# *Before You Start Your "Intro to Economics" Course…*

Let's start with the obvious question: why am I writing this for you, and ask that you read it before you start to study economics? There is an old saying, which has since been confirmed by neuroscience research that, "it takes twice as much information to change an opinion as it does to originally form it." When you initially encounter new subject matter, one of the first cognitive tasks you automatically perform is the construction of a new mental model, which summarizes the most important points and how they are related to each other. Mental models are therefore at the root of three critical cognitive processes: (a) they focus our limited attention on what we believe to be most important in a situation; (b) they provide the basis for explaining what we observe; and (c) they provide the basis for predicting how a situation is likely to evolve in the future. Researchers have repeatedly found that there is a strong relationship between the accuracy of one's mental model about an area, and one's relative performance. So to sum up: mental models are important, particularly the ones that we form when we initially encounter new subject matter.

Unfortunately, based on experience I fear that the mental model you will develop as a result of your "Intro to Economics" course is likely to be inaccurate. Over the past decade, the conceptual models and frameworks that are typically taught in such courses have been deeply challenged by two developments: the digitization of our economy and the global financial and economic crisis that erupted in 2007 and has yet to end. The good news is that these events have led to a substantial rethinking and evolution of some critical concepts in economics; the bad news is that few if any of these new developments have yet made their way into Intro Economics textbooks and the materials used by the people who teach these classes. Hence this note, the goal of which is to ensure that the mental model you develop over the next few months is as accurate as possible. If that leads to the occasional conflict with your teacher, so be it; given the long term cost of an inaccurate mental model, ruffling a teacher's feathers is a small price to pay.

So, with that in mind, I'll move onto a brief review of some concepts and issues I'd like you to keep in mind this year.

Micro vs. Macro

Micro is the study of the decisions and behaviors of individual agents in the economy. Sometimes these agents are taken to be individuals (e.g., consumers) while at other times (confusingly) they are taken to be organizations (usually assumed to be a profit seeking private sector company, rather, say, that a government or not-for-profit organization).

Macro is the study of the behavior of economic aggregates at the national or international level, such as the level and growth of the money supply, interest, and exchange rates, and changes over time in the level and composition of total output in an economy (e.g., gross domestic product) including private consumption, private investment, government consumption and investment, exports and imports.

The Way Micro is Traditionally Taught

Micro textbooks often make some very unrealistic simplifying assumptions in their analysis of the way individual agents make decisions. To wit:

> *Demand curves always slope downward (i.e., from the upper right to the lower left) on a Price (vertical)/Quantity(horizontal) graph, due to the diminishing satisfaction one obtains from consuming an additional unit of a given good.*
>
> *Supply curves always slope upward, over the time frame that is assumed to be relevant to making the decision. The underlying assumption here is that all workers are paid the same, and that additional workers can be easily added to produce more goods. At first, the marginal cost of producing an additional unit of a good is assumed to decrease, as workers gain experience using the*

*machines and managing the production process, and output per worker increases, while cost per worker does not. However, once the limit of this productivity growth is reached, additional output requires the addition of new workers, who are assumed to be less productive than the current workers. This causes average output per worker to decline, and the marginal cost of producing an additional unit of a good to increase.*

*Decision makers have access to perfect information at zero cost.*

*They have infinite cognitive resources, and are always rational.*

*They make decisions in isolation.*

*Any uncertainty they have about the future can be reduced to a set of possible outcomes that encompasses the full range of possible future conditions. The costs and benefits associated with each of these possible future outcomes are also known with certainty. Finally, the decision maker is also assumed to have an accurate understanding of the probability that each possible outcome will occur.*

*Competition is always perfect; no company has market power. This means that decisions can be made without regard to what competitors may do.*

*Decisions are usually made without regard to financing, which is assumed to always be available at an acceptable cost.*

<u>The Real World Doesn't Work Like A Micro Textbook. I'll Take Each Of The Above Assumptions In Turn:</u>

- Some goods provide status benefits that increase in proportion to their relative scarcity and the price that must be paid to obtain them. Hence, demand for a

good or service that confers status in the eyes of other people can actually increase as its price goes up. Put differently, the satisfaction you get from a good that confers status in the eyes of others may actually increase the more of it you consume, and the more your social status rises.

- First, the assumptions associated with rising marginal cost have always rested on shaky ground. For example, labor costs are usually not constant (e.g., the most productive workers can be paid more), nor is the amount of equipment (e.g., machines) per worker. And the quality of that equipment changes over time – think of a worker shaping a piece of metal with a lathe and a file, compared to one who programs a computer controlled machine tool. Moreover, with effective hiring and training, you can often quickly increase a new worker's productivity. Beyond that critique, it is also critical to recognize that we now have two economies – one physical, that is described in traditional economic textbooks, and one digital. For example, think of the way music used to be sold – as physical products, like vinyl albums, tapes, and CDs, in physical stores. In this case, there was clearly additional cost associated with producing and selling more copies of a popular song. But now think of iTunes, and ask yourself what is the marginal cost of producing more copies of a popular song today? Essentially, it is ZERO, as Apple's servers can keep downloading songs forever, with the only additional cost a tiny fraction of a cent to pay for the tiny amount of electricity used to perform the operation on a few chips in the server. In the digital world, supply curves do not typically slope upward to the right.

- Decision makers do not have access to complete and perfect information. And the information they have access to usually costs something to obtain. Granted, with the internet the marginal cost of obtaining information has fallen dramatically. But there is still an opportunity cost associated with the time spent to obtain it. And the quality and reliability of the information obtained from an internet search still varies widely.

- Human beings do not have infinite cognitive resources. Our bodies and brains get tired. In fact, research has shown that decision quality worsens as more decisions, and complex decisions in particular, are made in a fixed period of time. In addition, we are not perfectly rational thinkers. Our thinking naturally falls prey to a range of predictable cognitive biases, and is also influenced by our emotions. In fact, decisions inevitably result from both rational and emotional considerations, though the latter are often subconscious.

- Human beings do not make most decisions in isolation. Rather, the often consult with other people before making them, and are concerned with how other people will perceive the result. For example, researchers have repeatedly demonstrated that people will usually make different choices about which is the best of ten songs when they are evaluating them in isolation and when they can see other people's evaluations before they make their own. Or the CEO of a widget company may want to know his Board of Directors views about a decision before he makes it.

- You need to distinguish between situations and choices that are "risky" and those that are "uncertain". Frank Knight drew a critical distinction between the two in his 1921 book, Risk, Uncertainty, and Profit. In "risky" situations the full range of possible outcomes and their associated economic impacts and probabilities are all known; in "uncertain" situations, this is not the case. Most of the decisions you encounter in the real world involve uncertainty, not risk. And human beings have much more difficulty making decisions in the face of uncertainty than they do in the face of risk.

- In the real world, competition is not perfect. Some companies have advantages – e.g., patented technology, or access to cheaper resources – that enable them to earn higher profits than competitors. In fact, one view of business strategy is that its essence lies in identifying, implementing, and protecting barriers to

competition that enable a company to earn higher than average profits. For this reason, real world decisions often take the position and potential actions of competitors into account.

- Microeconomic textbooks like to assume that the financing required to implement any decision – e.g., building a new plant to increase widget production – will always be available at an acceptable cost. In the real world, this is most definitely not the case. For example, a company may have a much more optimistic view of the future demand for widgets (or the future evolution of widget technology) than potential providers of the financing needed to build a new factory. Beyond the simple availability of financing, its cost and its form are also usually important considerations in making an investment decision. Side note: financing comes in two broad flavors. Debt financing (i.e., loans or bonds) has an explicit cost (the interest rate on the loan or bond), and must be repaid on or before a fixed date (e.g., you borrow $100 from me for a year to buy a weed wacker to start a lawn care business. I charge you 5% interest per year, payable in arrears. At the end of 12 months, you owe me $105 -- $100 repayment of the loan "principal" and $5 of loan interest). Equity financing (stock) has neither an explicit cost nor does it have to be repaid. It is capital that is permanently provided to the company. Equity investors obtain their returns in the form of dividends (annual cash payments to stockholders that the company makes at its discretion) and the increase in the value of their stock as the company (hopefully) grows over time. To go back to our weed wacker example. Instead of a $100 loan, I could give you $100 for shares that provide me with a 50% ownership interest in your company (with you owning the other 50%). If you sell your company in two years for $1,000, we each walk away with $500.

- If after reading this you strongly suspect that most business people would say that microeconomics, as it is too often taught, has little or no relevance to how they actually run their business, then you are right.

The Way Macro is Usually Taught

The bad news is that the current state of macroeconomics education makes micro look like a paragon of intellectual clarity.

Macro courses too often still rely on two frameworks. The first breaks down economic output (i.e., total supply or total demand, since supply is assumed to equal demand) into its main components: private sector consumption spending; private sector investment spending (e.g., in business structures, inventories, equipment, and software; and household investment in new homes); government consumption and investment spending; exports (goods and services we produce in our country and sell to buyers in other countries) and imports (goods and services we buy from suppliers located in other countries). Exports and imports are usually netted against each other to produce "net exports" or (with a few technical adjustments) the so-called "current account balance" (the current account being one part of our overall balance of payments to and from other countries). The private sector can be further broken down into its three main subgroups: households, businesses, and financial institutions, like banks.

The second framework is called the "IS-LM" model. The LM part of the model tries to describe the monetary sector of the economy, in terms of the quantity of money supplied by the government and economic agents' demand for "liquidity" – that is, the quantity of money economic agents they want to hold in cash, to pay for transactions and to protect against uncertainty. As interest rates rise, the opportunity cost of holding cash that earns no interest increases, implying that rising interest rates should generate higher levels of savings (i.e., the investment of cash in interest paying accounts and securities). The IS part of the model tries to describe the so-called "real" side of the economy (e.g., the amount agents want to save compared to the amount agents want to invest). As the interest rate increases, the desired savings level should increase, but at the same time the desired level of investment should decrease

(because the cost of borrowing money is higher). The IS-LM model, with the intersection of the IS and LM curves, implies that there is an "equilibrium interest rate" that will balance desired holdings of cash with desired savings and investment levels (macro textbooks usually don't specify which interest rate they are talking about; rather, they just assume away the complexity of the real world and just call it "the rate of interest"). This simple interest rate is essentially the return demanded by savers to forego spending a given amount of money for one year. For example, ask your brother how much you would have to pay him in order for him to forego spending $100 on new Lego for one year.

The economy is assumed to rapidly adjust to a new equilibrium in response to changes in the interest rate, which in turn can be caused by the actions of government (via changes in the money supply, the exchange rate, and/or the amount of government spending and/or the tax rate) or the private sector (e.g., changes in people's desired level of savings or investment).

Either explicitly or implicitly, government is assumed to have a fairly strong ability to manage the level of interest rates and aggregate demand (total output) in the economy, through a combination of monetary, fiscal (tax and spend) and, occasionally, structural (e.g., regulatory and trade) policy. In essence, macro textbooks often treat the private and foreign sectors like sheep.

Discussion of inflation in many macro textbooks is relatively simplistic, and can be easily summed up. Inflation is driven by some combination of excess demand relative to supply ("demand pull") of a given good/service, or price increases by suppliers due to their expectation of future inflation ("cost push"). Inflation is also usually described as a monetary phenomenon, using the "MV=PQ" equation. "P" is the average level of prices, and "Q" is the real volume of output; hence, "P x Q" or "PQ" is the nominal value of output, or GDP. The first derivative of these (i.e., the annual changes) gives the annual change in nominal (so called "current dollar" GDP). On the other side, "M" refers to some measure of the supply of money in the economy, which tends to be a

fuzzy concept. For example, the narrowest measure of the "money supply" equals currency and coins plus deposit accounts against which checks can be written. A somewhat broader measure also includes other very liquid assets, such as short-term time deposits and money market funds, which can quickly be turned into cash that in turn can be used to buy goods and services. "V" refers to the "velocity of money", or how often the money supply turns over, so to speak, in a given year in order to purchase the nominal value of GDP (i.e., PQ). In essence, "V" is a residual, calculated by dividing PQ by M. Getting back to inflation, the basic textbook assumption is that V is relatively constant, and Q doesn't vary much in the short term. Hence, changes in M that are greater than the growth in real output (Q) should feed through relatively quickly to changes in P, or the rate of inflation. As for deflation, macro textbooks are relatively silent about it.

Again, finance (e.g., banking, bond markets) doesn't play much of a role in the macro textbooks.

Macro textbooks are also notably silent as to the processes by which millions of individual micro decisions combine to produce the macroeconomic aggregates we observe.

<u>Once Again, And Even More So Than in the Case Of Micro, Real World Macro is Very Different from the Textbooks</u>

Clearly, the distinction between risk and uncertainty applies even more strongly at the macro level, where you are hard pressed to find any decision that could be termed "risky". Virtually all macro policy decisions must be made in the face of uncertainty.

As I said, the "components" or "National Accounts" approach to describing the different components of GDP (Gross Domestic Product) is fine as far as it goes. But textbooks usually don't take this discussion far enough. The first mistake they make is to omit the critical difference between stocks and flows. Think of a waterfall flowing

into a pool, with a stream flowing out of it. In this model there are two "flows" – the volume of water passing over the waterfall and into the pool and the volume of water passing out of the pool into the stream. Flows are measured in amounts over some period of time. There is one "stock" in the model – the level of water in the pool. Stocks are measured at a single point in time. Let's apply this to economics. The annual rate of growth of GDP is measured over a period of time, and is therefore a flow. Confusingly, this flow is typically expressed four different ways, depending on (a) whether it is expressed using money or a percentage, and (b) whether the impact of inflation – the increase in the average level of prices – over the GDP measurement period is taken into account. Here's a concrete example: in 2011, US GDP was about $15 trillion in nominal terms (taking changes in the price level as well as the amount of real output into account), or about $13 trillion in "real" terms, which attempts to measure only output over the preceding 12 months, but not the change in price level (obviously, these measurements can only be made with a certain degree of uncertainty). When expressed in percentage terms, "GDP Growth" measures how much larger (or smaller) money GDP was in one period compared to the previous period of the same length (e.g., for the 12 months ended 31Dec2011 compared to the 12 months ended 31Dec2010). In 2011, US nominal GDP grew by 3.9%, while real GDP grew by 1.8%.

Now let's go back to the National Accounts approach to GDP, which, as I hope you will see, enables you to generate some very powerful insights about what is going on in the economy. Unfortunately, too few people, including many financial services industry professionals, have a good grasp of this model. Recall that GDP or total output, can be broken down into the private sector's spending on consumption and investment, the public sector's spending on consumption and investment, and the external or current account balance, which is the net of exports (the goods and services we sell to foreign buyers) less imports (what we buy from them). All of these can be expressed in dollar terms, or, more commonly, as percentages of GDP (note that this is not a growth rate; rather, it is akin to shares of the GDP pie).

Now let's take this framework another two steps. Private sector savings is equal to GDP less private sector consumption. The so-called private sector balance equals private sector savings less private sector investment. The public sector balance is equal to public sector spending on consumption and investment less the taxes collected from the private sector. And now we come to a very important point: the private sector balance plus the public sector balance MUST EQUAL the current account or external balance. Why? Because if your domestic consumption and investment is greater than the domestic output produced by the economy, you must be importing more than you export (and vice versa). Here is a concrete example: in 2011, the US private sector balance as a percent of GDP was 6.5%. That means that private sector (household and business) savings exceeded private sector investment spending by an amount equal to 6.5% of GDP. Also in 2011, the public sector balance equaled a negative (9.6%) of GDP. In other words, public sector consumption and investment spending exceeded taxes collected by an amount equal to 9.6% of 2011 GDP. Adding the private and public sector balances together gives you negative (3.1%) which is the amount (expressed as a percentage of 2011 GDP) by which domestic consumption and investment spending exceeded the domestic output of goods and services produced by the economy.

So far, so good, I hope. Now for a big insight: the fact that the private + public = external balance equation by definition MUST BALANCE means that a change in one balance must be offset by a change in one or both of the other two balances.

Now let's move on to the next big insight: the meaning of negative and positive private/public/current account balances in the context of flows and stocks. Remember this: a negative balance means that a sector is a net issuer of financial claims to the other two sectors. Put differently, a negative balance must be FINANCED by issuing, in broad terms, debt and equity to the other two sectors. The opposite of this is obviously that a positive balance means that you are a net lender to/investor in the deficit sector or sector. That is, you are providing the funds to finance the deficit in the other sector, and adding to your stock of financial claims (i.e., debt and equity

securities) on the future output of the deficit sector (or future tax revenues, in the case of the public sector).

On to the next big insight from our National Accounts/Stock-Flow model – what happens if a surplus sector – that is, a sector with a positive balance – doesn't want to keep accumulating additional claims on a deficit sector (say, because they suspect the other party may never make good on them)? Ask the Greeks. If nobody will buy your claims, your sector balance has to return to a zero balance or surplus very quickly, which is almost always extremely painful and disruptive for an economy. Here is the critical point: most of the economic statistics reported in the media and taught in "Intro to Economics" courses are about FLOWS. Rarely do we focus on sector balances and even less on changes in the STOCKS of financial claims held by different sectors, and what those stocks imply for future sector balances (in particular, the future ability of deficit sectors to continue to run those deficits). One of the biggest reasons the 2008 global financial crisis took so many people by surprise was that too few alleged experts were focusing on the changes in the stocks of financial claims and the willingness of different parties to keep accumulating them. As you can see, in macro as in microeconomics, finance is critical, and Intro textbooks and teachers do a great disservice by assuming it away and not acknowledging the role it plays.

Let me try to pull the insights we have discussed up to now into one final important point about macroeconomic dynamics (or the evolution of various macroeconomic measures/statistics if you prefer). We know that the private + public = external balance equation is an "identity" that must balance. We also know that, in practice, changes in the stocks of financial claims on the three sectors can act as a constraint, and sometimes trigger a crisis that forces rapid adjustment in the three sector balances. The final point is that these rapid changes in sectoral balances are frequently accompanied by collapses in the rate of GDP growth. Let's look more closely at why that is often the case.

To simplify the story, in 2008, other sectors decided that they wanted to drastically reduce the rate at which they had been accumulating claims on the US private sector (and in particular, the household and financial institution subgroups of the US private sector), because of growing fear that a significant portion of these claims would never be repaid (e.g., because of defaults on home mortgages and failing banks). In the absence of the ability to run a negative sector balance financed with the issuance of debt claims, households sharply cut back on their consumption spending and their spending on new houses. In turn, this led to a sharp reduction in private sector investment in residential structures (i.e., housing). Businesses, seeing those spending cutbacks by households, cut their spending on investment (no need to add a new factory if households are cutting their spending), and started to lay off workers to conserve their cash as their uncertainty about their future sales and revenue increased. The net result of all these changes was very substantial change in the US private sector balance, from a significant deficit to a significant surplus. The other result was a fall in total demand by the private sector, due to less spending on consumption and investment (and a rise in private sector savings, which equals GDP less private sector consumption). As we know, a change in one sector's balance must be offset by changes in the other two sectors' balances. In the case of the United States, both the public sector and the external sector balances changed. The deficit on the latter significantly declined, as American households spent less money consuming imported goods.

At the same time, the public sector deficit increased, for two reasons. First, spending on some government programs automatically increases when the private sector encounters problems – examples include increased spending on unemployment, healthcare, and food stamp benefits. Second, the government chose to enact new spending programs. Why? In order to prevent an even larger reduction in GDP/total output. Think of GDP in 2007 as equal to 100, which was composed of 70 in private consumption spending, 10 in private investment spending, 25 in government consumption and investment spending, and a net negative (5) in exports sales less spending on imports. In 2008, we know that private consumption and investment

spending fell sharply – let's say to 60 and 5, respectively. Let's also say that imports fell, which caused the external balance to be a negative (2). If government spending didn't change, total 2008 GDP would be 60+5+25+(2) or 88 – a very sharp decline from the 100 in 2007, which would have, in a vicious circle, quite likely triggered further reductions in private sector consumption and investment spending, and therefore GDP growth. By increasing its own spending, government was able to limit the fall in GDP growth. From a stock and flow perspective, this meant that the public sector deficit balance increased, and other sectors were asked to hold a larger number of claims on future government tax revenue (that is, they were asked to hold more US government debt). At some point, the stock of those claims may reach a point that other sectors are no longer willing to hold them (which is what happened to Greece). For that reason, in the medium term it is important to reduce the size of the public sector deficit. But in the short term, government leaders faced a choice of either not increasing the deficit and watching economic growth collapse and human suffering (e.g., due to unemployment) increase substantially, or increasing government spending and the public sector deficit to moderate the extent of this collapse. This is a critical point – people who claim that the recent increase in the size of the US government deficit is the cause of our current problems lack a clear understanding of how the macroeconomy functions, and particularly how different sectors interact with each other. Rising government deficits did not cause the fall in private sector growth; rather, they were a natural and justifiable reaction to a fall in private sector consumption and spending that was caused when we reached the limit on the stock of private sector claims other sectors were willing to hold.

That being said, I can also envision another scenario in which rising government deficits would be the root cause of an economic problem. Assume we have an economy where growth is ticking along nicely, and the public, private, and external deficits are all at zero. Let's say a new government is elected, which dramatically increases spending but not taxes, causing a sharp rise in the public sector deficit. We know that either the private sector balance and/or the external sector balance must adjust to accommodate this deficit, through some combination of the private sector

moving into surplus (i.e., reducing consumptions and/or investment spending, resulting in job losses and higher unemployment) or the external sector moving into deficit (due to higher imports or fewer exports, which would also increase domestic job losses and unemployment). Hopefully these two examples have helped you understand a critical point – it is wrong to say that government deficits are ALWAYS wrong or ALWAYS right. That may be good ideology, but it is bad economics. In the real world, you have to examine each situation on its own merits to understand the dynamics that are involved when private, public and external balances change.

By this point, you should have already gleaned that the assumption taught in many Intro courses that the macroeconomy is usually in equilibrium, and that departures from this state are rare and quickly corrected, is pure fiction. Disequilibrium in the macroeconomy is the normal state of affairs.

Inflation and deflation are also more important and complicated in the real world than they are in textbooks. Let's start with the MV=PQ equation, and what it neglects to take into account. First, Q can actually change rather quickly in the short term. Second, V is not constant – for example, when people are uncertain about the future, they will hold higher precautionary liquid savings, which, all else being equal, will reduce V. Third, and perhaps most important, the concept of "M" neglects the critical role of credit. People buy financial assets, and goods and services, using credit (i.e., "leverage") as well as cash. For example, increases in the availability of credit (i.e., mortgage loans) to individuals helped drive up the price of houses before the 2008 crash. Similarly, increases in the availability of credit to investors helped drive up the price of financial assets.

This highlights another important point – while textbooks assume that inflation is a phenomenon limited to the market for goods and services (i.e., "Q"), it also occurs in the market for financial claims. In fact, in the run up to the 2008 crash, inflation in the market for goods and services was being held down by the impact of globalization, which expanded the effective supply of goods and services relative to demand, while

lowering their prices because of lower wages in some countries (e.g., China) as well as rising labor productivity in the developed world (e.g., thanks to better use of information technology) that was not matched by rising wages (e.g., due to workers' fear that their jobs would be outsourced to a lower wage country). However, at the same time, money supply growth and expanding credit were causing historically large increases in the prices of many financial and real assets, like houses and mortgage backed financial securities.

As previously noted, deflation is hardly mentioned in many macro textbooks. Yet, as the history of Japan since its 1989 crash vividly shows, it can play a critical role in the real world. Let me start by saying that the dynamics of deflation remain somewhat mysterious, even to experts. What we do know is that credit can play a key role in a deflationary cycle. As noted above, an economy can reach a point where lenders become so uncertain about the prospect of getting repaid in the future that they drastically reduce the supply of credit, precipitating a sharp fall in private consumption and investment spending, rising unemployment, and falling household income. Assuming that the supply of goods cannot be sharply reduced in the short term, producers will usually cut prices in an attempt to make some sales and generate some revenue – e.g., to pay their workers and their loans. However, in the face of falling prices, consumers may well put off buying today, assuming that goods will be cheaper tomorrow. This only reinforces the deflationary cycle. Meanwhile, on the financial side of the economy, rising unemployment and falling business sales lead to rising bankruptcies and write-offs of uncollectable loans and bad equity investments by investors. As this shrinks their own capital, the potential supply of credit they can extend to households, businesses and governments also shrinks, further reinforcing the vicious deflationary cycle.

To be sure, central banks attempt to fight this cycle by increasing the money supply ("M"). However, in a period of deflation, increases in "M" can easily be offset by decreases in "V", as people hold onto their cash rather than spend it, as it will buy more in the future if prices continue to fall. In short, and as Japan has very painfully

demonstrated over the past twenty plus years, once a deflationary cycle gets underway, it can be very hard to end. Hence, in the real world, policy makers generally prefer to err on the side of higher than desired inflation, due to their very justified fear of getting caught in an extended deflationary cycle.

As you can see by now, government has an important role to play (via changes in its spending and tax policies, and in monetary policies which influence the level of interest rates and the rate of inflation) in offsetting the impact on economic growth caused by changes in the private and external sectors. However, and this is critical to remember, a government's economic policy decisions are always made in the face of a substantial and irreducible level of uncertainty about how other sectors will react to them, and the nature and timing of the causal processes involved. Government policy makers forget Murphy's Law at their – and our – peril.

This gets us to the point noted above, about the silence of most economics textbooks on the critical subject of how all those micro decisions get aggregated up into macro level changes in aggregate demand and sectoral balances. Addressing this issue requires me to introduce you to another concept: complex adaptive systems. A system can be thought of as a collection of agents and various feedback loops between them. Some feedback loops are positive, sending signals to an agent to continue or increase a given behavior. Others are negative, telling an agent to reduce or eliminate a behavior. A complex system is one in which there are multiple positive and negative feedback loops between multiple agents. Some of these operate linearly, and without time delays between cause and effect. Systems in which linear, "real time" feedback loops dominate are those whose future behavior is easiest to predict (and, in turn, influence). In other cases, however, feedback loops operate non-linearly (e.g., where 2 unit change in an independent variable produces a 4 unit change in the independent variable, but a 3 unit change produces a 10 unit result), and with time delays (e.g., if you put your hand on a hot stove, you get rapid negative feedback; however, the feedback from eating too much and exercising too little, as measured by changes in your health, takes place over years rather than milliseconds). As non-linear and time-

delayed feedback loops comprise a greater share of the loops in a complex system, its behavior becomes increasingly difficult to predict and control. In such systems, it is impossible to predict the behavior of a complex system simply by understanding the rules followed by a few of its agents. In economic terms, understanding microdynamics (the behavior of an individual or company) does not lead to understanding macro. In complex systems, macro level behavior is said, "to emerge" from the interaction of individual agents and the operation of the multiple connections between them.

To make explanation, prediction, and control even more challenging, in a complex ADAPTIVE system, the individual agents have intelligence, and are able to adapt their decision rules and feedback loops (and sometimes even the goals they are pursuing) on the basis of how different strategies perform over time.

To be sure, macroeconomists have long attempted to predict the behavior of the complex adaptive system that is our global economy. A discipline of economics called "econometrics" is dedicated to this quest, and historically has used the tools of regression analysis to try to explain and predict the behavior of certain macroeconomic aggregates (e.g., GDP growth). This technique is, at heart, quite straightforward: choose a dependent variable (e.g., GDP growth) and a group of independent variables that theory suggests should explain its variation over time (e.g., the level of private sector and public sector spending). Then use linear or non-linear regression to estimate an equation (i.e., an algorithm) that relates changes in the independent variables to changes in the dependent variable. In practice, however, over the years this approach has repeatedly failed to achieve an acceptable level of predictive accuracy.

In broad terms, in the face of real world forecasting problems, the good ship of econometrics has tended to run aground on two sets of rocks. The first is the adaptive nature of the economy, which over time causes the assumed relationships between

independent and dependent variables in an econometric model to become inaccurate. Put differently, explaining what happened in the past has proven to be much easier than predicting what will happen in the future. The second problem is one of reductionism. In most cases, the independent variables used in macro models are themselves emergent results. For example, when an econometrician says he or she is forecasting GDP growth in part on the basis of his or her estimate of future household consumption spending, the obvious question is the source of that estimate. Usually, the answer is another regression model, using another subgroup of estimated independent variables, which themselves beg the question of the source and logical justification for the estimates used. At some point, this line of questioning runs into unpassable bog created by the complex mass of feedback loops between adaptive agents. To be sure, another branch of economics is using one of the key tools of complexity science – agent based modeling – to attempt to unravel the complex dynamics at work. But while this approach holds great promise, it remains in its intellectual infancy and is still very much on the fringe of mainstream micro and macroeconomics.

A Final Point on Performance Metrics: The Holy Trinity of Effectiveness, Efficiency, and Adaptability

Particularly at the micro level, but also in macro, you are likely to encounter a discussion of the metrics used to measure a firm or an economy's performance. Your initial reaction is, if it is like mine, complete confusion over the profusion of terms and calculations. One of the great insights in my life was the realization that, when it comes to performance metrics, there are really only three categories, which are as applicable to a single cell bacteria as to a hunter/gatherer tribe eons ago on the East African plain, or to a modern corporation or government today.

Effectiveness measures the extent to which the results you have achieved match your goals. Setting these goals is another matter, but that is the province of leadership and strategy, rather than performance measurement. Suffice to say, that in an evolutionary

context, goals should align closely with short and long term selection criteria (that is, the measures that are used to determine which agents live and reproduce, and which die).

Efficiency measures the amount of resources used to obtain the results achieved.

Adaptability measures the change in Effectiveness and Efficiency for a given level of change in the agent's environment. In an environment that never changes, this metric is meaningless. But virtually all environments change, at least over some timescale. Also note that, when it comes to adaptability, success tends to carry with it the seeds of later failure. One of our great failings as human beings is our reluctance to acknowledge the full implications of living in a world of complex adaptive systems. Due to their constant evolution, the causes of yesterday's success are impossible to fully understand, and unlikely to be replicable to the same extent in the future. Put differently, we must be cautious about drawing unchanging lessons from the past, particularly about cause and effect relationships. And yet this is precisely what our pride and natural cognitive biases tempt us to do. Few people or organizations are willing to fully acknowledge the extent of their own uncertainty (which, after all, is related to the degree of fear we feel), or the role luck (more technically, randomness) has played in past successes. And so we naturally try to succeed again in the future, using the approach that worked in the past, with frequently disappointing and occasionally fatal results.

Perhaps the best antidote to this tendency is to subject any plan we make to a so-called "pre-mortem." Assume the plan has failed miserably, and list the factors which caused the failure. Then adjust the plan accordingly, for example, by adding contingencies or holding resources in reserve.

Note that effectiveness, efficiency, and adaptability are not necessarily financial or economic measures. In fact, since financial results (e.g., the profit earned by a company over a year, or the amount by which GDP grows over the same period) are

only produced with a time lag, they are notoriously poor indicators if you are trying to manage an organization. In that situation, you need to seek out, or devise, leading and real time indicators, that provide you with more timely performance information.

Note also that there are tradeoffs between these three performance measures. For example, adaptability usually requires some degree of resiliency (the ability to absorb surprising changes without great disruption or failure) – e.g., the troops a general holds in reserve to deal with the unexpected surprises that occur in every battle and war. However, efficiency sees resources that are not currently used to pursue results as waste, and seeks to eliminate them. For this reason, many organizations that are admired for their supreme efficiency later prove to be unadaptable and fail when their environment significantly changes. There is no algorithm or optimal way to make tradeoffs between the goals you set for effectiveness, efficiency, and adaptability. In complex adaptive systems, leadership will always remain an art. That said, in my experience the best leaders have been the ones who recognized the inescapable limitations of forecasting and planning in a complex adaptive system, and consequently sought to maximize the resiliency and adaptability of their organizations, subject to achieving the levels of effectiveness and efficiency needed to escape selection – i.e., to survive over a given period of time.

Finally, Never Forget the Intellectual Heritage of "Economics"

The area we today call "economics" was originally studied as one aspect of moral philosophy, which is "the branch of philosophy that involves systematizing, defending, and recommending concepts of right and wrong behavior." In the 19th century, "political economy" was separated from moral philosophy, and focused on the production and distribution of wealth at the level of the nation state. In the late 19th century, the study of "economics" was separated from "politics". The 20th century saw the study of economics further broken down into micro and macro, and the increasing dominance of mathematical approaches, as economics departments at universities did their best to imitate their fellow professors in the physics department.

Unfortunately, as I hope I have shown in this paper, economics will never be a so-called “hard science” like physics. The latter is governed by unchanging natural laws and experiments that can be replicated. In contrast, economics will always confront the irreducible uncertainties and emergent phenomena that are produced by the interaction of cognition and emotion, both within and across individuals and organizations. Just as important, economics will never be able to escape the “moral philosophy” questions that are inherent in the allocation of scarce resources and the distribution of wealth in a society. As you embark on your study of economics, never forget that it will always be an art as well as a science, and that the insights you generate and recommendations you make can have large and important impacts on the world in which you and your children will live.

Good luck!