JOHN T. HARVEY

Keynes's trade cycle: a system dynamics model

Abstract: *One of the core premises of post Keynesianism is that a proper understanding of the macroeconomy demands a dynamic analysis set in historical time. The default tool of mainstream economics, the general equilibrium framework, cannot possibly capture such a world. The goal of this paper is to suggest an alternative that allows for path-dependent, holistic analysis: system dynamics. Developed in the 1950s by Massachusetts Institute of Technology engineer Jay Forrester, not only does the worldview it embraces share much with that of post Keynesians, but there are many resources from which to draw including a wide literature, ready-made software, and numerous how-to guides. The paper provides an introduction to system dynamics and builds a model of Keynes's trade cycle theory.*

Key words*: Keynes, economic crisis, stock market, marginal efficiency of capital, post Keynesian, business cycle, system dynamics.*

JEL classifications: *C63, E12, E32.*

One of the core premises of post Keynesianism is that a proper understanding of the macroeconomy demands a dynamic analysis set in historical time. In Keynes's world, entrepreneurs hire workers and pay for inputs well before they earn income from sales of final output. They act in anticipation of profit and depend on bankers sharing their optimism if they are to find financing. All this takes place in an environment of uncertainty where agents' lack of a scientific basis for forecasting means that no significant forward-looking activity would take place were it not for animal spirits.

The default tool of mainstream economics, the general equilibrium framework, cannot possibly capture the world described above. This means that those taking the post Keynesian approach must look elsewhere. Among the

John T. Harvey is a professor of economics in the Department of Economics at Texas Christian University, Fort Worth.

ISSN 0160–3477 (print)/ISSN 1557–7821 (online)
DOI: 10.2753/PKE0160-3477360106

alternatives employed have been stock-flow consistent models (Godley and Lavoie, 2007), limit cycles (Blatt, 1983), "traditional econometric modeling, vector autoregressive (VAR) modeling, input-output analysis, simulation, differential equation modeling, and even static analysis" (Radzicki, 2008, p. 157). The goal of this paper is to suggest the addition of another technique to our toolkit: system dynamics (SD). Developed in the 1950s by Massachusetts Institute of Technology (MIT) engineer Jay Forrester, not only does the worldview it embraces share much with that of post Keynesians, but there are many resources from which to draw including a wide literature, ready-made software, and numerous how-to guides.

Mine does not represent the first attempt to make this argument.[1] However, our numbers remain relatively few, in part because of the effort necessary to learn SD. In addition, relatively less effort has been made to introduce SD to post Keynesians as compared to other heterodox traditions, particularly institutionalism. I hope to address both of these issues by building an SD model of a theory with which readers of this journal will already be very familiar: Keynes's trade cycle. In the process, I will endeavor to make SD understandable to the reader and highlight its many advantages.

History of system dynamics

The primary figure in the history of SD is Jay Forrester, an engineer by training but a professor of management at MIT.[2] While he started his work in the area of servomechanisms for military equipment during World War II, a series of chance events led him to the application of the principles learned there to business and eventually social problems. Implicit in his approach is the assumption that real-world social processes are governed by complex feedback mechanisms and that we cannot truly understand their character until we observe the operation of all the component parts simultaneously and in interaction with one another. The human mind, Forrester contends, is unable to manage this:

> The human mind is not adapted to interpreting how social systems behave. Social systems belong to the class called multi-loop nonlinear feedback systems. In the long history of evolution it has not been necessary until

[1] Michael Radzicki (1988, 2008) has been at the forefront of this effort, and Eric Tymoigne (2006) has also been a strong proponent.

[2] This section is based heavily on Forrester (2007).

> very recent historical times for people to understand complex feedback systems. Evolutionary processes have not given us the mental ability to interpret properly the dynamic behavior of those complex systems in which we are now imbedded. (Forrester, 1995, p. 3)

Thus, he claims that in the absence of computer simulation we are prone to recommend policies that are at best ineffective and at worst counterproductive.

His first major work based on these premises was *Industrial Dynamics* (Forrester, 1961), which examined the character of the processes contained within business firms (pricing, hiring, inventory adjustment, etc.). Instrumental in the development of that volume was his study of the General Electric plant in Kentucky, which led to the development of computer software and probably the most famous of his simulations, the beer game.[3] However, a bigger—and largely negative—splash was made by *Urban Dynamics* (Forrester, 1969). His first foray into social issues, it was the result of hours upon hours of interviews and painstaking research into the specific processes that were driving the deteriorating conditions in American inner cities. Among his conclusions was that the "obvious" solution of adding low-cost housing to these poverty-stricken areas was actually exacerbating the situation. This was so, he said, because the real problem was the lack of economic opportunity. The houses would build themselves if the local population had the income to afford them. Instead, the existing policies were attracting even greater numbers of people to an area that was already beyond its capacity to support them. To Forrester's surprise, this was met with often hostile reactions, particularly by those who had devoted considerable time and resources to working on solutions based on one-dimensional views of the issue. But, he remained convinced of the utility—indeed, the necessity—of his approach.

The reaction to his next major work, *World Dynamics* (Forrester, 1971), was again unexpected. A volume outlining the dangers of unrestrained resource use, Forrester was fully prepared for it to be totally ignored by the general public (particularly given the rather technical nature of much of the volume). On the contrary, it soon became the topic of editorials, major book reviews, and even a television documentary. A follow-up volume based on an extended version of his model but with a more accessible approach was published soon thereafter (Meadows et al., 1972). This only served to increase the popularity of his work and add to the credibility of SD modeling.

[3] The latter can be viewed and manipulated here: http://beergame.mit.edu/.

Forrester and his students soon began applying SD to a variety of topics, aided by the increasing sophistication of computer hardware and the efforts of those developing dedicated software programs. These included economics, wherein an explanation of the business cycle based solely on the time lag involved in inventories was developed.[4] Today, applications range from Shakespeare to physics, there is a journal (*System Dynamics Review*) and an annual conference, and one can even earn degrees and certificates in SD modeling.

Despite all this, applications to economics have been relatively few. Searching EconLit for "system dynamics" yields only 279 results for the past almost forty years! In addition, the majority of these are either in management or focused on natural resource issues. The outstanding exceptions are associated with heterodox journals and policy centers, and even these represent no more than a handful. This is not surprising given that mainstream economics is dominated by a paradigm that assumes that market systems are static, unicausal, moving toward resting points, and affected primarily by exogenous factors. This is not the world described by SD, nor is it that envisioned by post Keynesians.

Building a system dynamics model

SD is a computer simulation technique based on the premises that social systems are marked by complex feedback mechanisms. When one undertakes SD modeling, the first step is often the building of a boundary chart that the researcher uses to decide which concepts will be endogenous, exogenous, and excluded.[5] Obviously (and similar to any modeling process), what falls into each category depends as much on the purpose of the study as the phenomenon being examined, as does the level of detail involved. One must also consider at this stage the time horizon to be examined and the references that will be employed in considering the behavior of the variables.

Roughly speaking, the next step is the construction of maps or schematics of the causal structure of the phenomenon, noting in particular the direction of each link (i.e., positive or negative). This will be informed

[4] Although obviously pre-Forrester and SD, Keynes makes reference to such a possibility (Keynes, 1964, p. 322).

[5] John Sterman (2000) offers a comprehensive, textbook-style introduction to SD, including software and sample models. This explanation of how to set about an SD study is based on his chapter 3, "The Modeling Process." For those interested in learning SD, I would recommend first reading Radzicki (1988) and then Sterman (2000).

by a variety of sources: intuition, econometrics, statistics, existing literature, theory, and so on. With each, there are potential strengths and weaknesses and there is no universal guide as to which is best. Also, in specifying the relationships, SD assumes a post Keynesian-style concern with the real world. It is not the object of the SD researcher to build as-if models, but, as much as possible, to create a final product that reflects actual causal processes. If we have not done that, then how can we truly know which are the dominant lines of causation and where policy can be applied successfully?

Even before the above step is complete, the specification of the formal simulation model begins. This is so because it is recommended that one not try to build the entire system before any test runs. Errors in a 100-equation model will be much more difficult to diagnose and correct than those in a 10-equation model. Hence, some core process will be selected first and, if this appears to work, more details and structures are added. Important at this stage is consideration of which variables should be treated as stocks and which as flows, where stocks describe the state of the system and flows add to and subtract from those stocks (something SD has in common with the approach taken in Godley and Lavoie [2007]). Auxiliary variables are used to represent other factors (e.g., constants or exogenous forces).

Once a complete model of the system has been built, various tests are administered including comparison to the references that had been selected (does my model of the 1930s show the United States with appropriately high levels of unemployment?), robustness under extreme conditions (what does the banking sector do if interest rates are 500 percent?), and sensitivity to parameters and time horizon (is it critical whether the marginal propensity to consume is 0.7 or 0.8 and do I get markedly different behavior when steps are one-half month as opposed to one month?). Note that one is not operating under the assumption that model behavior *should* remain the same under all circumstances. One of the things that is discovered by this process is which changes are particularly significant relative to others. This is important because it reveals leverage points in the system, those key sectors whose structure dominates overall behavior. These are the places where it may be possible to apply policy most effectively. Once testing is complete, more runs can be carried out to see how the system behaves in various circumstances and other competing policy regimes. This obviously shares some factors in common with the sensitivity testing, but the goal is different.

The entire model-building process is highly iterative, with the researcher jumping back and forth between the various steps (even all the way back to the boundary chart) as more is learned about the system. This is to be expected and is one of the strengths of SD modeling because it means that one learns a great deal about the phenomenon simply as a consequence of creating the simulation. It is very easy to say, for example, that investment is a function of expected profits and the rate of interest, and that statement is not without significance. However, when one is forced not only to define that relationship mathematically, but to then place it into a context where it must interact over time with a larger economy containing various feedback structures, one suddenly becomes aware of the many other factors that play important roles in the real world. In addition, SD teaches that very few relationships are actually linear. Consequently, the model below will make frequent use of s-shaped functions wherein the effect of an increase in the independent variable on the dependent will vary by the magnitude of the former.[6]

SD analysis also suggests that there exists a set of fundamental modes of dynamic behavior, including exponential growth, goal seeking, s-shaped growth, oscillation, growth with overshoot, and overshoot and collapse (Sterman, 2000, pp. 107–133). Because each is related to a specific feedback structure, observing one of these modes of behavior in the real world can serve as an important guide to how it should be modeled. The business cycle, for example, is very similar to oscillation, which is created by a negative feedback loop in the presence of a time delay. This is not unlike the final structure of the model below, where several negative feedback loops are created by the processes identified by Keynes but investors react (overreact, in fact) to falling profits only after some quarters have passed.

With respect to the loops, as suggested above, they represent the key form in which relationships are specified in SD. This is not something the modeler tries to force onto the system, but is a result of the manner in which the real world operates. Feedback is the rule rather than the exception, and this, by definition, takes place over time. It can show how agents react to mistakes, change their plans in light of realized results, and so on. In other words, it allows the past to have a significant effect on the future.

In summary, SD is an ideal modeling technique for post Keynesians. It is set in time, focuses on real-world causal mechanisms, and is holistic. In addition, there exists a wealth of literature from which to draw, including many beginner guides. And there is a dedicated software package, the basic version of which can be downloaded free of charge

[6] Similar to the approach taken in Kalecki (1935), Kaldor (1940), and Hicks (1950).

for academic use (www.vensim.com/venple.html). Not every problem is best addressed by SD, but many are—particularly those analyzing systemic events over time.

An SD model of Keynes's chapter 22

To better illustrate how SD modeling works and to further show how it can be applied to post Keynesian theory, this section develops a simulation of the trade cycle outlined in chapter 22 of Keynes's *General Theory*. It is well suited because it is a process that creates not a static equilibrium, but a pattern through time. While one might normally start specifying the boundary chart first, in this case those decisions must be guided by how Keynes had approached the issue. Hence, it will wait until the next section.

On the first page of chapter 22, Keynes suggests that understanding the business cycle requires tracking a number of processes simultaneously, with one playing a dominant role:

> fluctuations in the propensity to consume, in the state of liquidity-preference, and in the marginal efficiency of capital have all played a part. But I suggest that the essential character of the Trade Cycle, and, especially, the regularity of time-sequence and of duration which justifies us in calling it a cycle, *is mainly due to the way in which the marginal efficiency of capital fluctuates.* (Keynes, 1964, p. 313; emphasis added)

The starting point for Keynes's argument is his belief that the economy's core problem is its inability to generate sufficient physical investment spending to offset saving. Over the expansion, the two primary determinants of investment, the rate of interest (r) and the marginal efficiency of capital (MEC), both move in directions that progressively reduce the incentive to add to the stock of physical capital. The former does so because the demand for financing will increase and at the point of the collapse, there may be a rush to liquidity as agents panic (Keynes, 1964, p. 316). Both of these drive interest rates higher. Meanwhile, the MEC necessarily declines over the expansion because it "depends, not only on the existing abundance or scarcity of capital-goods and the current cost of production of capital-goods, but also on current expectations as to the future yield of capital-goods" (Keynes, 1964, p. 315). Rising capital costs also play a role as bottlenecks emerge and the prices of various raw materials increase. These two factors, alone, are sufficient to create a systemic problem. But, as Victoria Chick suggests, there is more: "It is the special contribution of the financial side to transform a gradual downturn into a sharp crisis" (Chick, 1983, p. 289). That "special contribution" is a function of Keynes's "current expectations as to the future yield of capital-goods," and it represents the third and last determinant of the MEC. It works as follows.

While animal spirits allow agents to overcome the uncertainty of the real world sufficiently to engage in significant forward-looking activity, it is also true that such spontaneous optimism can lead to overreaction, disappointment, and collapse. This is especially likely when forecasts in the market for financial assets, which inevitably affect the expectations of profit in the market for physical capital, are dominated by those with a very short time horizon and who "do not manage and have no special knowledge of the circumstances, either actual or prospective, of the business in question" (Keynes, 1964, p. 153). In that case, the MEC can become unrealistically inflated during the expansion and as the gap between expectations and reality grows, disappointment is inevitable. Overreaction may occur in the sense that the earlier overly optimistic view "is replaced by a contrary 'error of pessimism'" (Keynes, 1964, pp. 321–322). The recession ends when the depreciation of capital and the fading memory of the collapse are sufficient to raise the MEC once again. Then the cycle starts over.

A simple causal map of Keynes's theory is shown in Figure 1. This is an important step as it will be used as the blueprint for the SD simulation. Investment (I) is a positive function of the MEC and a negative function of interest, with the former shown as a function of the three variables described above: stock of capital (K), cost of capital (cost of K), and speculative forecasts from the financial sector (speculative profit expectations). Both the stock and cost of capital are modeled as positive functions of the level of investment, while interest rates are modeled as a positive function of financing demand (which rises as investment rises) and the hoarding component of liquidity preference (which Keynes argued would spike with the fall in MEC at the start of the recession [Keynes, 1964, p. 316]). Since hoarding and the MEC are inversely related, the sign is negative.

Despite the multiplicity of obstacles to raising investment already evident by the number of negative feedback loops in Figure 1, Keynes points to the influence of the divergence of expectations from reality as the real key to crisis. As the economy expands and the contrast with the earlier recession is fresh in people's minds, animal spirits are encouraged (see "speculative profit expectations"). Consecutive quarters of high (if eventually moderating) realized profits lead agents to believe that the consequences of any rise in K, cost of K, or financing demand are trivial compared to the returns promised by the booming economy (Keynes, 1964, p. 315). Hence, the gap between realized profit from investment and

Figure 1 Causal map of Keynes's trade cycle

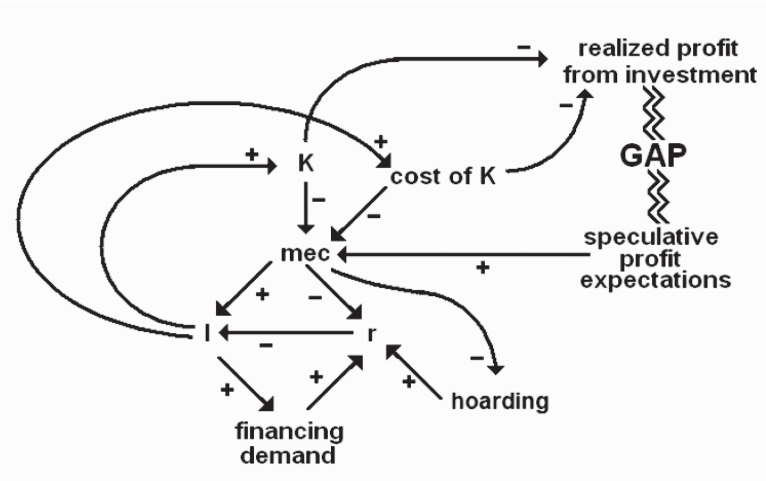

speculative profit expectations increases over time. Inevitably, the degree by which reality disappoints agents becomes so great that they despair, and their overly optimistic view "is replaced by a contrary 'error of pessimism'" (Keynes, 1964, pp. 321–322). Collapse and recession follow. Note that this will occur even if the realized profits would have been considered reasonable in the absence of the excessive forecast. The economy now stagnates until time erases the feeling of shock and the stock of physical capital depreciates sufficiently to buoy expectations once again.

One last important note is that Keynes is not arguing that this is a story of overinvestment. On the contrary, the crisis occurs in a manner that prevents the market for capital from truly being saturated in the sense that the community has all it wants or needs (Keynes, 1964, pp. 320–321).[7] Hence, the standard mainstream story does not apply. The collapse comes before the profit motive has provided sufficient capital to meet the level of demand that would have been forthcoming at full employment, not because it had already done so.

The complete SD model shown in Figure 2 is adapted from the causal map in Figure 1.[8] The central variable is the *marginal efficiency of capital,*

[7] Keynes does not deny that certain classes of investment may be overbuilt, only that this is not characteristic of the economy as a whole. He also agrees that inventory adjustment can cause "minor oscillations within the main movement of the Trade Cycle" (Keynes, 1964, p. 322).

[8] The complete model can be downloaded at www.econ.tcu.edu/harvey/Keynes_Business_Cycle.htm.

Figure 2 Complete SD Model of Keynes's Trade Cycle

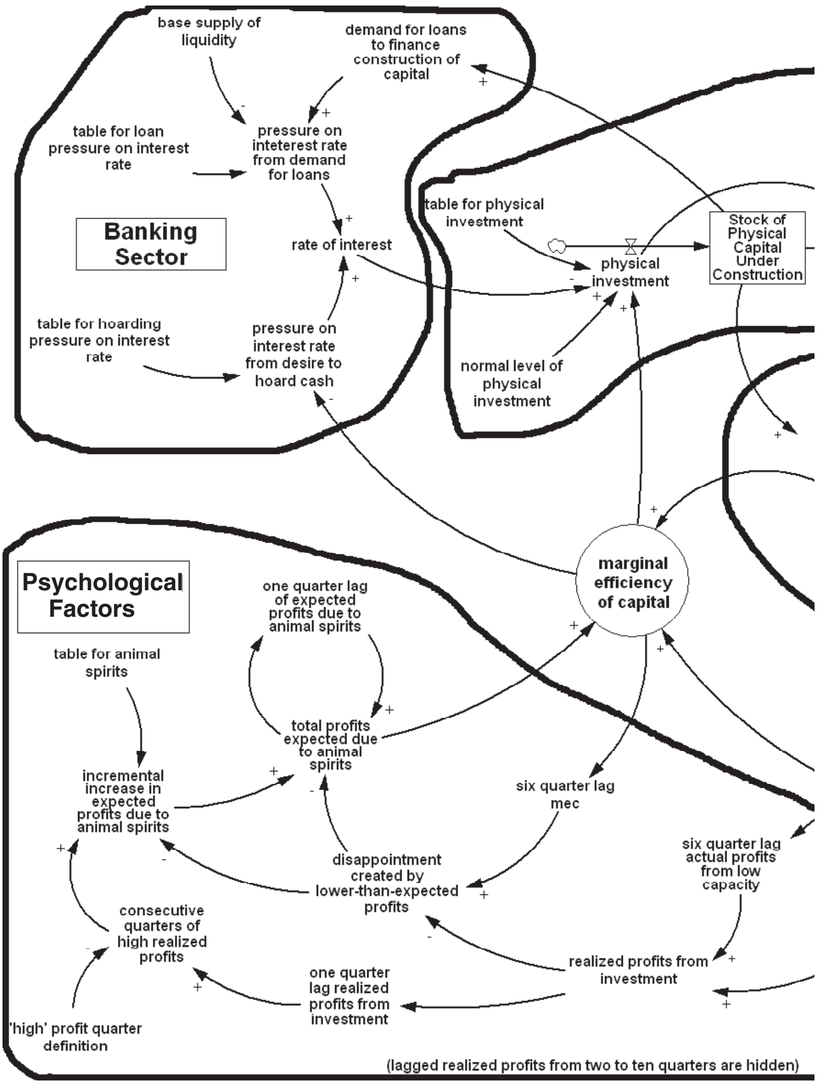

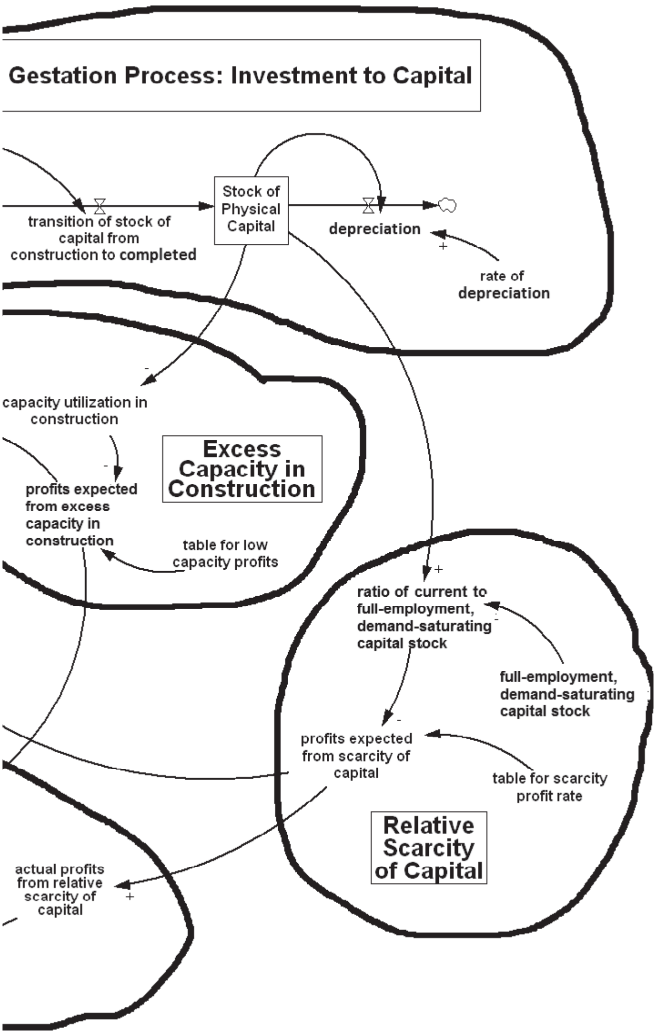
Gestation Process: Investment to Capital
Stock of Physical Capital
transition of stock of capital from construction to completed
depreciation
rate of depreciation
capacity utilization in construction
Excess Capacity in Construction
profits expected from excess capacity in construction
table for low capacity profits
ratio of current to full-employment, demand-saturating capital stock
full-employment, demand-saturating capital stock
profits expected from scarcity of capital
table for scarcity profit rate
Relative Scarcity of Capital
actual profits from relative scarcity of capital

with the rest of the economy organized into five parts: gestation process, excess capacity in construction, relative scarcity of capital, psychological factors, and banking sector (hereafter, when reference is made to a specific variable from the SD model, its name will be italicized). Each sector will be explained in turn below. In terms of the base range of values in the system, whenever a guide was available from Keynes, this was followed. Otherwise, current economic data were consulted and extrapolations made. For purposes of this analysis, an economy roughly the size of a midrange European country (e.g., Russia or Italy) was assumed. The model is quarterly and, assuming a business cycle that runs five to seven years from peak to peak, it was simulated for 100 periods (twenty-five years) so as to generate several repetitions of the pattern (since World War II, the United States has averaged 5.5 years from peak to peak).

Table 1 shows the boundary chart. In building it, the researcher attempts, as in any model-building process, to determine the minimum detail required to realistically portray the phenomenon in question. In this case, Keynes's approach in chapter 22 was especially important as the goal is to re-create his model. The variables listed as endogenous will fluctuate in response to changes within the system and will create the dynamic behavior. This is not to say that those marked as exogenous are less important. Indeed, checking how patterns change when these values are altered is an important step in the modeling process. Those listed as excluded, however, are not considered at all. This means, for example, that there is no government sector or international trade. Nor is employment formally represented, though it is assumed that it would vary directly with changes in investment. In fact, as the latter is taken to be the primary driver of consumption and gross domestic product (GDP), too, they are also excluded. Strictly speaking, there is some representation of the next item on the list, the financial sector, because the interest rate and speculative expectations are modeled. But, as there is no attempt here to build what post Keynesians would consider to be a full-scale financial market, it is listed as excluded. In addition, agent debt levels are not considered, even though they might otherwise be relevant in this context. The last important piece of information supplied by the boundary chart is that it assumes a stationary state in the sense that productivity levels do not change. This means that to achieve full employment, it is not necessary to sustain a particular rate of growth of GDP (as a result of a particular rate of growth of investment), only a given level. This greatly simplified the modeling process.

Investment is the key variable in Keynes's model of the trade cycle and so it will be explained first. Recalling that SD conceptualizes problems

Table 1
Boundary chart for system dynamics model of Keynes's trade cycle theory

Endogenous	Exogenous	Excluded
Investment	Rate of depreciation	Government
Stock of capital	Lag time for investment to become completed capital	Foreign sector
Target stock of capital	"High profit quarter definition	Employment
Marginal efficiency of capital	Gap between expected and actual profits necessary to induce panic	Consumption
Rate of interest	Full-employment, demand-saturating of capital stock	GDP
Cost of capital		Financial market
Profits from scarcity of capital		Agent debt levels
Speculative profit expectations		Productivity growth
Realized profits		

in terms of stocks and flows, investment clearly suggests the latter while the stock of capital is the former. Figure 3 illustrates these relationships. In terms of the time required for an investment project to become productive capital, empirical studies suggest that this averages eighteen months.[9] Therefore, *stock of physical capital under construction* is a level containing the previous six quarters of investment and representative of all current investment projects at various levels of completion. That level is reduced by *transition of stock of capital from construction to completed,* which is a six-quarter lag of *physical investment*. In this manner, in six quarters a project goes through the chain from inception at *physical investment,* to the construction phase at *transition of stock of capital from construction to completed,* to an active part of the economy's productive capital at *stock of physical capital*.

Note that the model-building process is already forcing us to address issues that Keynes never raised. The researcher must use judgment in deciding how to close any loopholes, but this is not unwelcome. It is all a part of learning which parameters and structures—whether implicit or explicit in the original—are key. And if it turns out to be impossible to

[9] Sterman (2000) cites Montgomery (1995) as suggesting that lags should be around 1.4 years.

Figure 3 Gestation process, from investment to the stock of capital (and depreciation)

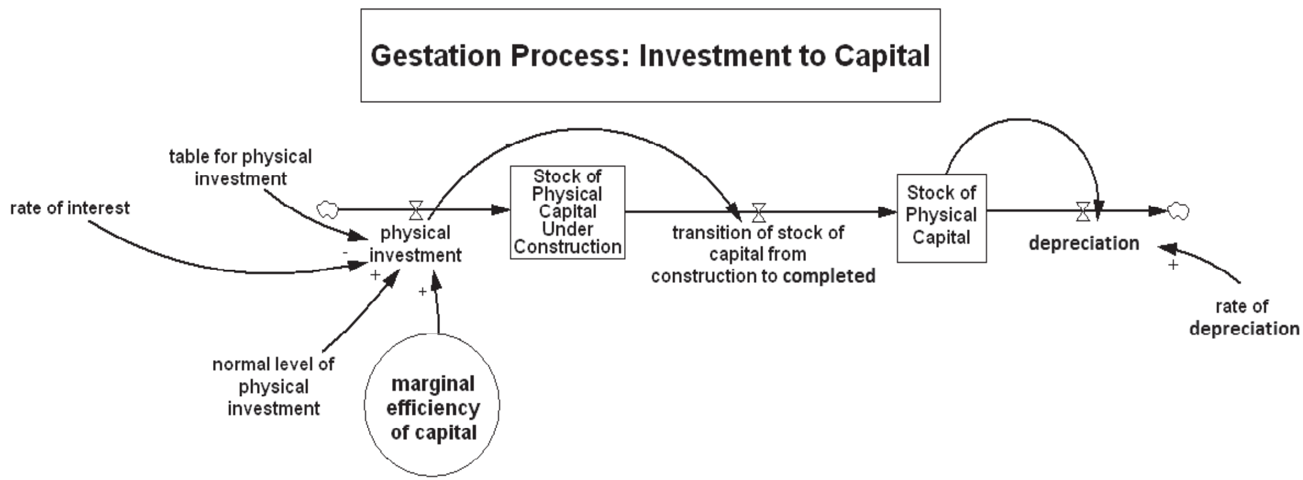

build an SD model of a theory once reasonable replacements for the missing pieces are included, then this could be an indicator that the original theory was not internally consistent.

Continuing the specification of the gestation process, *depreciation* removes from the pool capital what has worn out or become outdated (the primary force allowing the economy to recover from recession). The rate at which this occurs (set at 5 percent for the base run) is determined by *rate of depreciation*.[10] The main determinant of *physical investment* is the *marginal efficiency of capital*. When compared to the prevailing *rate of interest* via a table (where the latter assumes an s-shaped relationship with a range of values based on that witnessed in the U.S. economy over the past twenty years), a multiplier is generated that, when combined with the normal level of physical investment, determines the current level.

The marginal efficiency of capital itself is divided into three components: profits expected from excess capacity in construction, profits expected from scarcity of capital, and total profits expected due to animal spirits. These correspond to what was shown in the causal map in Figure 1 (where they were labeled as cost of K, K, and speculative profit expectations). To simplify the modeling process, each is defined so that it generates a positive number, the sum of which gives the MEC. Thus, for example, what Keynes called the cost of capital, which would have been negatively related to profits, is modeled as its inverse (the cheapness of capital, or the profits expected from excess capacity in construction) so that it becomes positively related. For that variable, the assumption was made that as construction on investment projects increased, bottlenecks would emerge. For simplicity, it was assumed that the ability of that industry to meet demand would be a direct function of overall capacity in the macroeconomy (i.e., capacity in construction would be a function of overall capacity). Figure 4 shows this relationship. Capacity utilization in construction is equal to the stock of physical capital under construction (i.e., the demand for construction) divided by the stock of physical capital (the ability of the construction industry to meet demand). To generate its contribution to expected profits, capacity utilization in construction is compared to an s-shaped table that then generates values ranging from 0 (at the lowest level of excess capacity) to 3.5 percent (at the highest level of excess capacity). It is important to note here that in the *General Theory,* there is no implication that this

[10] The magnitudes for the various capital stock and investment numbers were derived from British studies found at www.statistics.gov.uk/StatBase/Product.asp?vlnk=10730/.

Figure 4 Keynes's cost of capital determinant of the marginal efficiency of capital

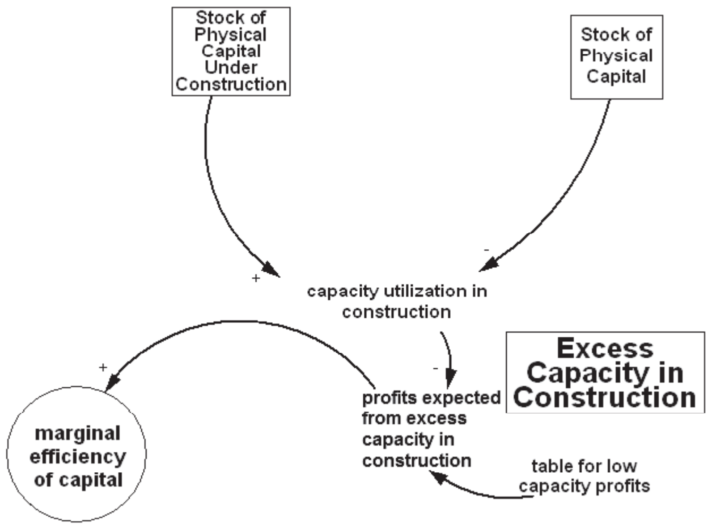

forecast might turn out to be wrong. Rather, agents know full well that the cost of capital is increasing; but in the euphoria generated by animal spirits they somehow expect future economic conditions to offset this (Keynes, 1964, p. 315).

The second component of marginal efficiency of capital is that based on relative capital scarcity, shown in Figure 5. Recall that Keynes believed that the more abundant capital is, the lower the MEC because the profitable investments will have been undertaken first and each successive addition to the stock of capital faces more and more competition. The basic measure of scarcity employed in this model is the ratio of the existing *stock of physical capital* with the level necessary to saturate demand under conditions of full employment (*full-employment, demand-saturating capital stock*). This number can vary from 0 (when there is no existing stock of capital) to 1 (when the stock of capital is equal to the full employment, demand-saturating level), and a straight-line function was assumed, with 0 being associated with a profit rate of 15 percent and 1 a profit rate of 0 percent. This number is then added to the *marginal efficiency of capital*. Keynes's view of the accuracy of this component of

Figure 5 Profits expected to be generated by the scarcity of capital

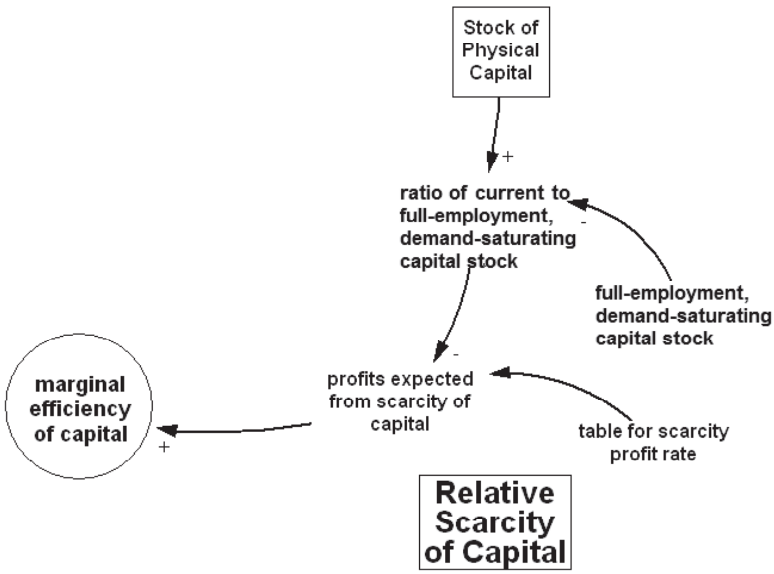

the forecast is more difficult to interpret. Since he offers no insights, the simple assumption that agents expect today's scarcity profits to continue indefinitely is adopted. In this way, the *profits expected from scarcity of capital* variable is used as the forecast of what will prevail in the future and it is the actual rate of contemporaneous scarcity profits. This will be useful later when a realized profits variable must be specified.

The last contributor to *marginal efficiency of capital,* profit expected by financial speculators, is the one Keynes claims is the most responsible for instability. While profits from the relative scarcity of capital and those from excess capacity in construction are rightly expected to decline over the course of expansion, agents, buoyed by the above-average returns, nevertheless believe that these negatives will be more than offset by the booming economy. Capturing this proved to be the most challenging part of the modeling process, the result of which is shown in Figure 6, "Psychological Factors."

Three things must occur: a cumulative increase in optimism as the expansion generates above-average profits, an eventual decline or lagging growth in realized returns over that same period, and panic when the gap between reality and expectations becomes too large to ignore.

Figure 6 Effect of animal spirits on the marginal efficiency of capital

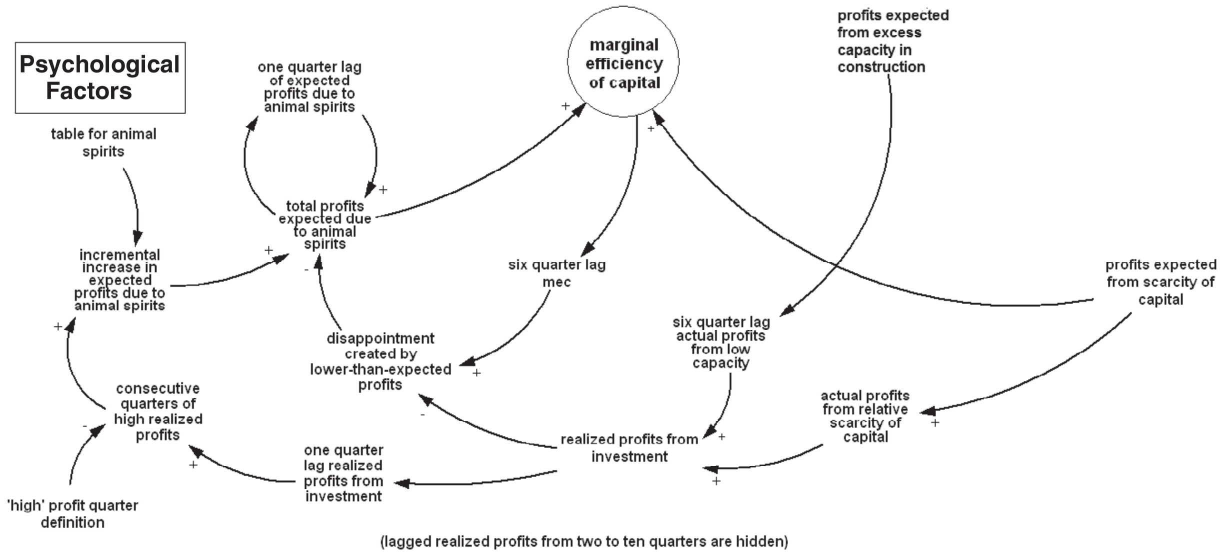

Showing this required the creation of a variable that reflected actual rates of return on investment, which, for reasons explained above, could be accomplished by combining *profits expected from excess capacity in construction* and *profits expected from scarcity capital* (this is also how it is shown in Figure 1; note that the former must be lagged six quarters so that it corresponds to costs when produced). With such a variable in place, it is then a relatively simple (if tedious, given all the lags that must be created) matter to create a counter to keep track of the number of consecutive quarters profits were "high." This variable used a set of nested if-then statements and lagged values of *realized profits from investment* (for clarity, all but one are hidden in both Figure 2 and Figure 6). The subjective value for "high" is defined in *"high" profit quarter definition,* and for the standard run was set at 4 percent (a number selected on the basis of Keynes's comments on page 321 of the *General Theory*).

Consecutive quarters of high realized profits (i.e., profit over 4 percent) feeds into *incremental increase in expected profits due to animal spirits* via an s-shaped table that generates a bonus to the profit forecast varying from 0 to 0.6 percentage points. This represents the quarterly rise in the overreaction of financial investors whose optimism is buoyed by recent events. This accumulates in *total profits expected due to animal spirits,* which includes all previous rises (except in the case of disappointment, which resets it to zero—more on that in a moment) by adding its lagged value. For example, if *incremental increase in expected profits due to animal spirits* generated 0.2, 0.3, and 0.4 over three quarters, then *total profits expected due to animal spirits,* assuming we had started at zero, would be 0.2, 0.5, and 0.9. None of this is in Keynes, of course, but it is consistent with his analysis.

This is then added to the *marginal efficiency of capital* as its third and last component. While the other two, *profits expected from excess capacity in construction* and *profits expected from scarcity of capital,* rightly predict that as the expansion continues returns must fall, the psychological component says just the opposite. Agents' animal spirits, without which Keynes argues there would be no investment at all, lead them to believe that good times now guarantee even better times in the future. This cannot be true, and looming below animal spirits in Figure 6 is disappointment.

Disappointment created by lower-than-expected profits compares the *marginal efficiency of capital* when the investment was undertaken (six quarters ago) with conditions when the capital project was completed (*realized profits from investment* in the current quarter). Using magnitudes suggested by Keynes's discussion in chapter 22 of the *General Theory,* it uses an if-then statement to generate a zero whenever actual profits fall short of expected by more than two percentage points and a one otherwise. Because *total profits expected due to animal spirits* is multiplied by *disappointment created by lower-than-expected profits*, and because the latter is zero whenever a project falls significantly short of what was forecast, this means that the psychological component of the *marginal efficiency of capital* suddenly collapses to zero whenever disappointment occurs. This will, of course, happen at the height of an expansion, and is sufficient (given the parameters I have assumed) to send the economy into recession. After that, we must wait for *depreciation* to reduce the current *stock of physical capital* sufficiently to raise hopes (and then investment), and the whole process starts over again.

This completes the specification of all the subsectors in the model except the one governing the interest rate. While Keynes clearly states that he sees this as a secondary factor, it can nevertheless have an effect over the business cycle and should be included particularly because so many mainstream economists view it as the key factor. There are always conflicting pressures at work in the banking sector. While during the upswing, the demand for financing may be high as new projects are undertaken, people are nevertheless confident and thus the hoarding of cash is low. Over the downswing, those positions reverse. To take these two distinct pressures into account, Figure 7 shows the rate of interest as the sum of the *pressure on interest rate from demand for loans* and the *pressure on the interest rate from desire to hoard*. The former relies on *demand for loans to finance construction of capital* (which is actually identical to the *stock of physical capital under construction,* but the effect on the banking sector is made more clear by renaming it) and compares this to the *base supply of liquidity* in the macroeconomy via a table that assumes that financing costs increase at an increasing rate. Note that this does not indicate that the model assumes exogenous money supplies. Any level of demand can be accommodated, but, ceteris paribus, the price of doing so (i.e., the interest rate) varies. Meanwhile, hoarding (taking a cue from Keynes) is an inverse, s-shaped function of the *marginal efficiency of capital*. The assumption was made that this factor is relatively less important than demand for loans and this is reflected in the values on the tables. However, given that it is inversely related to the *marginal efficiency of capital,* increasing the impact of hoarding would only increase the severity of the cycle predicted by Keynes, not eliminate it. Note that this sector was kept as simple as possible while re-creating the behavior described by Keynes. Further note that this is only the banking part of the

Figure 7 Determination of the rate of interest

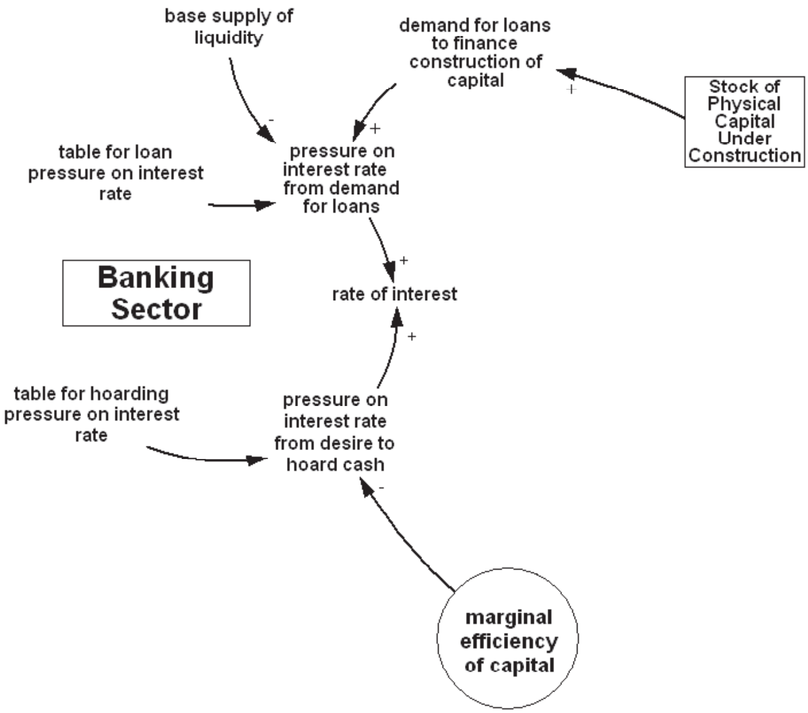

financial sector. The market for financial assets is not modeled, although its effect on MEC is reflected in the animal spirits generated in the segment of the diagram labeled "Psychological Factors."

Validation, sensitivity tests, and analysis

With the model complete, various tests need to be undertaken. First, in terms of basic behavior, the system structure is very similar to the theory laid out by Keynes and it generates a business cycle of the sort described in chapter 22 of the *General Theory*. Expansions are marked by an initially gradual then steep increase in investment, which is followed by a sudden collapse and a weak recovery. This is illustrated in Figure 8. Furthermore, the key variables fall into the expected ranges and—as predicted by Keynes—the stock of capital never reaches the full-employment saturation level. If we assume that the full-employment

Figure 8 Pattern created by investment in the system dynamics model of Keynes's theory of the business cycle

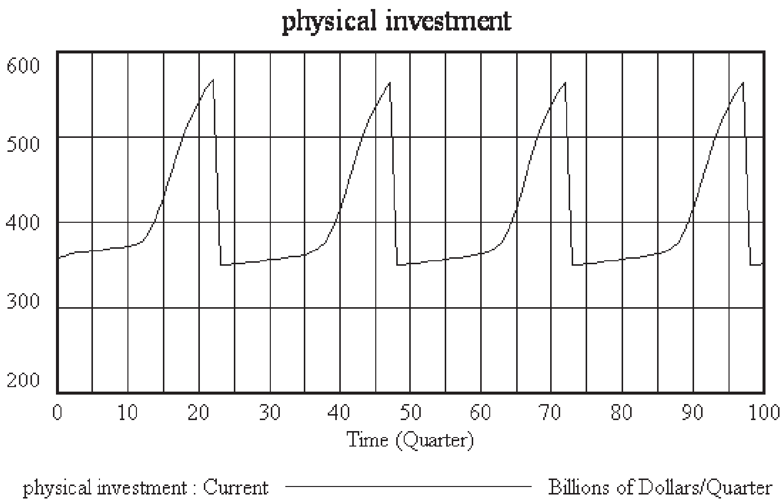

level of investment is around $500–550 billion, then Figure 8 shows an economy that can occasionally reach this level, only to collapse a short time later. The cycle here runs 6.5 years from peak to peak. Although this is one year longer than the United States since World War II, this economy lacks the automatic stabilizer created by the government's budget. Long slumps at lower than full employment are the rule. Note that, of course, there is no expectation that the periodicity, amplitude, frequency, or wavelength would be constant or that the parameters would remain stable. The goal here was only to build a representative business cycle from Keynes's chapter 22.

In terms of sensitivity tests, all five of the variables listed as exogenous in Table 1 were manipulated to see if significant changes in model behavior took place. Different values for the gap between expected and actual profits necessary to induce panic, the lag time for investment to become completed capital, *"high" profit quarter definition,* and *full-employment, demand-saturating capital stock* affected only the periodicity of the patterns and not their basic structure. The specification of *rate of depreciation,* however, proved to be very important. Increasing the rate only changed periodicity, but lower ones led to a rapid decline in the profitability of investment and therefore a collapse. Investment found an equilibrium at which new construction was exactly offset by depreciation and, because this was inevitably at a low level of investment, the economy became trapped in a never-ending slump. This fits Keynes's concerns well and indicates that increases in the longevity of physical capital, while obviously advantageous from a technological perspective, create even more problems for the economy. Integration-error tests were also conducted, but no problems were revealed (these are undertaken to ensure that the model behavior was not dependent on the time-step selected).

Key to understanding the structure of the model is to determine the effect of each of the four sectors in Figure 2 that affect investment (the fifth sector) and, therefore, the business cycle. Keynes clearly indicates, for example, that he sees interest rates as a potentially complicating but not primary factor. To test this, separate simulations were carried out while holding the effects of four of the major subsectors—banking sector, excess capacity in construction, relative scarcity of capital, and psychological factors—held constant. This was accomplished by replacing that sector's contribution with a constant equal to its average in the base run of the simulation.

For example, one run was made omitting *rate of interest*. While it remained in the investment equation as a determinant, it was set at 2.96 (its average value in the base run). This had an imperceptible effect, consistent with Keynes's argument (and contrary to what many economists argue today). The same was true when holding *profits expected from excess capacity in construction* and *profits expected from scarcity of capital* constant (at 1.12 and 2.91)—periodicity, time from peak to trough, and maximums and minimums changed slightly, but the basic patterns were the same. This was not the case with *total profits expected due to animal spirits,* however, as the business cycle was completely eliminated when that factor was replaced with a constant 1.09 in the equation for *marginal efficiency of capital*. In addition, instead of the rapid collapse of investment to an equilibrium level as was experienced in the sensitivity testing of the depreciation of the capital stock above, it rose. It nevertheless left the stock of capital well short of desired levels, but was otherwise preferable.

Policy lessons

The policy implications of this analysis are consistent with Keynes's argument. First, it is unclear that interest rates can significantly affect the process. Mainstream economists believe that expansions should be dampened by raising interest rates. In this way, the "overinvestment" and overconfidence that lead to catastrophic collapse can be avoided. Keynes cautions, however, that not only are interest rates a relatively weak tool in general, but that overinvestment in a realistic sense never occurs. This is supported both by the results of the sensitivity tests and by a fresh run of the simulation conducted specifically to put this expansion-dampening theory to the test. In it, interest rates were raised whenever the *marginal efficiency of capital* went above 4 percent by adding to the original rate (already determined by the demand for hoarding and loans) the *marginal efficiency of capital* minus 4 percent. In this way, overexcitement would be exactly offset by a jump in the cost of financing projects. While there is no question that this succeeded in keeping investment from fluctuating as wildly (the standard deviation falls from $71 billion in the unadjusted simulation to only $5 billion in the expansion-dampening one), it also kept it lower on average (declining from a mean of $407 billion to $369 billion) and caused the stock of physical capital, after an initial surge, to go into a mild long-run decline. This is hardly a solution to economic crises because, as Keynes warned in the *General Theory,* such a method for "abolishing booms" keeps "us permanently in a semi-slump" (Keynes, 1964, p. 322).

Instead, the real problem is the market for financial assets and the mindset that it creates. Uncertainty, animal spirits, and the ability to divest oneself of a financial asset within moments of acquiring it combine to create what amounts to a legitimized casino. People ignorant of the fact that profits cannot continue their upward track not only bet otherwise, but they influence (and in some cases are) those undertaking physical investment. When disappointment inevitably arrives, it can do so catastrophically. The sensitivity tests above already showed the effect of removing speculative expectations from the equation: average investment falls only slightly to $397 billion (just under the $403 generated by the unadjusted model) and its standard deviation is lower even than that created by the expansion-dampening policy ($2.6 billion rather than $5 billion). This model suggests, consistent with Keynes's analysis, that reducing the liquidity of financial investments and, therefore, the effect of speculation and animal spirits is key.

However, this still leaves the problem of the overall stock of capital. On average, while under the speculation-constrained scenario it is higher than with an expansion-dampening policy ($7,975 billion versus $7,559 billion), it is actually lower than in the unadjusted scenario ($8,061 billion). This is so because one positive side effect of the animal spirits-fueled speculation was to raise investment (and therefore capital) longer than cool-headed, profit-maximizing entrepreneurs might have done. This points to a key failure of our current system: to push the stock of physical capital to the point that it is equal to the full-employment, demand-saturating level, it would be necessary for firms to continue to invest *up to* the point that they were expecting zero profit. So long as there is a positive rate of interest, that is impossible. This is consistent with Keynes's primary policy recommendation that, along with interest rates being kept as low as possible at all times and speculative influences in financial markets being reduced, the process of investment should be socialized in the sense that the government take over the responsibility of pushing the stock of capital to the saturation point. Only in this way, he believed, could we experience both sufficient investment to create full employment and a stock of physical capital that would satisfy current demand. All of this is supported by the SD model developed here.

Conclusions

In chapter 21 of the *General Theory,* Keynes wrote: "The object of our analysis is, not to provide a machine, or method of blind manipulation, which will furnish an infallible answer, but to provide ourselves with an organised and orderly method of thinking out particular problems" (Keynes, 1964, p. 297). Those who are supporters of SD have not deluded themselves into thinking that it is a machine that furnishes infallible answers. It is, instead, simply another "organised and orderly method of thinking out particular problems." As such, it is as strong or weak as the premises upon which it is built. But, what makes it unique and particularly valuable to post Keynesians is its ability to show processes through time and to allow the modeling of whole systems at once. Indeed, the latter reduces the need to isolate "the complicating factors one by one," which is important because processes operating side by side may produce behaviors that differ from what we observe studying them in isolation.

The basic model here is useful not only as a means of analyzing and understanding Keynes's trade cycle theory, but it could be used as a core to which could be added a more sophisticated financial sector, international trade and capital flows, a central bank, export-constrained growth, and so on. Equilibrium models can be useful in limited contexts, but we should bear in mind that they are the preferred method of a school of thought that does not really understand how modern, capitalist economies work. SD is a much tighter and more logical fit with post Keynesianism.

References

Blatt, J. "Economic Policy and Endogenous Cycles." *Journal of Post Keynesian Economics,* 1983, *5* (4) (Summer): 635–647.

Chick, V. *Macroeconomics After Keynes: A Reconsideration of the General Theory*. Cambridge, MA: MIT Press, 1983.

Forrester J.W. *Industrial Dynamics*. Waltham, MA: Pegasus Communications, 1961.

———.*Urban Dynamics*. Waltham, MA: Pegasus Communications, 1969.

———. *World Dynamics*. Waltham, MA: Pegasus Communications, 1971/1995.

———. "System Dynamics: A Personal View of the First Fifty Years." *System Dynamics Review,* 2007, *23* (2/3) (Summer/Fall), 345–358.

Godley, W., and Lavoie, M. *Monetary Economics: An Integrated Approach to Credit, Money, Income, Production, and Wealth*. Basingstoke, UK: Palgrave Macmillan, 2007.

Hicks, J.R. *A Contribution to the Theory of the Trade Cycle.* Oxford: Clarendon Press, 1950.

Kaldor, N. "A Model of the Trade Cycle." *Economic Journal,* 1940, *50,* 78–92.

Kalecki, M. "A Macrodynamic Theory of Business Cycles." *Econometrica,* 1935, *3* (3) (June), 327–344.

Keynes, J.M. *The General Theory of Employment, Interest, and Money*. San Diego: Harcourt Brace Jovanovich, 1964.

Meadows, D.H.; Meadows, D.L.; Randers, J.; and W.W. Behrens III. *The Limits to Growth*. New York: Universe Books, 1972.

Montgomery, M. "'Time to Build' Completion Patterns for Nonresidential Structures." *Economics Letters,* 1995, *48* (2) (May), 155–163.

Radzicki, M.J. "Institutional Dynamics: An Extension of the Institutionalist Approach to Socioeconomic Analysis." *Journal of Economic Issues,* 1988, *22* (3) (September), 633–666.

———. "Institutional Economics, Post Keynesian Economics, and System Dynamics: Three Strands of a Heterodox Braid." In John T. Harvey and Robert F. Garnett (eds.), *Future Directions for Heterodox Economics*. Ann Arbor: University of Michigan Press, 2008, pp. 156–184.

Sterman, J.D. *Business Dynamics: Systems Thinking and Modeling for a Complex World*. Boston: McGraw-Hill Higher Education, 2000.

Tymoigne, E. "The Minskyan System, Part III: System Dynamics Modeling of a Stock Flow-Consistent Minskyan Model." Levy Economics Institute of Bard College, 2006, Economics Working Paper 455 (available at www.levyinstitute.org/pubs/wp_455.pdf.