A System Dynamics View of the Phillips Machine: "Chasing some Hares!"[1]

William H. Ryder[a] & Robert Y Cavana[b]

[a] 6735 Allview Drive, Columbia, Maryland 21046, USA. onebyke2ryders@gmail.com
[b]Victoria Business School, Victoria University of Wellington, NZ. bob.cavana@vuw.ac.nz

Abstract

In 1949, A.W.H. (Bill) Phillips, a New Zealand electrical engineer studying sociology as an undergraduate at the London School of Economics, built a hydraulic simulator (called the MONIAC – Monetary National Income Analogue Computer) of the British economy based on contemporary economic theory. It illustrated stocks, flows, employed feedback among variables, drew sharp distinction between exogenous and endogenous variables, and plotted its outputs as behavior over time graphs. All of the "Principles of Systems" articulated by Jay W. Forrester more than a decade later have examples in this simulator. Although the term "system dynamics" had not yet been coined in 1949, this paper suggests that Phillips' simulator could have been the first true system dynamics model applied to economics! The paper observes that this Phillips's model possesses the following characteristics common to most system dynamics models:

- Operational orientation.
- Feedback as a central design feature.
- Isolation of outflows from the associated stock.
- Use of non-linear functions.
- Use of the model as a means to explain a complex system to other people.
- Use of the model as a spur to additional models.

The paper includes a mapping of the physical components of the machine to system dynamics notation and provides an overview of a system dynamics (Vensim) simulator of the Phillips machine. A number of simulations with the model are presented. Finally some concluding comments are provided comparing the Phillips machine with other mainstream system dynamics models.

Introduction

In 1949, A. W. H Phillips and W. Newlyn designed and built an electro-hydraulic dynamic model of the macro-economy of a country. It was the first analog computer to solve the nonlinear coupled differential equations of mid-twentieth century economic theory (Bollard 2016) and led to the use of control theory to study stabilization issues in macroeconomics

[1] Paper presented at the 2nd Asia-Pacific Region System Dynamics Conference of the System Dynamics Society, February 19-22, 2017, National University of Singapore, Singapore.

(Phillips 1953), (Bissel 2007). Later known as the “Phillips machine”, the “Philips-Newlyn machine” or the “MONIAC”, the device demonstrated difficult macroeconomic concepts so clearly that it became a mainstay of instruction at the London School of Economics for more than 15 years (Barr 2000) and at Leeds University for more than 20 (Newlyn, 2000). Descriptions of the machine and its underlying macroeconomic theory have been given by multiple authors (Phillips 1950), (Newlyn 1950), (Moghadam and Carter 1989), (Newlyn, 2000), (Swade 2000), (Vines 2000), (Bissel 2007).

Facing a Phillips machine for the first time, the system dynamicist must feel its gentle seductions: the charismatic sloshing of the turbulent water, the animated movements of the machine as it responds to some external shock in its model economy, or the seamless union of engineering practice with economics theory. Here is a device that uses real bathtubs for stocks and real drains for flows, yet it is a puzzle in plastic. How does it work, and what does it mean?

Previous work described the Phillips machine in terms of elementary system dynamics components (Ryder 2009). The present work synthesizes the components into a complete model to simulate the machine. Starting from the detailed information in the cited papers, we reverse-engineer the Phillips machine to obtain simulator equations that both describe the machine and represent its intended theory.

Puzzles will confront us. The machine is really a portrait of its ingenious designer and renowned economist, Bill Phillips. Phillips once remarked that “I did not do very much. I just threw out a few hares for other people to chase” (Bollard 2016). The Phillips machine harbors quite a few such hares. We will chase some of them here. We begin with an overview of the machine’s structure, then explore its various sub-systems. Along the way we expose some of the implementation issues missing in the standard references. Finally, we review the significance of the Phillips machine as an important specimen in the evolutionary development of system dynamics models.

Overview

“By the late 1960s the two original MONIAC machines were left unused in the basement of the LSE. They stayed there until 1987 when the LSE donated one of the MONIACs to NZIER, where it was restored. Dr Alan Bollard, the then Director of NZIER, and present Reserve Bank governor, was instrumental in finding the machine at LSE and organising a fundraising drive to relocate it back to NZIER. He then proceeded to spend countless hours getting the Moniac back into working order. “ (NZIER, 2016)

Figure 1 shows the New Zealand Institute of Economic Research’s (NZIER) MONIAC machine. It stands over six feet tall, comprises of a number of plastic tanks and tubes through which coloured water flows. Linked to the tanks are gauges, sluices, pulleys and felt tip pens. The falling water accumulates in tanks containing floats. The floats drive sliding valves elsewhere in the machine through strings, pulleys and slot-cams, known also as functions, in white plastic frames. Together, the tanks, sluices, floats, pulleys, cams and valves comprise a stock and flow system with non-linear feedback control. Mechanically driven plotters at the top of the machine move charts sideways to record behavior over time graphs of important variables. The user allows the machine to equilibrate, then switches on the plotters and changes some variable. The plots capture the machine’s transition to a new equilibrium.

Figure 1: NZIER's Moniac machine is now on loan to the Reserve Bank museum in Wellington
Source: NZIER, http://nzier.org.nz/about/monia-machine/ downloaded 30 December 2016

We focus our overview with the comparative stock and flow diagrams in figures 2a and 2b. Figure 2a shows the stocks and flows in the notation of hydraulics, figure 2b shows the same model in the notation of system dynamics. Both diagrams omit the feedback structure. Figure 2a traces the path of water flowing around the Philips machine and uses the same or similar labels for the corresponding stocks, flows and flow nodes in figure 2b.

The Phillips machine forms a recirculating waterfall fountain. Water, pumped to the top of the machine, flows by gravity through a set of divides and gateways to either accumulate in one of two side tanks or to fall to the bottom tank where a pump sends it back to the top. The side tanks drain through controlled flows back to the main cascade of the waterfall. The heavy black bars in figure 2a are the horizontal sliding valves that control the various water flows. The valves with white arrows are normally controlled by the pulley-cam mechanisms, but can be unhooked from their cams at any time during machine operation and moved manually. With the feedback function-cams present, this mechanical arrangement simulates the incremental general macroeconomic variables in a country's economy.

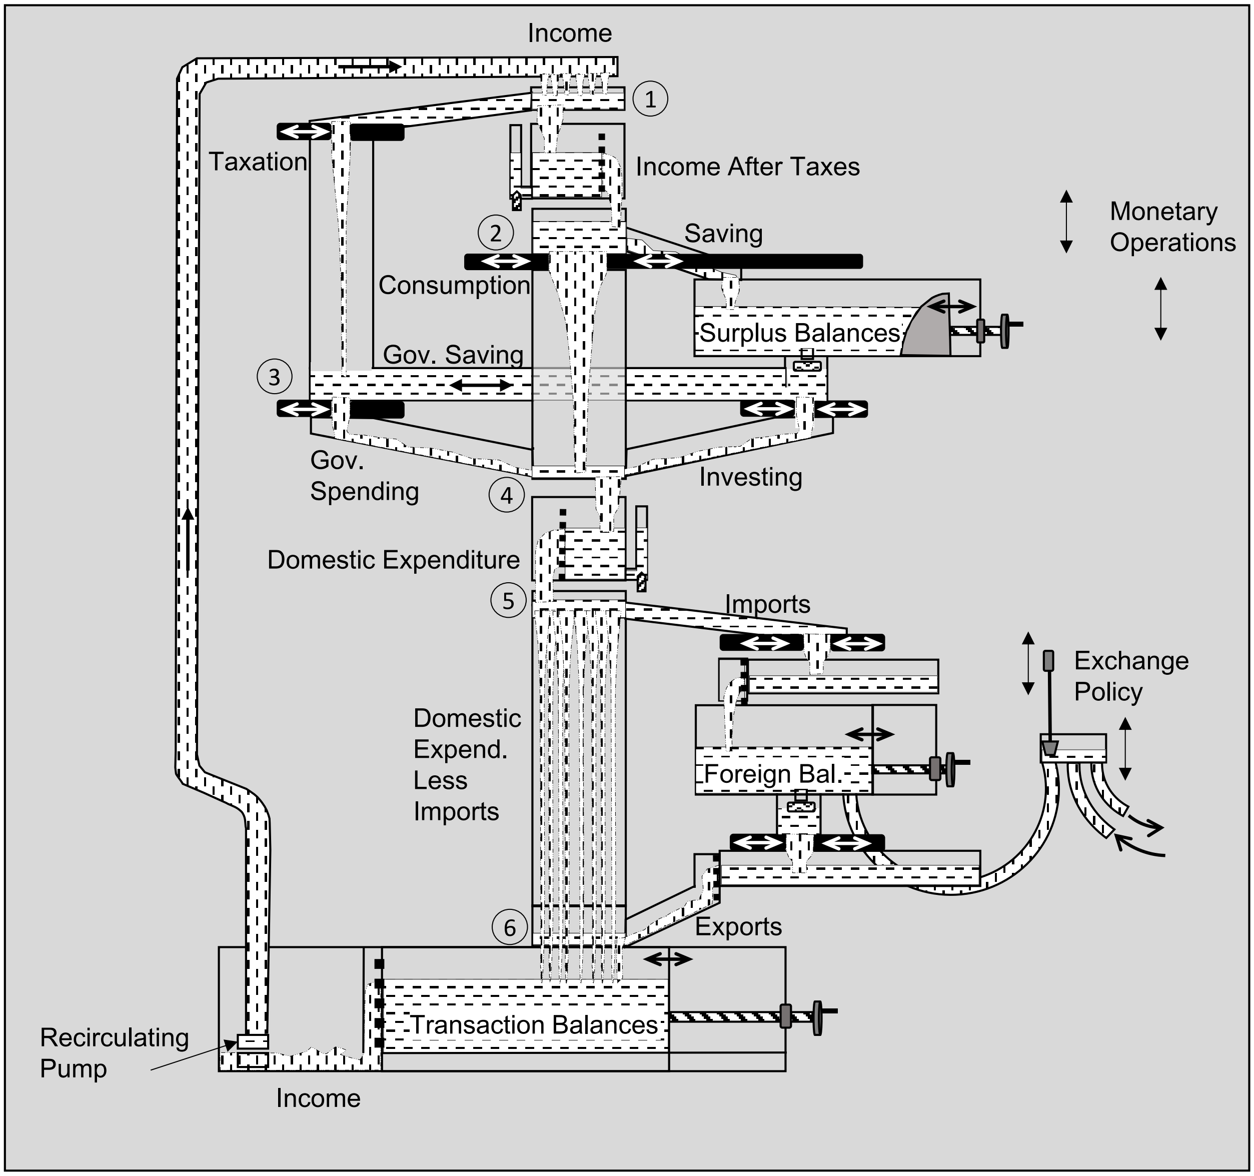

Figure 2a: Flows in Phillips' Model in Hydraulic Notation (Feedback Omitted)

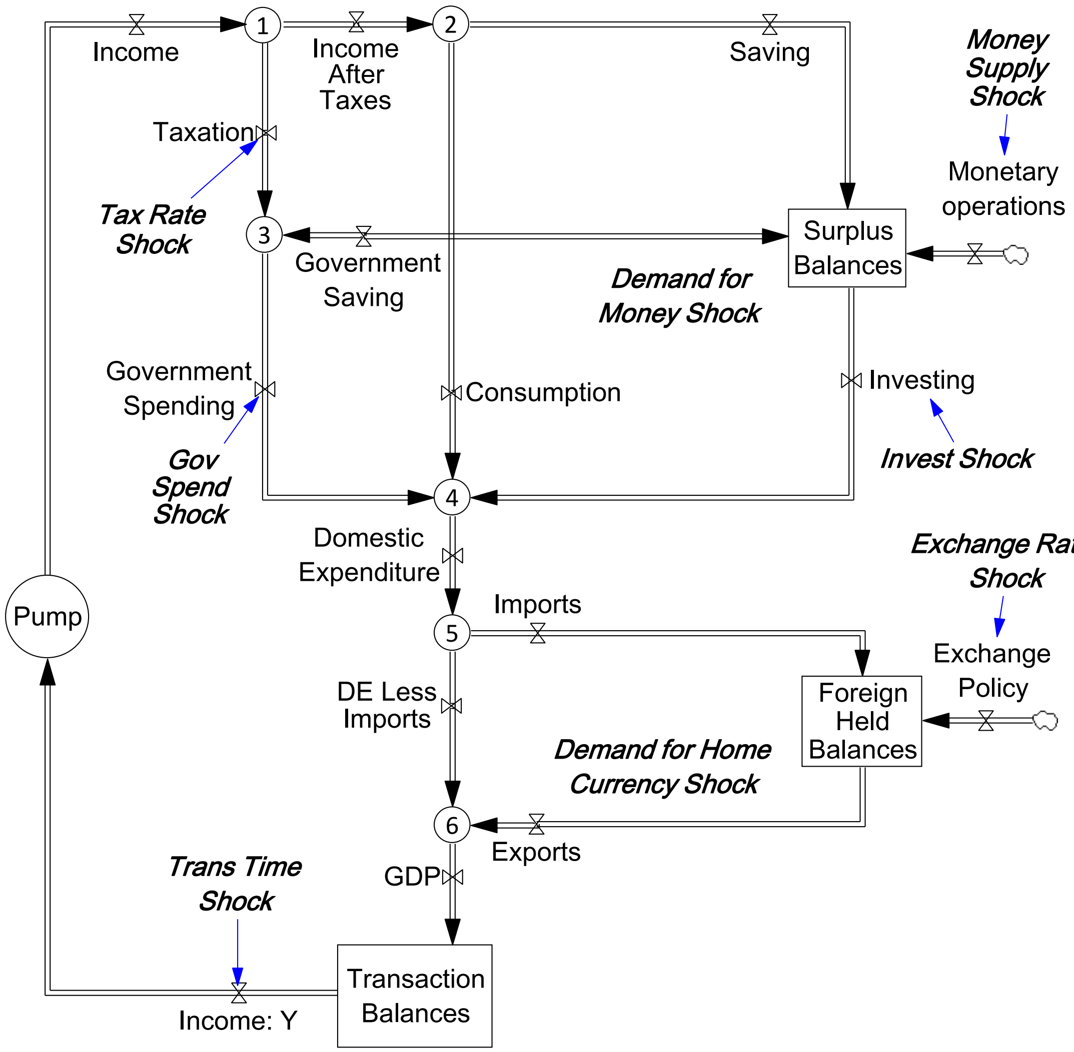

Figure 2b: Flows in Phillips' Model in System Dynamics Notation (Feedback Omitted)

The two side tanks represent different markets for money in the currency of the home country. The upper "Surplus Balances" tank represents the domestic money market where the price of money, or interest rate, is determined. The market includes a central bank that can fix interest rates to a desired level by increasing or decreasing the money supply from a hidden tank behind the machine. The lower "Foreign Bal." tank is similar but represents the foreign exchange market that determines the exchange rate. Foreign exchange market operations of the home country can influence the foreign exchange rate by supplying or removing additional home currency in the market from a second tank behind the machine.

The bottom "Transaction Balances" tank represents so-called active money in circulation. It is money changing hands but waiting for transactions to complete. The income flow due to the execution rate of transactions starts its circulation at the top of the machine where some of it is taxed to be spent by the government. The remainder is either spent on goods and services or accumulates in the "Surplus Balances" accounts. Simultaneously, another money flow, equal to the savings money flow at equilibrium, is withdrawn from "Surplus Balances" and combines with the money flow spent by the government and the money flow spent on goods and services to form "Domestic Expenditure". Some of this combined flow may be spent on imports while additional money flow may be recovered by selling exports. The net flow then returns to the "Transactions Balances" tank to wait, on average, the mean time between transactions before recirculating. In modeling terms, all the flow circulation occurs instantly and the delays between all transactions are lumped together in the single material delay stock called "Transaction Balances".

The user calibrates the machine to an economy in two stages: first setting the timescale for circulation, then fitting the various function cam shapes to match empirically measured relationships between the related independent and dependent variables. One begins with the knowledge of the total volume of water in the machine. This equates to the total money supply of the country, so the conversion between water volume and currency is established. The machine is designed to circulate all of its water in a few minutes, but the time can be adjusted by changing the tank capacity of "Transaction Balances" with a hand crank. The adjustment sets the relation of machine time to real economy time, because the ratio of total water volume to machine circulation time must equal the ratio of total money supply to observed circulation time. One strives to obtain an adjustment so that one minute of machine time equates to integral years of real economy time. The cams require a knowledge of the relevant variables for the country. For example, the function cam controlling "Taxation" flow as a function of "Income" would replicate the empirical relationship of actual tax receipts to actual national income of the country.

The openings of the horizontal sliding valves determine how the water splits through the various paths it can take. Because the head of water above all valves in the machine is the same, the same linear relationship between valve movement and water flow change applies to all valves. (McRobie 2009). Since the feedback function cams control the valve openings, the machine becomes a different economic model for each choice of function cam shape and control line engagement. Removing a cam will break its associated control feedback loop. Shifting a cam within its frame will shift the functional relationship it implements.

Stuff in the flows

The Phillips machine makes clear which variables are stocks and which are flows. Moving water signals a flow, stationary water signals a stock. So what is the meaning of the water that is in the main cascade and sluices and <u>not</u> in the stocks of the Phillips machine? Some of it appears to move while some of it appears stationary. There are eight places in the flows of the machine where stationary water accumulates. None of them are recognized as stocks in any of the literature. Let us consider these places:

- *The "Income After Taxes" box*

This is a small box holding a small amount of water at any one time. It has a tapered slot in one side that allows water to flow out quite rapidly and a manometer tube outside the box to permit a quiescent (averaged) value of the level in the box to be sensed. An adjusting screw in the passage between the manometer and the box controls the averaging time. Because the tapered slot has a particular shape, the level of water in the box and the rate of flow out of the box are linearly proportional. Thus, the level in the box estimates the rate of flow through the box – a brilliantly conceived flowmeter with no moving parts. Thinking in system dynamics terms we really have two stocks here. The small box with the tapered slot is a material delay with a short averaging time. It is in the flow of money and will delay the flow very slightly. The manometer compares its current level with the rapidly fluctuating level in the box and accumulates the difference with an adjustment time determined by the screw. This is an information delay with a potentially long delay time. It is not in the flow of money. However, the screw-determined information delay drives a function cam and could be very significant in the machine's dynamics.

The remaining places are:

- *The "Domestic Expenditure" box* - A flow meter with information delay similar to the "Income After Taxes" box.
- *The pool of water at node 2* - Forms the head above the consumption gate. Is an artifact of the machine and has no model significance.
- *The pools of water in the sluices at nodes 1 and 5* - Form the heads above the Taxation and Imports gates respectively. Artifacts of the machine.
- *The thin, long box between "Imports" and "Foreign Balances"* - A flow meter measuring the "Imports" flow. A float in this box drives a recorder pen.
- *The thin, long box between "Foreign Balances" and "Exports"* - A flow meter measuring the "Exports" flow. A float in this box drives a recorder pen.
- *The "Government Savings" connector tube between Government and Financial sectors* - A clever representation of the Government's sale or purchase of it own debt securities to or from the financial markets in order to balance its budget.

If we ignore the delays of the four flowmeter boxes and eliminate all the other non-stocks, we end up with figure 2b as our abstract SD model of the machine. It is an odd model, for most of its flows are "hooked", meaning they originate and terminate at the same stock. Absorbing those flows into the "Transactions Balances" stock and lumping "Saving" with "Government Saving", we arrive at the much simpler figure 3.

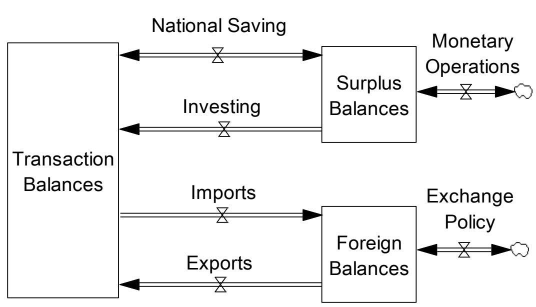

Figure 3: Simplified SD Model

From figure 3, one might guess that the Phillips machine simply allocates a total volume of water into three states, and that it will seek an equilibrium with the flows "National Saving and "Investing" equal and the flows "imports" and "Exports" also equal. By adding in the feedback control loops later we will reinforce this view.

Stock and Flow equations

We summarize the stock and flow structure of the Phillips machine by listing the stock and flow equations for an equivalent simulator in Table 1. Newlyn gives a complete set of equations for the first model (the Mark I) of the Phillips machine (Newlyn 1950, pp. 118-119). Our equations follow his, but for the Mark II model. Like Newlyn, we adopt abbreviated names for some variables. Most of the stock and flow equations can be formulated by inspection of figures 2a and 2b. Variables in **bold font** are defined in the feedback equations given later. Table 1 also omits definitions for constants such as time constants, and initial values.

Variable Name	Abbrev.	Equation	Units
Transaction Balances	TB	TB = Integral(GDO-Y) + TB0	Money
Surplus Balances	SB	SB = Integral(S + GSV + MO – **Inv**) + SB0	Money
Foreign Balances	FB	FB = Integral(**IM** – **EX** + EP) + FB0	Money
Income	Y	Y = TB/Transaction Time	Money/Time
Income After Taxes	IAT	IAT = Y – **T**	Money/Time
IAT Delayed	IATD	IATD = Integral(IAT-IATD)/IATtime + IATD0	Money/Time
Savings	S	S = IAT – **C**	Money/Time
Government Savings	GSV	GSV = **T** - **G**	Money/Time
Domestic Expenditure	DE	**G** + **C** + **Inv**	Money/Time
DE Delayed	DED	DED = Integral(DE-DED)/DEtime + DED0	Money/Time
Gross Domestic Product	GDP	GDP = DE – **IM** + **EX**	Money/Time
Monetary Operations	MO	(Target SB – SB)/SB Time	Money/Time
Exchange Policy	EP	(Target FB – FB)/FB Time	Money/Time

Table 1: Non-Control Stock and Flow Equations

Feedback controls

The economic actors' decisions and policies show up in the Phillips machine as feedback controls. Their implementation follows a common pattern. An observed variable, always a water level, is sensed by a float. The float's vertical movement translates via strings under tension, pulleys, levers, and cams to a sliding valve's horizontal movement thereby always regulating a flow variable. There are ten such control relations implemented. Nine of them translate from vertical to horizontal movements using the nonlinear function cam mechanism. This captures the known or assumed macroeconomic relationships between the independent float variable and the dependent valve variable. Conveniently for economists, the function cams have the independent variable on the vertical axis. The tenth, known as the accelerator, uses a bent lever that rotates around a fulcrum to perform the translation.

There are three types of float-sensors used: conventional buoyant floats, power-assisted sensors, and a unique leaky-float box used in the accelerator. The conventional floats are large and generate enough force as they ride the levels in their respective tanks to operate two function cams with force to spare for a plotter scribe. The power-assisted sensors are actually electrodes. They are embedded in a weight suspended on thin wires from a fixed pulley that is driven by a reversible torque source. The torque source, switched by the electrodes, will pull them up out of the water when they are immersed but drops them back into the water when they are not. Thus, the weight behaves like a float in that it follows the water level. The rotational position of a second fixed pulley on the same shaft then drives one or more cams. This arrangement is used to sense the levels in the manometers on the "Income After Taxes" and "Domestic Expenditure" flow boxes. Small conventional floats in the manometers would not produce enough force to operate the cams. On the original machines the torque sources were "vaned wheels with water falling in different amounts on either side" (Moghadam and Carter 1989 p. 25). Modern restorations have used servo motors instead.

The accelerator deserves special attention. Phillips wanted a mechanical method to represent the fact that firms build inventory and make other investments based on expectations of future activity. An estimate of the trend of recent Income, in addition to the value of Income itself, ought to be one of the components of Investing. His extraordinary solution was a box on a string in the "Transactions Balances" tank. The box has a hole in the top and an adjustable hole in the bottom. Such a box will sink as water enters from the bottom and air escapes through the top. However, it does not sink because the string is suspended from a fixed point high on the machine through a spring.

Control equations – Closing the loops

We now formulate the definitions of the flow control variables missing from table 1. These equations connect the observed variables sensed by floats to the valves to implement feedback in the model. For clarity and brevity we adopt three conventions: First, function names will include the abbreviated name of the control variable they regulate, possibly with subscripts as needed. For example, the equation defining "Consumption" will be written $C = C_1(IAT) + C_2(i)$. Second, the variables used in the equations will reflect the intention of the machine design, not the physical design. Thus, the equation for "Government Spending" will appear as $G = G(Y)$ rather than $G(TB)$, even though the sensing float is clearly in the TB tank. This interpretation assumes that Phillips designed the "Transaction Balances" tank as a material delay flow meter where the tank estimates a flow by virtue of its linear flow slot. Phillips himself validates this view (Phillips 1950 in Leeson, p 73). Third, we assume that functions contain the appropriate conversion factors inside them to translate the units of their input variable into the units of their output variable. The functions as listed here do not have normalized dimensionless inputs and outputs as recommended by Sterman. (Sterman 2000) Table 2 lists the feedback equations.

Given the equations in tables 1 and 2, one can set up a working simulator. The sketch for such a simulator is given in Figure 4.

Variable	Abbrev.	Equation	Input Units	Output Units
Taxation	T	T = T(Y)	Money/Time	Money/Time
Gov. Spending	G	G = G(Y)	Money/Time	Money/Time
Consumption	C	C = C1(IAT) + C2(i)	Money/Time, 1/Time	Money/Time
Investing	Inv	Inv = MIN(Inv(i) + Y*PGAIN + DY*DGAIN, SB/mintime)	Money/Time	Money/Time
Income Change	DY	DY = Y – Recent Y	Money/Time	Money/Time
Averaged Y	Recent Y	Recent Y = Integral(DY/AccTime) + Y_0	Money/(Time^2)	Money/Time
Imports	Im	IM = MIN(Im(DED) + Im(Fxch), FB/mintime)	Money/time	Money/Time
Exports	Ex	Ex(DED) + Ex(Fxch)	Money/Time	Money/Time
Interest Rate	i	iLPF(SB)	Money	1/Time
Exchange Rate	Fxch	Fxch = iDCP(FB)	Home Currency	$\frac{Forgn\ Currency}{Home\ Currency}$

Table 2: Feedback Equations

In table 2, the abbreviations iLPF and iDCP stand for the inverse functions to the Liquidity Preference Function and Domestic Currency Preference functions respectively. AccTime, mintime, DGAIN, PGAIN, and Y_0 are constants. The MIN() functions in the equations for Investing and Exports prevent the stocks they drain from having negative values.

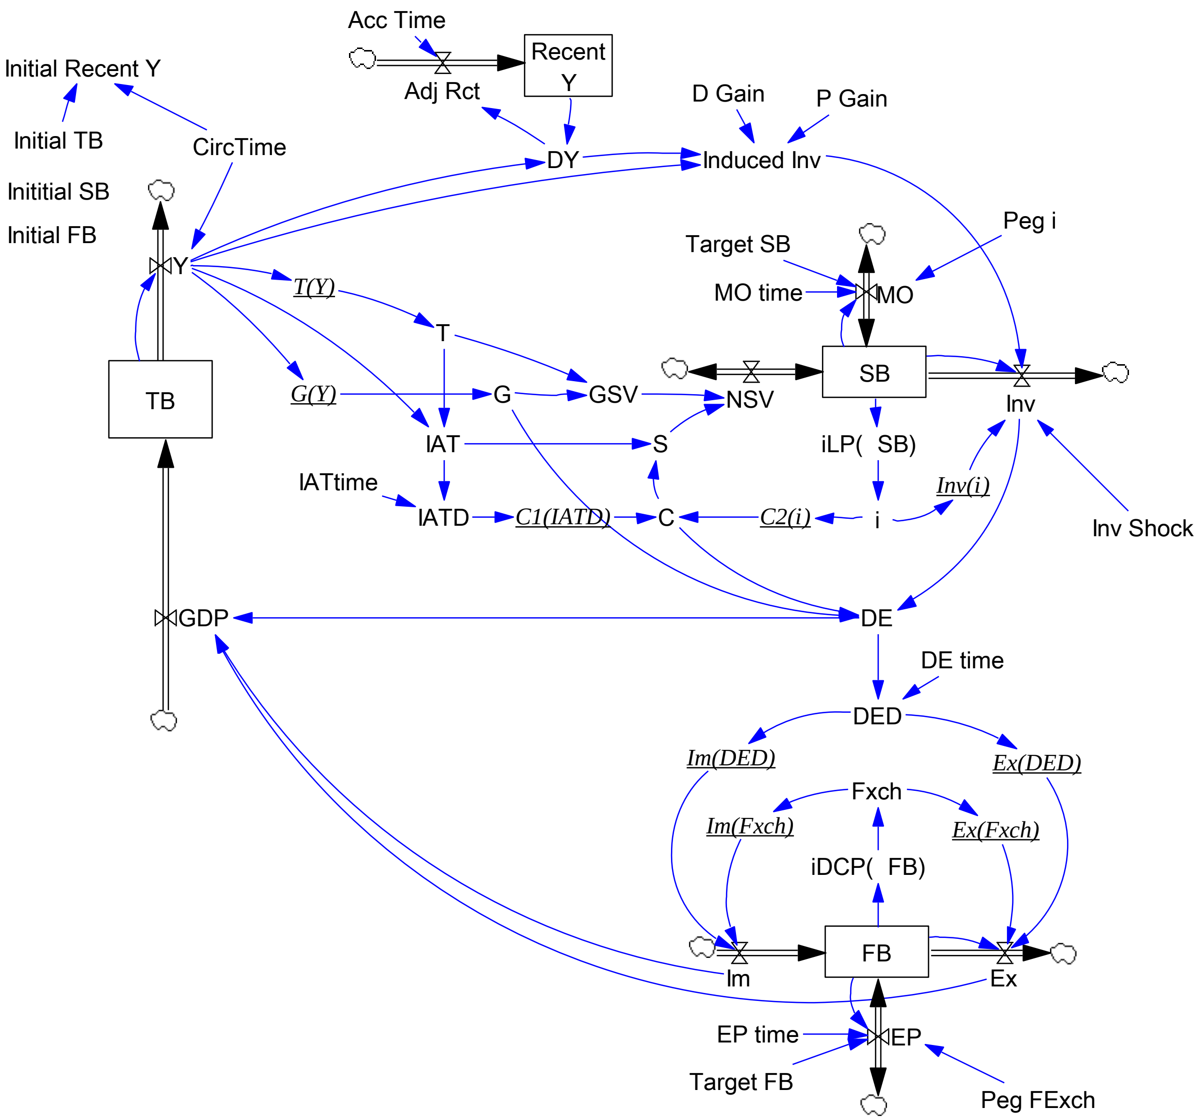

Figure 4: Sketch of a minimal Vensim simulator for the Mark II Phillips machine.

In figure 4, the variables “Peg I” and “Peg FExch” are switches that enable or disable the “MO” and “EP” flows respectively. They each correspond to stoppers that disconnect or connect tanks behind the machine to fill or empty the two stocks so as to hold their levels constant. The “Inv Shock” variable is a step function applied at time = 25, after the simulator’s startup transient has died out. To apply shocks at other places in the machine, one must add similar step function variables to the above list. The variables in *underlined italicized font* represent the function cams.

In designing the cams, it is essential that one understands the role of each one in the larger system. The cams and the accelerator act in pairs, either in a forcing role or an equilibrating one. We offer here a few observations concerning each pair’s role.

The “Taxation” and “Government Spending” functions act as a forcing pair that operate differentially in a “push and pull” manner on the rest of the system. When the government’s budget is balanced, i.e., both functions are identical in shape and position, only the tax function drives the system through “Saving” and “Domestic Expenditure” flows. Specifically, raising taxes lowers “Saving” and could increase “Domestic Expenditure” if “Investing” does not decrease also. With a deficit, i.e., the government spending function exceeds the taxation function for some income value, raising both an equal amount will decrease ”Saving” and drain the “Surplus Balances” tank as if “Investing” had increased. In classroom use, the government spending cam, representing an automatic government fiscal policy as a function of income, would be removed and a student would supply the government’s fiscal policy decisions as the machine operated.

The “Consumption” cam and the accelerator act as the forcing pair for the domestic money market. “Consumption” regulates the left side of “Saving” (inversely because “Saving” is a residual). The accelerator operates the left side of “Investing” directly. Together, they act to move “Surplus Balances” to a state other than the “natural” equilibrium set by the equilibrating pair for “Surplus Balances”. The “Consumption” cam need not be linear, for its slope, the Marginal Propensity to Consume, will probably be a function of income.

The two left hand cams for “Imports” and “Exports” are the forcing pair for the foreign exchange market. Unlike the other forcing functions, they are functions of “Domestic Expenditure” rather than “Income”. If identical, the cams have no effect on the state of “Foreign Balances” and do not change the total expenditure of the economy. If they differ for any value of “Domestic Expenditure” their action will drive “Foreign Balances” away from its “natural value” implied by the equilibrating pair. The size of the difference determines the strength of the influence. Typically, the imports cam will increase “Imports” while the exports cam will decrease “Exports” as “Domestic Expenditure” increases.

The remaining two pairs of cams equilibrate their respective markets with the strength of the effect determined by the slope of their production – consumption difference function. The stronger the effect, the faster the markets will adjust to a shock pushing away from equilibrium. Moreover, the difference slope must act to fill the underlying stock if the level is below the natural equilibrium (the intersection point of the two functions) and drain it if the level is above natural equilibrium for the system to be stable. The functions do not have to have opposite slopes – the slopes could both be positive or negative. They just have to intersect. Phillips’ diagram of the machine shows the equilibrating pair in the foreign exchange market with the same slope sign. (Phillips in Leeson 2000, pp. 60 and 84). It is important to remember that the graph of the upper equilibrating function in the money market is flipped (horizontally on the machine and vertically in the simulator) to account for the fact that it acts through consumption and thus inversely on saving.

The equilibration strength of the “Surplus Balances” cam pair must be small relative to other effects. Phillips points out that:

> “The effect is weaker and slower in action than was thought before the work of Keynes, and, moreover, in time of wide fluctuations in income it may be almost completely swamped by the effects of the accelerator and shifts in the liquidity preference and marginal efficiency of capital functions. But if sudden changes in the level of income are avoided by fiscal or other policies, the equilibrating influence of the rate of interest

becomes relatively stronger, so monetary policy becomes a necessary supplement to fiscal policy." (Phillips in Leeson 2000, p.82)

We speculate that the two functions are relatively flat (for the simulator) or vertical (for the machine).

Using the Phillips machine

Before running a Phillips machine or a machine's simulator, one first configures it. This means: 1) designing the specific shapes for each of the nine function cams and deciding which of them to install for the experiment at hand, 2) setting manually adjusted valves to desired positions, 3) adjusting the hole in the accelerator, 4) adjusting the attachment point of the accelerator string to the bent lever, 5) applying or removing the plugs connecting the spare tanks to the market stocks, 6) setting the heights of the two boxes containing the plugs, 7) setting the timescale by adjusting the moveable side in the "Transactions Balances" tank, and 8) laying out a plan for operating the machine during the experiment.

We give here a few examples of simulator operation. The examples were suggested by Phillips in his user manual for the machine.(Phillips in Leeson 2000, pp59-62) We will consider three experiments: 1) illustrating the investment multiplier for cases of fixed amount of money and fixed interest rate, 2) illustrating the effect of a shift in the demand for liquidity, and 3) showing the effect of the accelerator in creating instability in the economy. In each case we modified the equations above to add step shock variables.

Investment multiplier

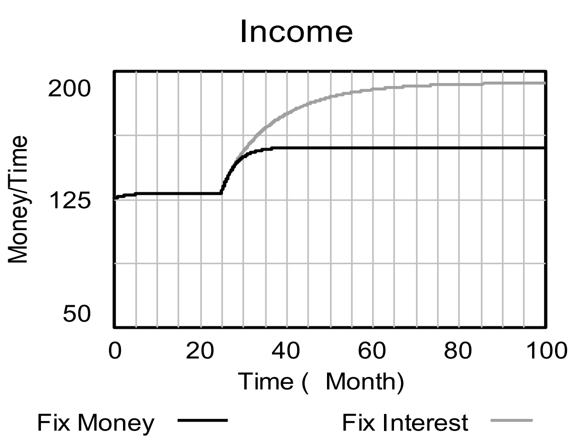

Figure 5: Investment multiplier experiment

This experiment superimposes two runs, Fixed Money and Fixed Interest. For Fixed Money, we ran the simulator with no central bank intervention and a step increase of 30 monetary units per month applied to Investing at time 25 months. The Fixed Interest rate run is the same except the central bank begins to pursue a fixed interest rate policy at time 25 months. The target interest rate is the same as at the beginning of the simulation. All functions are present for both runs, so all the listed feedback equations operate except the accelerator, whose two gains were set to zero. The graph, figure 5, shows that although the multiplier value is larger with central bank intervention, the full effect takes longer to achieve.

Shift in the demand for money

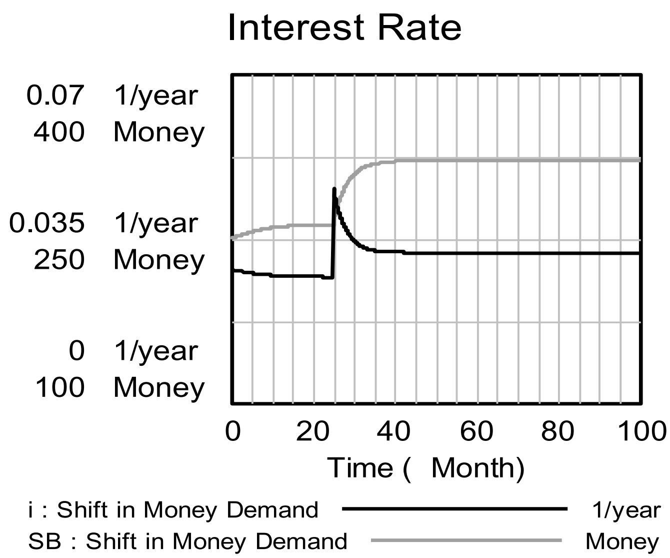

Figure 6: Response to a shift in the demand for money

This experiment shows the effect on interest rate when the demand for money shifts suddenly. On the Phillips machine, this corresponds to shifting the curved moveable side of the "Surplus Balances" tank, outward for increases in demand and inward for decreases. For the simulator it means adding a constant to the input of iLPF(), the inverse of the Liquidity Preference function. The constant will be negative for increases in money demand and positive for decreases. Here we shift demand for money by 100 monetary units at time 25. The result in figure 6 shows that the interest rate rises immediately when demand increases, but falls as the higher rates attract more money into the "Surplus Balances" accounts.

Destabilizing effects of the accelerator

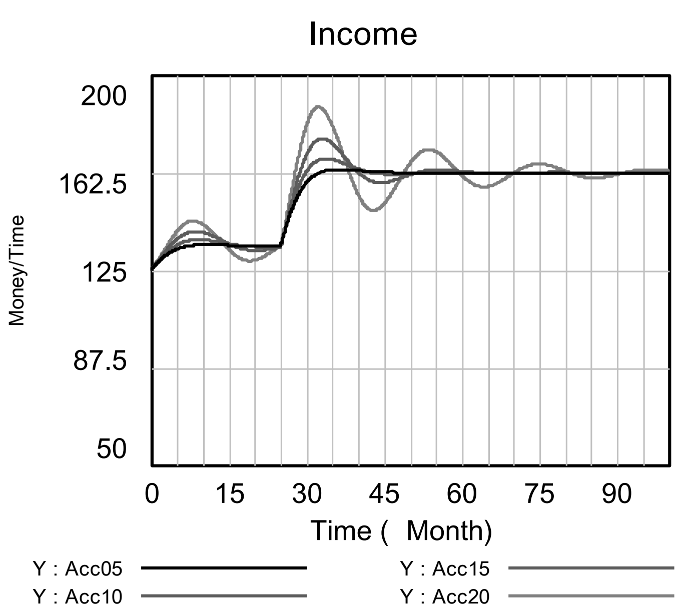

Figure 7: Connecting the accelerator

This experiment explores the effects of induced investment when a sudden change occurs in "Investing" at time 25 in an economy with a fixed amount of money. Four runs with different parameters for the accelerator gains were performed. In each run, the lag time for the accelerator ("Acc Time" in the simulator) was set to 3 months to represent quarterly estimates of trend. The "P Gain" variable was set to .05, while the "D Gain" variable ranged from 0.5 to 2.0 in increments of 0.5. Therefore, the feedback control from the accelerator was almost exclusively from the differential component. Each of the runs show the response to a step change in “Investing” of 30 monetary units per month. All the runs are superimposed in figure 7. The large variations of “Income” suggest that the economic system captured by the model is quite sensitive to small changes in feedback parameters.

Concluding comments

Because the Phillips machine preceded Jay Forrester’s founding of the field of system dynamics by a decade, it has not been considered an early example of that discipline. Yet the pieces of system dynamics models are all there: the stocks as tanks, the flows as flowing liquid, the feedback control involving actors’ decisions as the nonlinear functions, the exogenous assumptions as manual settings. Even the behavior over time graphs appear.

Subsequently Haywood and Cavana[2] (1986) outline a small medium term system dynamics model (SDMACRO) of the New Zealand macro economy developed at the New Zealand Planning Council in the mid 1980's. SDMACRO was developed to provide likely trend movements and scenarios, some 10-15 years into the future, in the key macro-economic aggregates including gross domestic product, capita1 formation, population, employment, exports, imports, and the current account balance. This model has also been discussed at earlier International System Dynamics Conferences (Cavana & Haywood, 1988; Cavana, 2005).

More recently, Wheat has used system dynamics to teach macroeconomics (Wheat 2007). Portions of Wheat’s economic model closely resemble portions of the Phillips machine. What separates the Phillips machine from mainstream system dynamics models?

Saeed offers a possible criterion (Saeed 2014). Sketching a simplified version of Forrester’s economic model from Urban Dynamics, Saeed argues that it is operational thinking – the direct representation of physical world activities, decisions and entities by model components -- that distinguishes system dynamics economic models from other economic models. Operational Thinking sets a limit on the amount of abstraction in the models, apart from simple aggregation, and encourages the development of realistic models. Saeed uses this criterion to distinguish classical equilibrium models and models based on perfectly rational agents from those of system dynamics. We suggest that the Philips machine passes Saeed’s test. The water in the Phillips machine represents circulating money denominated in the home currency, the tanks represent aggregated bank accounts, and the water flows from one state to another under the direction of decisions made by not necessarily rational actors whose actions are modeled by the non-linear functions. Decisions are present in all feedbacks involving flow gates. If we add the decisions and their feedback loops to the simplified diagram of Figure 3 (without the foreign sector or Government saving), we arrive at Figure 8. Here the upper “saving” decisions are made by households and investors, while the lower “investing” decisions are made by firms and investors. The degree of aggregation stems from the two purposes of the model: 1) to illustrate how the rate of interest is determined and 2) how instabilities in the economy might arise. Agency is clearly apparent.

[2] Dr Bob Cavana is a former student of Prof Bill Phillips (at the University of Auckland, New Zealand), on his Masters course in ‘Economic Dynamics’.

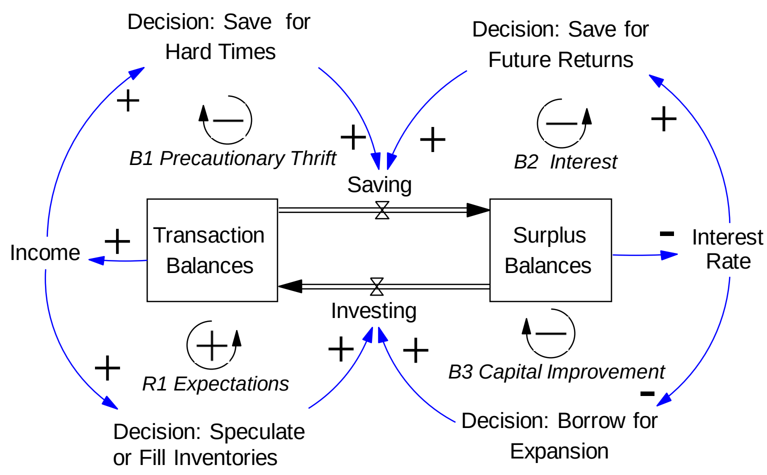

Figure 8: Principal Feedback Loops in the Domestic Economy

Olaya offers a stronger test (Olaya 2015). He requires not only that the model represent the decisions and actions of actors of a system (its agency), but also that the model structure include all of the main components needed to explain how the system actually works in producing the observed behaviors of the key variables. To restate his example, if you want to build a model about milk production, you had better include cows as an endogenous variable. Stated another way: the model boundary must enclose the parts of the system that are relevant to the purpose of the model. The Phillips machine fails this test. Because central bankers often adjust interest rates in response to inflation, there must be an operational causal link between the two. Yet the Phillips machine does not represent prices, the production of goods and services, labor, or the inventory of goods. The major sectors of the economy such as the goods market, the labor market, the inventory of capital, the productivity of the factors of production and the availability of raw materials lie outside the model boundary of the Phillips machine. Without operational determination of wages or prices, it cannot represent inflation or a reasonable endogenous model for the decisions of households and firms to consume. Yet the decisions to consume strongly influence how money circulates in the model's economy. By Olaya's test, the Phillips machine represents a partial economic model whose restrictive model boundary compromises its putative standing as a true system dynamics model derived through operational thinking.

Yet we must make some allowances for the technology Phillips had to use. In order to include stocks and flows of workers, goods, and raw materials, Phillips would have required a separate circulating fountain with its own circulating pump for each type of "stuff". The variables in the separate circulating sectors would then have to be cross-connected via (possibly) electro-mechanical actuators similar to those in the flow boxes of the present machine. Perhaps this is why there are so few hydraulic-based system dynamics simulators in the world. It is too difficult to represent multiple types of "stuff".

Would Phillips have extended his model boundary through continued use of the operational thinking evident in the present machine? We only know he did not do so. He could have done it in principle, for he had already solved all the technical problems standing in the way. However, the machine, so adept already at teaching major aspects of the circulation of money, would have been at least 4 times larger, unwieldy and expensive beyond Phillips's budget. Classroom apparatus should be simple and to the point. By employing extraordinary ingenuity, Phillips obtained maximal explanatory power from a minimal number of hydraulic components within the present machine. Having accomplished his objective of an operational demonstration of contemporary economic theory, he moved on to other important topics.

We are fortunate to be able to reflect on Professor Bill Phillips' pioneering work on economic stabilization modeling by offering a modern day system dynamics interpretation of this seminal work.

References

Barr, Nicholas 2000. "The History of the Phillips Machine", in *A.W.H Phillips: Collected Works in Contemporary Perspective*, R. Leeson, Ed. Cambridge, UK: Cambridge Univ. Press, 2000, pp 89-114.

Bissell, Chris 2007. "The Moniac", IEEE Control Systems Magazine, February 2007 pp 69-74.

Bollard, Allan 2016. *A Few Hares to Chase – the economic life and times of Bill Phillips*, Oxford, UK, Oxford University Press 2016.

Cavana RY, Haywood E. 1988. "A System Dynamics Model of the New Zealand Economy". *Proceedings of the 1988 International System Dynamics Conference.* San Diego, USA. pp27-38

Cavana RY. 2005. "Revisiting medium term macro-economic scenarios (1985 - 1995) generated by a system dynamics model of the New Zealand Economy". *International Conference of the System Dynamics Society*, Boston, SDS

Haywood E, Cavana RY. 1986. *A Macro-Economic Model and Scenarios to 1995*. New Zealand Planning Council: Wellington (Planning Paper No. 24).

Moghadam, R. and Carter, C. 1989. *Economic Affairs*, October/November 1989 pp 21-27.

McRobie Allan 2009, Private Communication.

Newlyn W. 1950. "The Phillips/Newlyn Hydraulic Model", *Yorkshire Bulletin of Economic and Social Research*, 2.2, /September 1950 pp 111-127.

Newlyn W. 2000."The Origins of the Machine in a Personal Context", in *A.W.H Phillips: Collected Works in Contemporary Perspective*, R. Leeson, Ed. Cambridge, UK: Cambridge Univ. Press, 2000, pp 68-88.

NZIER (2016). "The Moniac machine is a dynamic model of a working economy". New Zealand Institute of Economics Research, Wellington, New Zealand. Available at http://nzier.org.nz/about/monia-machine/ downloaded 30 December 2016

Olaya, C 2015. "Cows, agency, and the significance of operational thinking", System Dynamics Review, Vol 31, No 4, 2015

Phillips, A.W.H. 1950. "Mechanical Models in Economic Dynamics", in *A.W.H Phillips: Collected Works in Contemporary Perspective*, R. Leeson, Ed. Cambridge, UK: Cambridge Univ. Press, 2000, pp 68-88.

Ryder W. 2009. "A System Dynamics View of the Phillips Machine" Presented at the 27th International Conference of the System Dynamics Society, available on line at http://www.systemdynamics.org/conferences/2009/proceed/index.html , last referenced 2 March 2016.

Saeed, K. 2014. "Jay Forrester's Operational Approach to Economics", *System Dynamics Review*, Vol 30, No. 4, 2014.

Sterman J.D. 2000. *Business Dynamics – Systems Thinking and Modeling for a Complex World,* McGraw Hill. 2000 Chapter 14.

Vines, D. 2000. "The Phillips Machine as a 'Progressive' Model" in *A.W.H Phillips: Collected Works in Contemporary Perspective*, R. Leeson, Ed. Cambridge, UK: Cambridge Univ. Press, 2000 pp 39-69.

Wheat, D.I. 2007. "The Feedback Method of Teaching Macroeconomics: Is It Effective?", *System Dynamics Review* Vol 23, No. 4, 2007