	44	46	48	50
18		11	11	13	15	17	19	21	23	25	27	29	31	33	35	37	39	41	43	45	47	49	51
19		12	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40	42	44	46	48	50	52
20		13	13	15	17	19	21	23	25	27	29	31	33	35	37	39	41	43	45	47	49	51	53
21		14	14	16	18	20	22	24	26	28	30	32	34	36	38	40	42	44	46	48	50	52	54
22		15	15	17	19	21	23	25	27	29	31	33	35	37	39	41	43	45	47	49	51	53	55
23		16	16	18	20	22	24	26	28	30	32	34	36	38	40	42	44	46	48	50	52	54	56
24		17	17	19	21	23	25	27	29	31	33	35	37	39	41	43	45	47	49	51	53	55	57
25		18	18	20	22	24	26	28	30	32	34	36	38	40	42	44	46	48	50	52	54	56	58
26		19	19	21	23	25	27	29	31	33	35	37	39	41	43	45	47	49	51	53	55	57	59
27		20	20	22	24	26	28	30	32	34	36	38	40	42	44	46	48	50	52	54	56	58	60
28																							
29		Use Copy and Edit/Paste Special/Paste Link sequence; then delete contents of cell in upper left corner (B31)																					
30																							
31			0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
32		0	0	2	4	6	8	10	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40
33		1	1	3	5	7	9	11	13	15	17	19	21	23	25	27	29	31	33	35	37	39	41
34		2	2	4	6	8	10	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40	42
35		3	3	5	7	9	11	13	15	17	19	21	23	25	27	29	31	33	35	37	39	41	43
36		4	4	6	8	10	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40	42	44
37		5	5	7	9	11	13	15	17	19	21	23	25	27	29	31	33	35	37	39	41	43	45
38		6	6	8	10	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40	42	44	46
39		7	7	9	11	13	15	17	19	21	23	25	27	29	31	33	35	37	39	41	43	45	47
40		8	8	10	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40	42	44	46	48
41		9	9	11	13	15	17	19	21	23	25	27	29	31	33	35	37	39	41	43	45	47	49
42		10	10	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40	42	44	46	48	50
43		11	11	13	15	17	19	21	23	25	27	29	31	33	35	37	39	41	43	45	47	49	51
44		12	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40	42	44	46	48	50	52
45		13	13	15	17	19	21	23	25	27	29	31	33	35	37	39	41	43	45	47	49	51	53
46		14	14	16	18	20	22	24	26	28	30	32	34	36	38	40	42	44	46	48	50	52	54
47		15	15	17	19	21	23	25	27	29	31	33	35	37	39	41	43	45	47	49	51	53	55
48		16	16	18	20	22	24	26	28	30	32	34	36	38	40	42	44	46	48	50	52	54	56
49		17	17	19	21	23	25	27	29	31	33	35	37	39	41	43	45	47	49	51	53	55	57
50		18	18	20	22	24	26	28	30	32	34	36	38	40	42	44	46	48	50	52	54	56	58
51		19	19	21	23	25	27	29	31	33	35	37	39	41	43	45	47	49	51	53	55	57	59
52		20	20	22	24	26	28	30	32	34	36	38	40	42	44	46	48	50	52	54	56	58	60

Figure 3. Getting the Data Ready for the 3-D Surface Chart Option

or shadings denote combinations of *x* and *y* falling into certain utility level ranges.[12] To view the graph at different angles, click on and drag a corner of the graph or adjust the settings in the Chart/3-D View window. Figure 6 shows the 3-D surface chart associated with the data in Figure 3.

As the legend in Figure 6 illustrates, each shaded region represents a set of (*x, y*) coordinates that generate utility values falling into a certain range. For example, the lowest region shows all the combinations that are assigned a utility index between 0 and 10, the next region shows all the combinations of *x* and *y* that are assigned a utility index between 10 and 20, and so on. Therefore, the dividing line between these first two regions shows all the combinations of *x* and *y* that are assigned a ranking of 10.

Plotting the data in Figure 3 using the lower-left option in the Chart subtype matrix shown

[12] Once the Chart Wizard is completed and the 3D surface chart appears, it is necessary to make a minor adjustment to the way in which the x-axis is constructed. Select (click on) the x-axis and then choose the Format/Selected Axis command. Look under the Scale tab of the Format Axis window, and remove the check mark from the "Value (Z) axis crosses between categories" option.

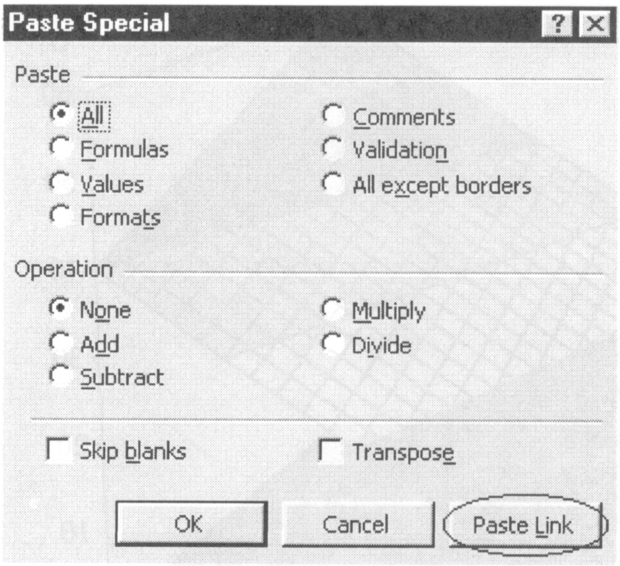

Figure 4. The Edit/Paste Special Command Window

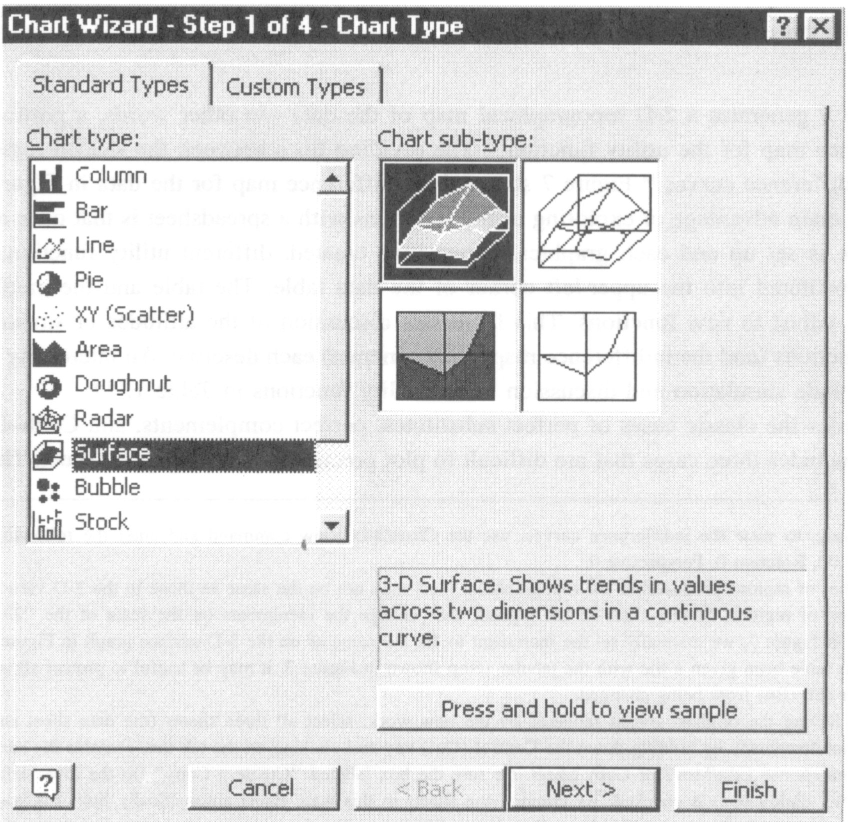

Figure 5. Window Associated with Insert/Chart Command

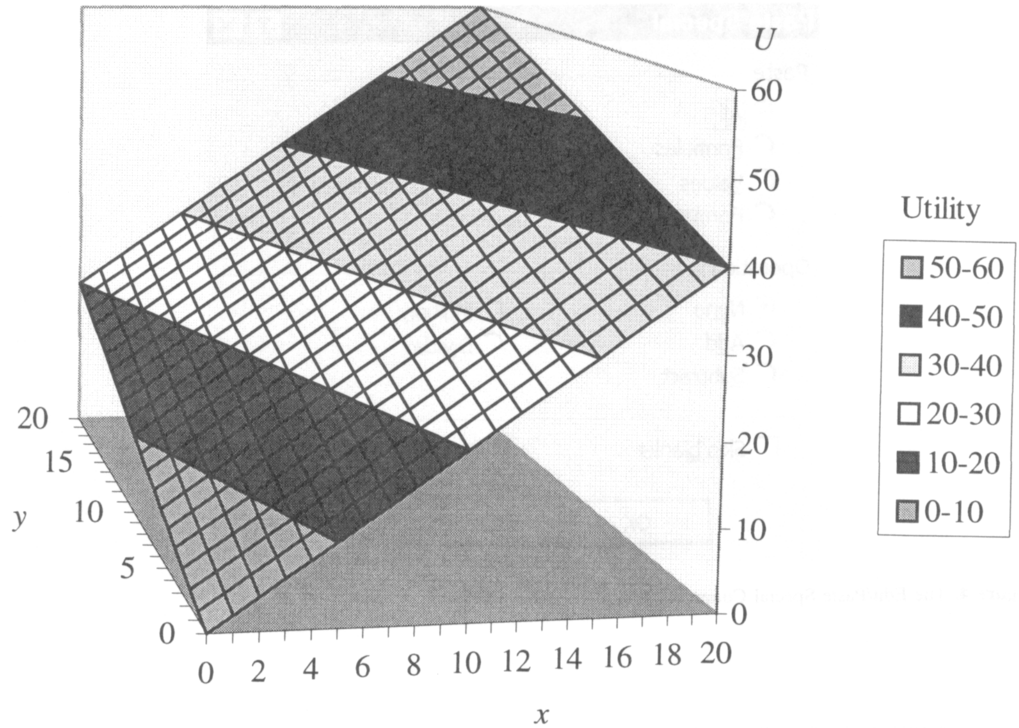

Figure 6. 3-D Plot of the Utility Function $U = 2x + y$

in Figure 5 generates a 2-D topographical map of the data—in other words, a portion of the indifference map for the utility function.[13] The dividing lines between the shaded regions represent indifference curves.[14] Figure 7 shows the indifference map for the data in Figure 4.

The main advantage of exploring utility functions with a spreadsheet is that once an initial data table is set up and each graphical perspective created, different utility functions can be easily substituted into the upper-left corner of the data table. The table and the graphs automatically adjust to new functions. This facilitates discussion of the attitudes or situations that utility functions (and the indifference maps they generate) each describe. An interesting exercise might include simulation and discussion of the utility functions in Table 1.[15,16]

Besides the classic cases of perfect substitutes, perfect complements, and Cobb–Douglas, Table 1 includes three cases that are difficult to plot precisely without a spreadsheet. The quasi-

---

[13] Alternatively, to view the indifference curves, use the Chart/3-D View command and enter the following settings: Elevation 90, Rotation 0, Perspective 0.

[14] The number of regions produced in the topographical view may not be the same as those in the 3-D view. To adjust the number of regions, double-click on the legend, then change the increments on the scale of the "Z-axis." For example, in Figure 7, we manually set the increment to 10, the same as on the 3-D surface graph in Figure 6.

[15] If students have been given a file with the tabular setup shown in Figure 3, it may be useful to protect all cells except the corner cell (B6) from being changed.

[16] To avoid having the original graphs replaced by the new work, select all three sheets (the data sheet and the two graphs) simultaneously by holding down the Control (Ctrl) key and clicking on the tab that contains the name of each sheet. Then choose Edit/Move or Copy Sheet. Be sure the box labeled "Create a Copy" (in the lower-left corner of the resulting dialog box) is checked. By copying the sheets in this way, Excel automatically links the graphs to the new (copied) data sheet, not the old sheet. Entering a different utility function into the corner of the new data sheet will cause the graphs to adjust.

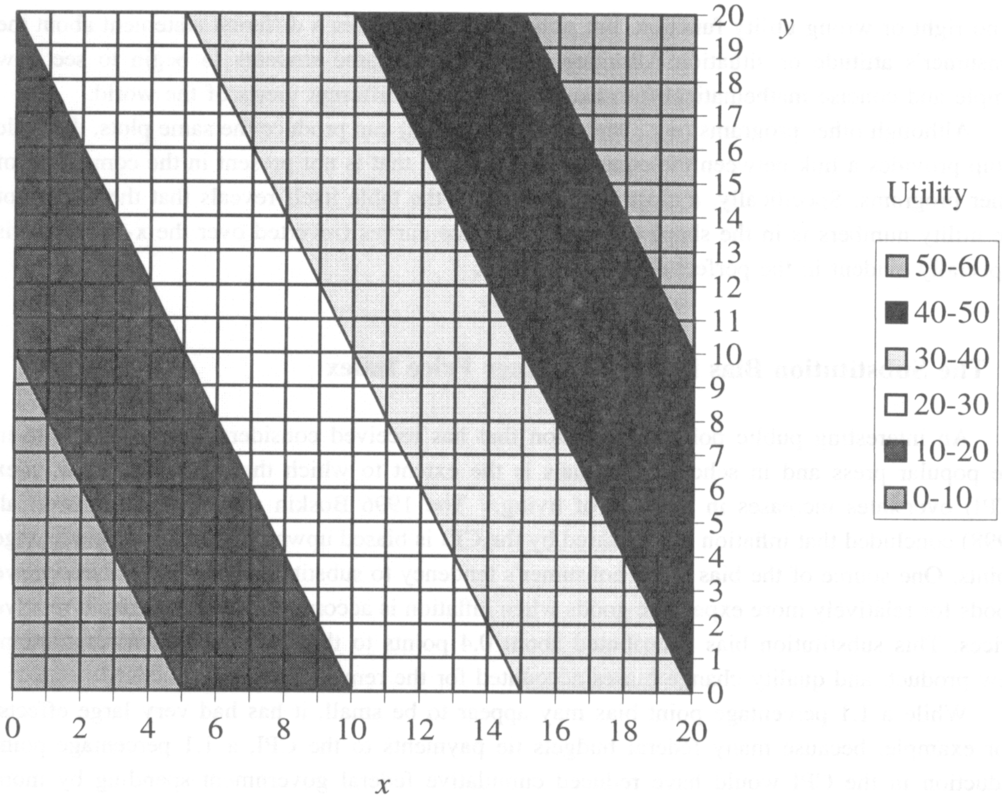

Figure 7. Indifference Map for $U = 2x + y$

linear function is an interesting alternative to the standard Cobb–Douglas because the preferences are convex but the marginal rate of substitution (MRS) is constant at any given x value. Such preferences yield a demand function for x that is independent of income. The concave preferences exhibit an increasing MRS and thus may be useful in introducing situations where an individual chooses to behave in an extreme manner. The last function in Table 1 explores the idea of x being undesirable while still maintaining the assumption of convexity. All of the functions listed in Table 1 will plot smoothly for the x and y range used in Figure 3.

This exercise has the potential to give students a better idea of the origin of indifference curves as well as the terminology of an indifference map. It also highlights the idea that there

Table 1. Other Utility Functions to Try in this Exercise

Utility Function	Description
$U = x + y$	Perfect one-to-one substitutes
$U = \min(x, y)$	Perfect complements ($y = x$)
$U = \min(5x, y)$	Perfect complements ($y = 5x$)
$U = xy$	Cobb–Douglas (MRS $= -y/x$)
$U = x^2y$	Cobb–Douglas (MRS $= -2y/x$)
$U = 2\sqrt{x} + y$	Quasi–linear preferences
$U = x^2 + y^2$	Concave preferences
$U = [(-1/20)x^2] + y$	Convex preferences but x is undesirable

is no right or wrong utility function, but rather that each makes a different statement about the consumer's attitude or situation. Ultimately it may lead some students to begin to see how simple and concise mathematical formulas can articulate different views of the world.

Although other programs (e.g., Maple, Mathematica) can produce the same plots, the table setup provides a link between the equation and the plot that is not present in the commands of other programs. Specifically, a close examination of the table itself reveals that the pattern of the utility numbers is in the shape of the indifference curves (inverted over the x-axis); this is especially evident in the perfect complement case.

3. The Substitution Bias in the Consumer Price Index

An interesting public policy application that has received considerable attention both in the popular press and in scholarly journals is the extent to which the consumer price index (CPI) overstates increases in the cost of living.[17] The 1996 Boskin report (see Boskin et al. 1998) concluded that inflation as measured by the CPI is biased upward by about 1.1 percentage points. One source of the bias is the consumer's tendency to substitute relatively less expensive goods for relatively more expensive goods when inflation is accompanied by changes in relative prices. This substitution bias contributed about 0.4 points to the bias; the outlet substitution, new product, and quality change biases accounted for the remainder.[18]

While a 1.1 percentage point bias may appear to be small, it has had very large effects. For example, because many federal budgets tie payments to the CPI, a 1.1 percentage point reduction in the CPI would have reduced cumulative federal government spending by more than $1 trillion over just 12 years. Just as importantly, over the last 25 years, average real earnings would be measured as rising, not falling, and real median income would be measured as rising rather than stagnating (Boskin et al. 1998).

Provided students are made aware of some basic demand functions, Excel can help students explore this issue of the substitution bias from a numerical perspective. For example, suppose an individual has Cobb–Douglas preferences represented by the utility function $U = x^a y^b$, where a and b are positive constants. For an individual facing fixed prices (p_x and p_y) and a fixed income (I), Equations 3.1 and 3.2 give the utility maximizing quantities x^* and y^*:[19]

---

[17] For an overview of the research in this area, see Moulton (1996) and the symposium entitled "Measuring the CPI" in the Winter 1998 issue of the *Journal of Economic Perspectives.* Boskin et al. (1998) appears in this symposium.

[18] The outlet substitution bias accounts for about 0.1 percentage points of the total bias and is derived from the fact that there has been a trend toward consumers shopping at discount stores and on weekends (when there are sales). The CPI methodology uses the same stores and weekday data over the life of the original basket of goods. The new-product bias results from the fact that new products are often not included in the CPI for several years after they penetrate the market (e.g., VCRs, microwave ovens), and during this time the prices of new products usually fall considerably. Furthermore, statistics do not accurately subtract the value of quality improvements from price increases, especially with regard to services (e.g., health services). These latter two factors together contributed 0.6 percentage points toward the total estimated bias.

[19] For students without a calculus background, it is possible to derive these formulas using the results from the utility function exercise discussed in the previous section. The essential step in the derivation is finding the marginal rate of substitution (MRS) expression and then setting it equal to the price ratio ($-p_x/p_y$). Solving this optimum condition simultaneously with the budget constraint yields the demand functions. To justify the MRS expression, the indifference map associated with the $U = xy$ function could be used to show that the slope of a tangent line to any indifference curve is $-y/x$. Then the indifference map associated with the $U = x^2y$ function could be used to show that the slope becomes $-2y/x$. This provides some justification for using the expression $-ay/bx$ as the MRS in the most general case.

	A	B	C	D	E
1		Original values	copy & paste,	New values	
2	*a*	1	then	1	
3	*b*	1	update prices	1	
4	p_x	$4.00	⟶	$4.20	
5	p_y	$1.00	⟶	$1.40	
6	*I*	$400		$400	
7	*x**	50		47.62	
8	*y**	200		142.86	
9	*U*	10,000		6,802.72	
10					
11				Laspeyres price index	
12				1.225	
13					
14					
15		Cost of living raise		Holding *U* constant	
16	*a*	1		1	
17	*b*	1		1	
18	p_x	$4.20		$4.20	
19	p_y	$1.40		$1.40	
20	*I*	$490.00		$484.97	…by adjusting income
21	*x**	58.33		57.74	
22	*y**	175.00		173.21	
23	*U*	10,208.33		10,000	Goal Seek drives this cell to original U…
24					
25				Ideal price index	
26				1.212	

Figure 8. Substitution Effect Bias in the Laspeyres Price Index

$$x^* = \left(\frac{a}{a+b}\right)\frac{I}{p_x} \tag{3.1}$$

and

$$y^* = \left(\frac{b}{a+b}\right)\frac{I}{p_y}. \tag{3.2}$$

Suppose initially $a = b = 1$ and that $p_x = \$4$, $p_y = \$1$, and $I = \$400$. Figure 8 shows this original information organized into an Excel worksheet. Cells B7 and B8 contain the optimal expressions for *x* and *y*, while cell B9 contains the utility function formula. After entering this information, it is possible to copy and paste this original scenario to another location on the worksheet, change the prices, and compare the results.

In Figure 8, for example, it is assumed that inflation occurs in such a way that the price of *x* rises from \$4 to \$4.20 and that the price of *y* rises from \$1 to \$1.40 (so that the relative price ratio falls from 4 to 3). With the new optimum established, compute the price index using the basic Laspeyres formula (the cost of original bundle at the new prices divided by cost of the original bundle at the original prices).[20] In this example, the Laspeyres price index shows an increase of 22.5% in the cost of living.

[20] In January 1999, the Bureau of Labor Statistics moved away from a simple Laspeyres formula in order to reduce the

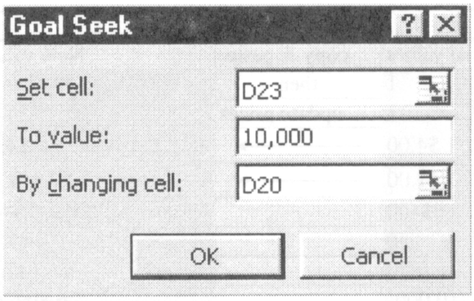

Figure 9. The Tools/Goal Seek Window

After the price index is computed, a third what-if scenario can be created by copying the second situation to a new location on the worksheet and giving the person a cost-of-living raise as measured by the Laspeyres price index (by multiplying income by 1.225). Note that the utility level is higher after the cost-of-living raise. This confirms that the Laspeyres price index overstates the increase in the cost of living associated with the price changes. The overstatement occurs because of the substitution effect associated with the lower relative price of x (note that x^* is higher and y^* lower than in the initial scenario).

Given that a cost-of-living raise based on the Laspeyres price index makes the consumer better-off, a natural follow-up question is, "How much of an increase in income would be needed to leave the consumer with the same level of utility as in the original scenario?" To answer this question, copy the second scenario to another location on the worksheet, and then use Excel's Tools/Goal Seek command. In response to the Goal Seek prompts shown in Figure 9, instruct Excel to set cell D23 (the utility cell) to a value of 10,000 (the original level of utility) by changing cell D20 (the income cell). As shown in Figure 8, the result is an income level of $484.97. To recap, this means that an income level of $484.97, given the new prices, will just return the consumer to the original level of utility. Taking the ratio of this adjusted income level to the original income level of $400 yields an ideal price index of 1.212. Comparing the ideal price index with the Laspeyres price index leads to the conclusion that the Laspeyres price index overstates the true increase in the cost of living by 1.3 percentage points. The steps in this exercise highlight one of the main advantages of Excel. While the Goal Seek feature lets students find a numerical solution to a complex problem without doing the computations, setting up the spreadsheet requires students to have a clear understanding of the analytical structure of the model and the linkages between the variables.

It is possible to analyze other price-index formulas (e.g., the Paasche index) using the basic approach taken in Figure 8.[21] To further emphasize the role that relative price shifts play in the analysis, however, a good exercise is to have students repeat the Laspeyres analysis just described and attempt to find the degree of overstatement associated with inflationary scenarios exhibiting different relative price shifts. Table 2, for example, clearly shows that the greater the change in relative prices (recall that the original relative price ratio was four), the more the Laspeyres price index overstates the increase in the cost of living. Alternatively, students could

---

degree of substitution bias in the consumer price index. For more on this change, see Dalton, Greenlees, and Stewart (1998). The change does not completely eliminate substitution bias.

[21] The Paasche index should exhibit opposite effects. As a further application, compare the results for the Laspeyres, Paasche, and calculated ideal indices with those of a chain-weight index, which is the geometric average of a Laspeyres and Paasche index. The chain-weight index should be close to the ideal.

Table 2. Sensitivity of the Laspeyres Price Index Bias to the Degree of Relative Price Change

New p_x	New p_y	New Relative Price Ratio (p_x/p_y)	Value of the Laspeyres Price Index	Number of Percentage Points the Price Index Overstates the Cost-of-Living Increase
$4.20	$1.40	3	1.225	1.3
$3.20	$1.60	2	1.200	6.9
$5.00	$1.25	4	1.250	0.0

be asked to contrast the overstatement results from the Cobb–Douglas case with results generated by a perfect complement utility function (exhibiting no substitution effect) or a quasi-linear utility function of the form $U = 40\sqrt{x} + y$, which exhibits a larger substitution effect. Repeating the analysis in Figure 8 with this quasi-linear function, for example, leads to an overstatement of 2.9 percentage points.

The web site that accompanies this paper provides two examples related to the discussion in this section. The first supplementary example explores the extent to which the traditional (fixed-weight) GDP deflator suffers from substitution bias, and the second supplementary example provides an application of the Tools/Goal Seek command in an industrial organization context.

4. Hicksian Decomposition and Consumers' Surplus

By using the basic spreadsheet setup shown in Figure 8 and the Tools/Goal Seek command, it is possible to compute the size of the substitution and income effects associated with a price change using the Hicks method of holding utility constant. While this is a relatively technical exercise, it may be of some help in enabling students to make the transition from the purely graphical analysis used in most textbooks to a more numerically oriented approach. Figure 10, for example, shows the substitution and income effects associated with an increase in the price of good x from $2 to $4 for a consumer with a Cobb–Douglas utility function of the form $U = x^2y$ and a fixed income of $300 ($p_y$ held constant at $1). The Tools/Goal Seek command

	A	B	C	D	E	F	G	H	I
1		Original equilibrium		New equilibrium		Adjusted equilibrium			
2	a	2		2		2			
3	b	1		1		1			
4	p_x	$2.00		$4.00		$4.00			
5	p_y	$1.00		$1.00		$1.00			
6	I	$300		$300		$476.22	income adusted using Goal Seek		
7	x*	100		50.00		79.37			
8	y*	100		100.00		158.74			
9	U	1,000,000		250,000.00		1,000,000	to return utility to original level		
10									
11						Substitution effect	-20.63		
12						Income effect	-29.37		
13						Total change in x*	-50.00		

Figure 10. Calculation of Substitution and Income Effects (Hicks Method)

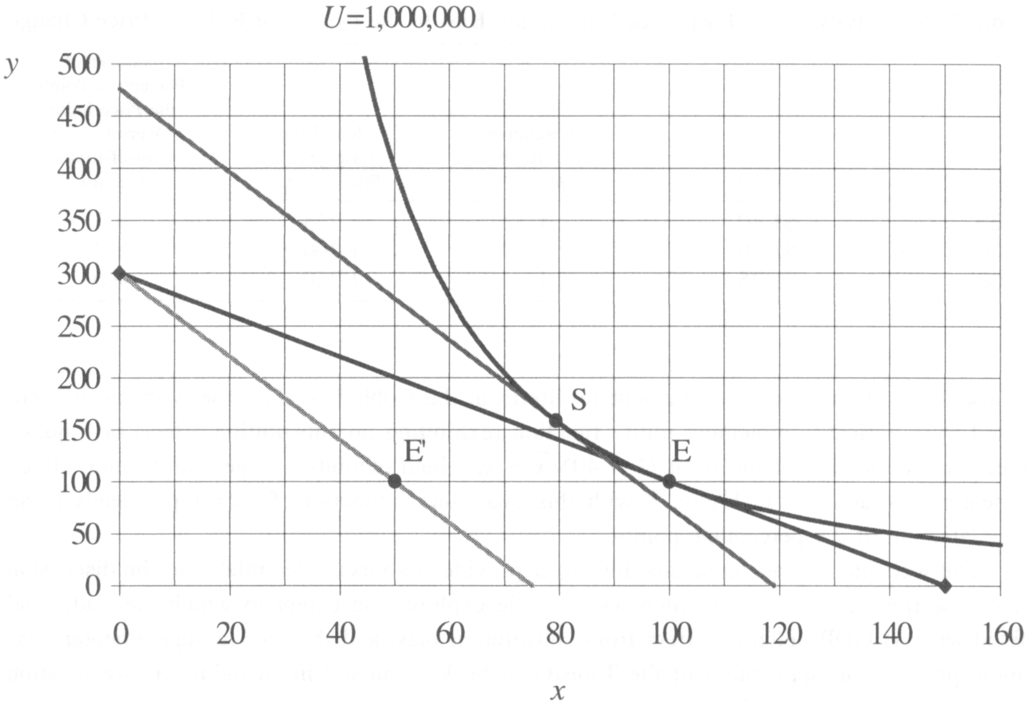

Figure 11. Substitution and Income Effects (Hicks Method)

derives the adjusted income level; Goal Seek sets cell F9 to the original utility level of 1,000,000 by changing cell F6.

The Insert/Chart command (see Figure 5, above) can generate a graphical look at the results by choosing the XY (Scatter) plot option. The eventual result is shown in Figure 11, where the point E denotes the original optimum, E′ is the new optimum, and the point S denotes the adjusted optimum from which the substitution and income effects are derived. Figure 12 shows the worksheet data underlying Figure 11. Complex charts such as the one shown in Figure 11 must be built from the worksheet data curve by curve. The web supplement contains complete instructions for constructing Figure 11.

Using the Hicks method to find the substitution and income effects associated with a price increase leads naturally to the topic of consumer's surplus. Comparing the original and new optima in Figure 10, it is clear that the consumer is worse off after the price increase (U = 250,000 < 1,000,000). However, because utility numbers have a purely ordinal interpretation (as a ranking), nothing about the degree to which the consumer's well being changed can be determined. Instead, one way to gauge how much the consumer has been hurt by the price increase is to determine the minimum amount of additional income a person would willingly accept to be subjected to the price increase. This amount is the compensating variation (CV) in income and can be determined by comparing the adjusted income and the original income in Figure 10 ($476.22 vs. $300 means that the CV is $176.22). An alternative way to generate a monetary measure of the degree to which the consumer is harmed by the price increase is to find the maximum amount the person would be willing to give up to avoid being subjected to the price increase. This amount is the equivalent variation (EV) in income. The essential step

	A	B	C	D	E	F	G	H
1	Original budget line			Original indifference curve			Equilibrium points	
2	x	y		x	y		x	y
3	0	300		40	625.00		100	100
4	150	0		45	493.83		50	100
5				50	400.00		79.37	158.74
6	New budget line			55	330.58			
7	x	y		60	277.78			
8	0	300		65	236.69			
9	75	0		70	204.08			
10				75	177.78			
11	Adjusted budget line			80	156.25			
12	x	y		85	138.41			
13	0	476.22		90	123.46			
14	119.06	0		95	110.80			
15				100	100.00			
16				105	90.70			
17				110	82.64			
18				115	75.61			
19				120	69.44			
20				125	64.00			
21				130	59.17			
22				135	54.87			
23				140	51.02			
24				145	47.56			
25				150	44.44			
26				155	41.62			
27				160	39.06			

Figure 12. Worksheet Data Underlying the Substitution/Income Effect Graph

in finding the EV is to use the Tools/Goal Seek command and determine the level of income that would bring about the new level of utility at the original prices.

An alternative way to approximate the CV and EV for small price changes is to compute the change in consumer's surplus (CS) association with the price change. Willig (1976) derived this point in his widely cited paper. Consider, for example, a consumer with a Cobb–Douglas utility function $U = xy$ and a fixed income (I) of \$300. Substituting this information into the Cobb–Douglas demand function for good x (Eqn. 3.1), the consumer's demand curve for x becomes

$$x^* = \left(\frac{1}{1+1}\right)\frac{300}{p_x} = \frac{150}{p_x}. \tag{4.1}$$

Figure 13 plots this demand curve. Suppose that the price of good x (p_x) is initially \$2, but then p_x rises to \$2.50. For students with a background in integral calculus, it is possible to compute the reduction in the area of CS as follows:

$$\int_2^{2.50} \frac{150}{p_x}\,dp_x = 150 \ln p_x]_2^{2.50} = 150(\ln 2.50 - \ln 2) = 33.47.$$

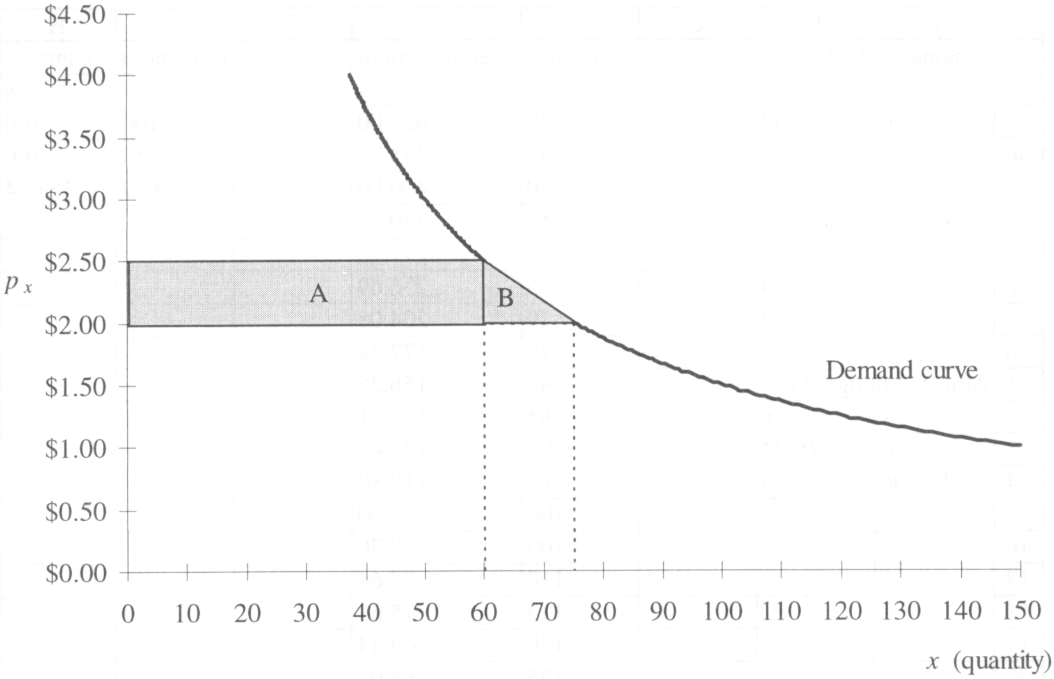

Figure 13. Computing the Reduction in Consumer Surplus Associated with a Price Increase

For those without such a background, it is possible to approximate the area by the shaded region in Figure 13. It can be computed by adding the area of rectangle A (60 × 0.5) and the area of triangle B (½ × 15 × 0.5) for a total reduction in CS of $33.75.

Figure 14 shows how to set up an Excel worksheet to compute the CV and EV associated with the increase in p_x (the price of good y (p_y) is held constant at $1). The Tools/Goal Seek command computed the adjusted income levels. Note that the absolute value of the CS reduction is a good approximation to the more exact monetary measures of the change in the consumer's well being. Note also how the value of the CS reduction lies between the value of the EV and CV. Relating the change in the area of CS to the EV and CV may be useful in helping students to interpret the economic meaning of a change in CS better.

The idea that the change in CS is a monetary measure of the change in the consumer's well being can be further developed by exploring a quasi-linear utility function such as $U = 60 \ln x + y$. The demand functions associated with this utility function are

$$x^* = \frac{60p_y}{p_x} \quad \text{and} \quad y^* = \frac{I}{p_y} - 60.$$

Repeating the same steps as in the Cobb–Douglas case, it can be shown that $|\Delta CS| = |CV| = |EV| = |\Delta U|$. In other words, for quasi-linear utility functions, the change in consumer's surplus is an exact monetary measure of the impact of the price change on the consumer's well being.

5. Optimal Aggregate Consumption and Saving

Another important issue in consumer theory is the trade-off between present and future consumption. From the standpoint of an individual consumer, an increase in saving by $1

	A	B	C	D	E	F	G	H
1		Original		New		Hold Utility at Original Level		
2	a	1		1		1		
3	b	1		1		1		
4	p_x	$2.00		$2.50		$2.50		
5	p_y	$1.00		$1.00		$1.00		*CV*
6	I	$300.00		$300.00		$335.41	…by adjusting income	$35.41
7	x^*	75.00		60.00		67.08		
8	y^*	150.00		150.00		167.71		
9	U	11,250.00		9,000.00		11,250.00	Goal Seek drives this cell to original U…	
10								
11								
12		Original		New		Reduce Utility to New Level		
13	a	1		1		1		
14	b	1		1		1		
15	p_x	$2.00		$2.50		$2.00		
16	p_y	$1.00		$1.00		$1.00		*EV*
17	I	$300.00		$300.00		$268.33	…by adjusting income	$31.67
18	x^*	75.00		60.00		67.08		
19	y^*	150.00		150.00		134.16		
20	U	11,250.00		9,000.00		9,000.00	Goal Seek drives this cell to new U…	

Figure 14. Computing the CV and EV Associated with a Price Increase

reduces present consumption by $1; in exchange, the consumer receives future consumption of $1 plus the interest on the dollar. For an economy, a $1 increase in saving results in a loss of $1 of present consumption, but because the dollar saved becomes capital investment, the ability to produce consumption (and capital) goods in the future grows.

Many principles texts present this trade-off in a discussion of production possibility frontier (PPF) diagrams. In these examples, the choice is between producing capital goods (by saving) and producing consumer goods. The choice exists because an economy that produces a relatively large quantity of consumer goods suffers a slower expanding PPF; the slow expansion of the PPF limits the possibility of consumption in the future. The exercise below explores the saving rate (and thus the capital accumulation rate) that maximizes total consumption over time. This saving rate must be greater than zero and less than unity. The model used to study this issue is the Solow (1956) long-run model of an economy.

The Solow (1956) model hinges on the concept of balanced growth. Balanced growth is the technical condition in which capital (K) grows at the same rate as the effective labor supply (LE). The effective labor supply is an index of both the number of workers and the productivity of the workforce.[22] In his famous stability proof, Solow showed that economies naturally tend toward a state of balanced growth, so this condition describes the long-run equilibrium for the economy.[23] Consequently, to study total long-run consumption, we need only look at consump-

[22] For example, if the productivity of each worker grows by 10% (and the labor force remains the same size), or the number of workers in the labor force grows by 10% (and productivity remains the same), the effective labor supply grows by 10%.

[23] For a discussion of the stability proof, see Mankiw (1997). Other intermediate macroeconomics textbooks also derive the proof.

tion under balanced growth. Phelps (1961) first discussed the properties of the saving rate that maximizes long-run consumption in the context of the Solow model. Phelps dubbed the optimal saving rate the "golden rule."[24]

Because the purpose of this example is to explore the concept of an optimal saving rate, and not the growth model, the general discussion of the Solow dynamic balanced growth model is left to the third web supplement. The exercise below uses the Solver feature to find the golden rule saving rate for numerical examples of a static (one period) model; the results are the same as in a dynamic framework.

Consider the following variant of the Solow (1956) model. The labor force grows at rate n, and labor productivity grows at rate g. Thus, the effective labor supply grows at rate $(n + g)$. Equation 5.1 defines labor supply growth; Equation (5.2) defines the balanced growth rate of capital:

$$\Delta(LE) = (n + g)(LE) \tag{5.1}$$

$$\Delta K = (n + g)K. \tag{5.2}$$

Households either consume or save income, where the saving rate is s. This defines consumption (C) in terms of the saving rate (s) and income (Y):

$$C = Y - sY. \tag{5.3}$$

All saved output becomes capital and is put into production the following period. Thus, for each period, total saving(s) equals gross investment (I):

$$S = sY = I. \tag{5.4}$$

Capital depreciates at rate d. Net investment (ΔK) is thus gross investment (I) less depreciation (dK).

For capital to grow at the balanced growth rate of $n + g$, gross investment must be equal to

$$I = (n + g + d)K. \tag{5.5}$$

Substituting Equation 5.4 into Equation 5.5 forms Equation 5.6.

$$sY = (n + g + d)K. \tag{5.6}$$

By solving Equation 5.6 for the capital/output ratio (K/Y), we see that K/Y is a constant:

$$K/Y = s/(n + g + d). \tag{5.6'}$$

Equation 5.6′ is the balanced growth capital/output ratio. It defines long-run capital as a function of output, and vice versa. When paired with a production function, Equation 5.6′ and the production function jointly define long-run capital and GDP.

The aggregate production function assumes that output depends on a technology parameter (A), capital (K), and effective (adjusted for productivity) labor (LE). A Cobb–Douglas form is used, with parameters $a < 1$ and $A > 0$:

$$Y = AK^{a}(LE)^{1-a}. \tag{5.7}$$

---

[24] Some intermediate macroeconomics textbooks, including Mankiw (1997), derive the golden rule. Specifically, consumption is maximized when the marginal product of capital equals the sum of the labor force growth rate, labor productivity growth rate, and depreciation rate. Footnote 29 goes into more detail on the derivation.

	A	B	C	D	E	F	G	H	I	J	K	L
1	**PARAMETERS**											
2	s (saving rate)	0.000	0.100	0.200	0.300	0.400	0.500	0.600	0.700	0.800	0.900	1.000
3	n (labor force growth rate)	0.015	0.015	0.015	0.015	0.015	0.015	0.015	0.015	0.015	0.015	0.015
4	g (labor productivity growth rate)	0.010	0.010	0.010	0.010	0.010	0.010	0.010	0.010	0.010	0.010	0.010
5	d (capital depreciation rate)	0.050	0.050	0.050	0.050	0.050	0.050	0.050	0.050	0.050	0.050	0.050
6	A (C-D technology index)	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
7	a (C-D parameter)	0.300	0.300	0.300	0.300	0.300	0.300	0.300	0.300	0.300	0.300	0.300
8	LE (Labor force)	10.000	10.000	10.000	10.000	10.000	10.000	10.000	10.000	10.000	10.000	10.000
9	**VARIABLES**											
10	K (capital) [K=Y (s/(n+g+d)]	0.000	15.083	40.600	72.458	109.287	150.319	195.042	243.090	294.180	348.087	404.627
11	Y (output) [Y= $AK^a(LE)^{(1-a)}$]	0.000	11.312	15.225	18.114	20.491	22.548	24.380	26.045	27.579	29.007	30.347
12	C (consumption) [C = Y - sY]	0.000	10.181	12.180	12.680	12.295	11.274	9.752	7.814	5.516	2.901	0.000

Figure 15. Solow Balanced Growth Model

Equations 5.6′ and 5.7 together define the long-run levels of output and capital, which are a function of the effective labor supply (*LE*), total factor productivity (*A*), and parameters. Depending on the value of *a,* solving Equations 5.6′ and 5.7 simultaneously will likely require a numerical method. Because long-run consumption is a function of long-run income, the solutions to Equations 5.6′ and 5.7 define long-run consumption.

Figure 15 displays an example of a spreadsheet that has been set up to solve the balanced growth economy. The sheet is best set up in columns, with the first column displaying parameter and variable names and the other columns containing parameter values and equations. The necessary rows include each of the parameters (*s, n, g, d, A, a*), the (initial) size of the effective labor supply (*LE*), and the three endogenous variables (*K, Y,* and *C*). After typing in the parameter and variable names, type in values for the parameters and *LE,* as well as the equations defining *K, Y,* and *C.* The formula used to define *K* is a rearrangement of Equation 5.6′, $\{K = Y[s/(n + g + d)]\}$. The formula used to define *Y* is the right-hand side of the production function (Eqn. 5.7), and the right-hand side of Equation 5.3 provides a formula for *C.* Each of the formulas refers to the appropriate cell (*s, Y,* etc.) for its values. For example, in Figure 15, the equation for capital (in B10) is =B11*(B2/(B3+B4+B5)).

When the formula for *Y* is entered (in B11, above), it creates a circular reference with the *K* formula cell (in B10), and Excel displays an error. To allow Excel to solve the equations simultaneously, select Tools/Options. . ., then the Calculation tab. After the window in Figure 16 appears, check the Iteration box (as in Figure 16); from this window, it is also possible to change the iteration options from default values.[25] When OK is pressed, Excel will solve the equations simultaneously using an iterative numerical method.[26] This Iteration option allows Excel to solve a variety of complex systems, including those with several variables. For an example of using the Iteration option to solve the IS–LM/AS–AD model, see the fourth web supplement.[27]

---

[25] The Maximum iterations option sets the maximum number of times Excel will search for a solution in each round of calculations. Pressing the F9 key (the "recalculate spreadsheet" key) reinitiates the solving process. The Maximum change option specifies the precision of the calculation procedure. For example, a value of 0.001 specifies that Excel will cease searching for a solution when the candidate solution is 0.001 or less from the true solution.

[26] Be aware that sometimes, due to a bug in the Excel solving algorithm, Excel does not find a solution and returns a #NUM! error. This problem appears to be sensitive to parameter values and the value that resides in the cell before entering the formula.

[27] For an early example of simultaneous equation solving without the iteration function, see Johansson (1985). See Goddard, Romilly, and Tavakoli (1995) for a spreadsheet example of the Mundell–Flemming (IS–LM–BP) model. Although they do not specify how the spreadsheet solves the model, we assume that an iterative solution is used. See

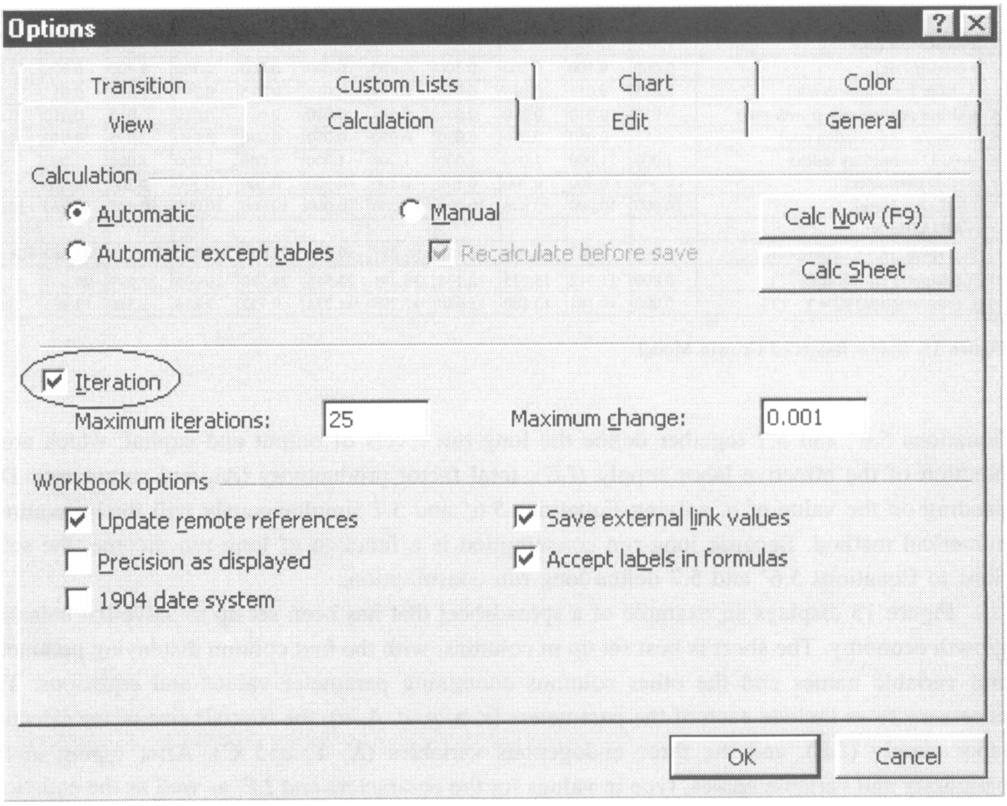

Figure 16. The Tools/Options. . ./Calculation Window

Figure 15 displays 11 columns of data. To create columns C through L, copy column B into this range, and then change the value for *s* in each. This displays a full range of *s* values in order to explore the nature of the maximum.[28] The consumption data in row 12 (and the plot of such data) show that a saving rate around 0.3 maximizes long run consumption. Figure 17 clearly displays this result. At this point, the trade-off between present and future consumption is clear. Higher consumption (lower saving) rates result in higher current consumption at the expense of lower future capital, and consequently lower future consumption. Next, we use the Tools/Solver add-in to calculate the exact saving rate that maximizes long-run consumption.[29]

The Solver allows the user to do a variety of constrained minimization, maximization, and

---

Adams and Kroch (1989) for examples of the Keynesian Cross, IS–LM, and Phillips Curve models using Lotus 1-2-3. Judge (1990a) develops a multiplier-accelerator model using Lotus 1-2-3. Judge (1990b) develops an IS–LM model using Lotus 1-2-3 and solves it by using matrix commands. Jones and Judge (1990) develop a dynamic macroeconomic model.

[28] For other examples, it may be desirable to use additional columns to create a time series of data. In this case, after column B is copied to columns C through L, the cell containing the labor force data would be changed to a formula to reflect Equation 5.1; for example, cell C8 would contain a formula =B8*(1+B3+B4), and this formula should be copied to cells D8 through L8. See web supplement 3 for more details on setting up such a time series model.

[29] Phelps (1961) showed that the golden rule condition is the marginal product of capital (MP_K) = $n + g + d$. This may be algebraically derived by maximizing consumption ($Y - sY$) subject to the balanced growth condition $sY = (n + g + d)K$ by choosing K. For more advanced classes, it is suggested that lines to track the MP_K ($aA(LE/K)^{1-a}$) and ($n + g + d$) be added to the spreadsheet in Figure 15 to verify this result. Note that the MP_K changes one for one with changes in n, g, and d in the given Cobb–Douglas formulation, so a is the only parameter that affects the golden rule s.

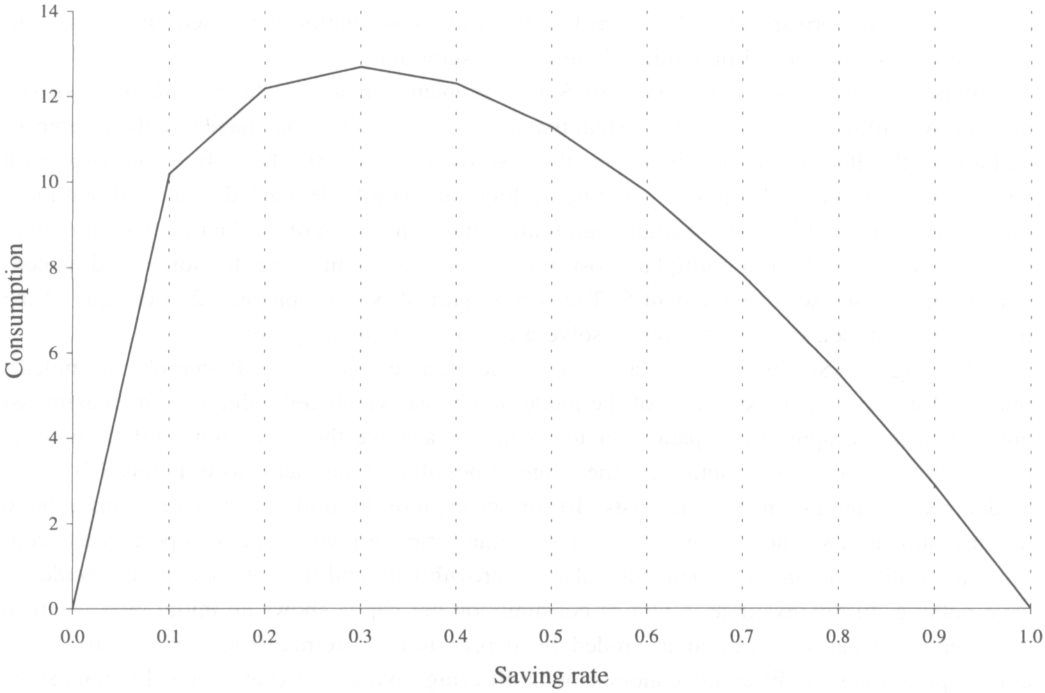

Figure 17. Consumption for a Range of Saving Rates

goal-seek problems. To access the Solver, it must be installed as an add-in (Tools/Add-Ins. . .). When installed once, it permanently appears in the Tools menu. For the problem above, we wish to find the saving rate that maximizes consumption, where consumption is a function of income and the saving rate. After selecting the Solver, the window in Figure 18 appears.

As implied by Figure 18, the Solver maximizes, minimizes, or goal seeks a target cell (which must contain a formula) by changing another designated cell or cells. Figure 18 displays

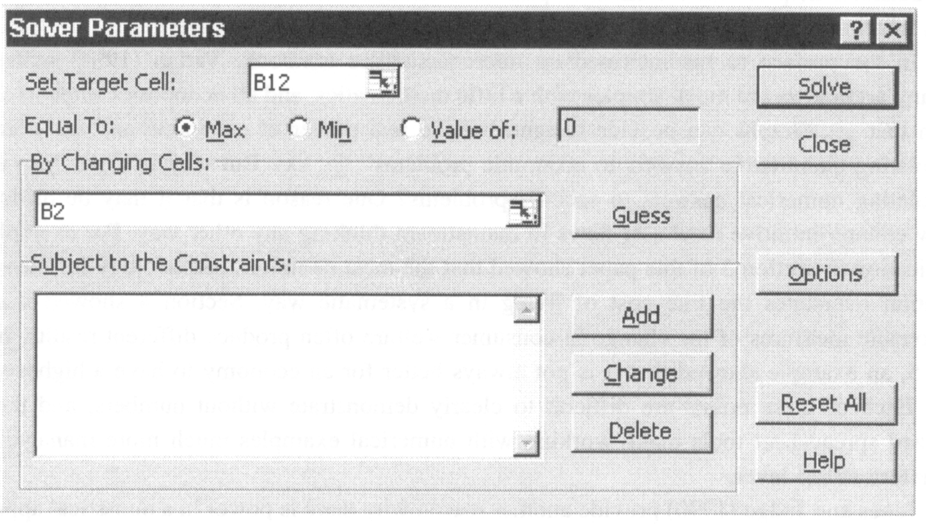

Figure 18. The Solver Window

cell addresses to correspond with Figure 15. When the Solve button is pressed, the Solver will find that $s = 0.300$ indeed maximizes long-run consumption.

It is possible to add constraints to Solver problems in a straightforward, menu-driven manner. Any of the formulas in the system that are to be optimized may have circular references as long as the Iteration option is active. Because of its flexibility, the Solver can assist in a variety of economic applications, including finding the quantity demanded of a good that maximizes utility given a budget constraint and finding the distribution of production that minimizes costs. For an example of a multiplant cost-minimization problem using the Solver and its constraints option, see web supplement 5. The second part of web supplement 2, a critique of the limit-pricing model, uses the Solver to solve a common oligopoly problem.

Although the Solver acts as a black box to hide the algebraic optimization problem, students must still understand the structure of the model to choose which cell value is to be maximized and which is the appropriate parameter to change to achieve the maximum. Further, plotting the level of long-run consumption for the range of possible saving rates (as in Figure 17) shows students that a unique maximum exists. To further explore the trade-off between consumption and investment, a spreadsheet with a dynamic (time series) growth model can portray an economy that initially is on its golden rule balanced growth path and then at some point decides to save nothing. In this example, a plot of consumption per capita shows an initial increase, then a decrease (to zero) as capital is eroded by depreciation. Alternatively, it is useful to plot consumption rates for different economies with differing saving rates (but identical initial capital endowments) over time on the same chart. For a description of setting up a dynamic growth model, see the third web supplement.

6. Conclusion

While this paper offers no evidence that the use of spreadsheet software increases student learning in economics, the applications presented above suggest to us that spreadsheets may be able to help instructors delve more deeply into traditional topics or tackle more advanced topics. Spreadsheets help to make difficult topics in economics more accessible because of the ease with which they generate a wide array of quantitative solutions.

In the preface to his intermediate microeconomics textbook, Varian (1996) writes that "Many arguments are much simpler with a little mathematics, and all economics students should learn that. . . . Graphs can provide insight, but the real power of economic analysis comes in calculating quantitative answers to economic problems" (p. xx). But why is there "power" in calculating numerical answers to specific problems? One reason is that it may be difficult to show counter-intuitive results or flaws in mainstream thinking any other way. For example, the application in section 3 of this paper showed that the most common methodology for measuring inflation overstates the true cost of living in a systematic way. Section 4 showed that two reasonable measures of the change in consumer welfare often produce different results. In section 5, an example showed that it is not always better for an economy to have a higher saving rate. Each of these points are difficult to clearly demonstrate without numbers, and the specialized spreadsheet tools make working with numerical examples much more manageable at the intermediate level.

Jones and Judge (1990) provide another reason why there is power in a numerical approach, although they also provide a note of caution:

> Rather than replacing formal analysis, the objective of this kind of experimentation is to provide an intuitive base upon which to develop a more formal treatment. Psychologically the benefits of working from simple examples to general principles in teaching are well established. The major danger in such an approach is that students may draw incorrect inferences from the perusal of a smaller set of examples. (p. 94)

Jones and Judge then note, "The risk is lessened when, as with the spreadsheet, additional examples can be generated virtually instantaneously with little cost in time or effort" (p. 94).

In summary, spreadsheet software enables the instructor to harness the clarifying power of quantitative solutions by eliminating the need for tedious calculations or extensive algebraic manipulations, yet the software is far from a black box. Instead, it requires students to have a clear definition of the economic concepts involved in the applications and be able to explore carefully the linkages between these concepts.[30] For these reasons, spreadsheets strike us as a good way to reach out to the increasingly large number of students who claim "I understand the concepts but just can't do the math." Varian (1996) emphasizes in the preface to his intermediate micro text that "It is perfectly possible to be analytical without being excessively mathematical" (p. xix). The challenge is that for many students, even basic algebraic manipulations represent an excessively mathematical approach. The spreadsheet applications presented in this paper are an attempt to deal with this challenge while meeting the need for increased analytical depth and breadth of coverage.

References

Adams, F. Gerard, and Eugene Kroch. 1989. The computer in the teaching of macroeconomics. *Journal of Economic Education* 20:269–80.

Bartlett, Robin L., and Paul G. King. 1990. Teaching economics as a laboratory science. *Journal of Economic Education* 21:181–93.

Bateman, Bradley W. 1992. The education of economists: A different perspective. *Journal of Economic Literature* 30: 1491–5.

Becker, William E. 1997. Teaching economics to undergraduates. *Journal of Economic Literature* 35:1347–73.

Benninga, Simon. 1997. *Financial modeling*. Cambridge, MA: MIT Press.

Blecha, Betty J. 1991. Economic pedagogy and microcomputer software. *Social Science Computer Review* 9:541–57.

Boskin, Michael J., Ellen R. Dulberger, Robert J. Gordon, Zvi Griliches, and Dale W. Jorgenson. 1998. Consumer prices, the consumer price index, and the cost of living. *The Journal of Economic Perspectives* 12:3–26.

Boyd, David W. 1998. On the use of symbolic computation in undergraduate microeconomics instruction. *Journal of Economic Education* 29:227–46.

Breece, James H. 1988. Computerized economic modeling: From laboratory to classroom. *Social Science Computer Review* 6:498–514.

Chizmar, John F., and Mark S. Walbert. 1999. Web-based learning environments: Issues of best teaching practice. *Journal of Economic Education* 30:248–59.

Dalton, Kenneth V., John S. Greenlees, and Kenneth J. Stewart. 1998. Incorporating a geometric mean formula in the CPI. *Monthly Labor Review* 121:3–7.

Goddard, J. A., P. J. Romilly, and M. Tavakoli. 1995. Macroeconomics and information technology applications. *Computers in Higher Education Economics Review* 9(1):6–10.

Johansson, Jan-Henrik. 1985. Simultaneous equations with Lotus 1-2-3: An example from macroeconomics. *Byte* February:399–405.

Jones, Robert, and Guy Judge. 1990. Using spreadsheets in the teaching of dynamic macroeconomic models. *British Review of Economic Issues* 12:93–112.

Judge, Guy. 1990a. *Quantitative analysis for economics and business using Lotus 1-2-3*. Hemel-Hempstead: Harvester-Wheatsheaf.

[30] See Wilkins (1992) for a nonspreadsheet-based assignment designed to facilitate the transition from graphical to algebraic models.

Judge, Guy. 1990b. Spreadsheets: Flexible tools for data analysis, model analysis and problem-solving. In *Information technology for social scientists,* edited by Richard Welford. Shipley, UK: European Research Press, pp. 39–74.

Judge, Guy. 1996. Student created spreadsheet models for teaching and learning economics: An initial self-evaluation of an experimental assignment. *Computers in Higher Education Economics Review* 10(2):16–8.

Kasper, Hirschel, Geoffrey Woglom, Vernon Dixon, et al. 1991. The education of economists: From undergraduate to graduate study. *Journal of Economic Literature* 29:1088–109.

King, Paul G., and Ross M. LaRoe. 1991. The laboratory-based economics curriculum. *Journal of Economic Education* 22:285–92.

Mankiw, N. Gregory. 1997. *Macroeconomics.* 3rd edition. New York: Worth Publishers.

Moulton, Brent R. 1996. Bias in the consumer price index: What is the evidence? *Journal of Economic Perspectives* 10:159–77.

Phelps, Edmund. 1961. The golden rule of accumulation: A fable for growthmen. *American Economic Review* 51:638–43.

Porter, Tod S., and Teresa M. Riley. 1995. *Computer exercises in microeconomics to accompany microeconomic theory.* 6th edition, by Walter Nicholson. Fort Worth: Dryden Press.

Solow, Robert. 1956. A contribution to the theory of economic growth. *Quarterly Journal of Economics* 70:65–94.

Varian, Hal R. 1996. *Intermediate microeconomics.* 4th edition. New York: W. W. Norton.

Welford, Richard, ed. 1990. *Information technology for social scientists.* Shipley, UK: European Research Press.

Whigham, David. 1998. *Quantitative business methods using Excel.* New York: Oxford University Press.

Wilkins, Allen J. 1992. Facilitating the transition from graphical to algebraic models: A teaching challenge for intermediate microeconomics. *Journal of Economic Education* 23:317–31.

Willig, Robert D. 1976. Consumer's surplus without apology. *American Economic Review* 66:589–97.

Yohe, William P. 1989. Spreadsheets in teaching economics. *Social Science Computer Review* 7:431–45.

Yohe, William P. 1995. *Interactive money and banking.* Cambridge: Blackwell.