The Effect of Changes in Economic Growth on Stock Prices

HCR-Lecture 2002, Prof. Paul M. Romer, Ph.D.

Introduction

In this lecture, I will focus on a somewhat narrow topic, the connection between changes in economic growth and changes in stock prices - even though most of my research has been on the forces that might cause changes in economic growth. My focus will be narrow in the sense that I will not ask about where those changes in economic growth come from. I will merely ask what effect it would have on stock prices if we were to see an increase in the rate of economic growth. The answer that I will give is one that I think may surprise some of you: What I will try and convince you of is that faster economic growth - for example caused by faster technological change - will not lead to an increase in stock prices measured in terms of a standard measure like the price earnings ratio. At the end of the lecture I will talk a bit about the question of what - if not economic growth - might explain the very pronounced prices and price earnings ratios we saw in worldwide asset markets in the late 1990s. Despite all prevalent counter-arguments, I think the explanation is that we were going through a financial bubble, and I will give some insight in the implications of that.

But the deeper issue behind this topic has to do with how economists use mathematics, how it is that we use the formal language of mathematics to help us reason about the world. Those who are familiar with economics know that it now is a very mathematical subject, that we rely heavily on mathematics, that it is the language in which we as economists speak to each other. It is so much a characteristic of economics that we even have jokes which are all based on the idea that economists use too much mathematics. Yet, the issue which many people have some discomfort about is whether or not the use of mathematics has actually aided our understanding of the world. What becomes clear if we look from a broader historical perspective in the history of economic thought is that our transition to the use of math has been a very recent one, and a very sharp and dramatic one, and in some ways a very turbulent one, too. It is interesting to think about this transition from the perspective of a German university because the German tradition of economics was one which had deep roots in the more verbal, institutional, contextrich, historical tradition of economics. I think this is a tradition or an approach which modern economics is now returning to from a digression or an exploration of the world of mathematics.

My own work is heavily mathematical and my understanding of technological change is deeply influenced by the use of mathematics, so I am a big believer in mathematics as an aid to understanding the world. But the deeper message that I would like to convey is that math is productive in science - social science as well as physical science - only if we use it in two important ways. First, we have to bring the abstraction of math back to the world: we connect the conclusions of math with what we see in the world. This is where the formal traditions of economics in the last half of the 20th century are reconnecting with the earlier historical and institutional traditions of German economics. So bringing the abstraction back to the world is one critical part of that process. The second essential part of using mathematics is that we must take it serious, we must believe in our equations, and take the applications seriously. When we map back to the world, we must ask whether they actually describe what we see in the world.

One danger with math is that it can be used almost as sport. It is a way to compete with each other, to try and beat a colleague in math just the way you beat him on the tennis court. And it can become in that sense a kind of a parlor game, a showing off. One of my colleagues, Brian Arthur, described mathematics as being like antlers on elk or moose, that people use them to butt heads with each other and try to see who the alpha male is. That use of mathematics should be criticized, and I think it is sometimes not criticized enough. The flipside is that if we take it seriously, I think math can help us improve our intuitions and deepen our understanding. What I will try to convince you of in this lecture is that if you take some relatively simple mathematics seriously, you can disprove the intuitively plausible conjecture that fast technological change should lead to high price earnings ratios. And I hope along the way I will even show you where intuition leads us astray in this issue.

Now, let me focus on the narrow topic: The conjecture is, if growth is faster, does it therefore follow that price earnings ratios will be higher? I want to start with a quote from Paul Krugman, the most recent recipient of this prize[1]. It is fortunate for us as economists

[1] Paul Krugman is the year 2000 recipient of the biennial H.C.Recktenwald-Prize.

and people interested in the history of thought that Paul writes a weekly column in the New York Times, because this gives us a written record which is like the conversations that economists have at the lunch table. And that lets us refer back to things that we were saying to each other. So I will use a quote from his column as representative of what very thoughtful sophisticated economists were saying as of the 1990s or early in this new century. The quote is one where he is discussing a book by Robert Shiller[2] where Shiller argues the case that in fact the high stock prices represented a bubble. And what Paul says is that he is sympathetic to this conjecture but he is not entirely convinced. The key sentence is in the middle, it says that some marks of a bubble are plain to see, but so is the spectacular pace of technological progress:

> *"Mr. Shiller believes that the whole stock market is inflated by a speculative bubble...I'm sympathetic but not entirely convinced. The social and political hallmarks of a bubble [...] are plain to see, but so is the spectacular pace of technological progress. I'm not sure that the current value of the Nasdaq is justified, but I'm not sure that it isn't."*
> *Paul Krugman, Feb. 2000"*[3]

What he is setting up is a kind of opposition. He states that it might be driven by an - at least temporarily - self-fulfilling expectation of higher stock prices, a phenomenon we call a bubble, but it could also be faster technological change that is causing the higher price earnings ratios. This is representative of things that many economists and many commentators were thinking and saying. I will try to argue that if we use our mathematics and our models carefully, we can see that that second intuitively plausible conjecture does not follow logically. That is the setting and the agenda for this lecture.

The Basic Model

To set up the model, we have to think about what is involved in making a mathematical model of stock prices. The first thing we have to address is that asset prices involve time. They inherently involve thinking about things that go into the far future. And to help us understand the implications of the factor time, let us think about a simple case where there are some flows of income in the future, and let us see if we can put a price on those. Once we have done that, we will transfer that to thinking about stock prices. So, here is the situation: You just won the million dollar lottery - that is the good news. The bad

[2] Robert Shiller, *Irrational Exuberance*; Princeton, N.J.: Princeton University Press, 2000.
[3] Paul Krugman, *The Ponzi Paradigm*; in: NYT, 12.03.00.

news is that it is a dollar a year for a million years. So, the question is: What would you pay to have this stream of income? Or, if you were given this and you wanted to sell it to a friend, how much would your friend be willing to pay for a dollar a year for a million years? Your friend could presumably turn around and sell it to someone else later in time, so we have to think of this series of payments going on into the future. Now, surely it is plausible that you would not pay a million dollars for this million dollar prize, but what would you pay for it?

In the following, I will use the letter p to stand for the price[4]. And - just to generalize slightly - instead of saying it is a dollar a year, let us suppose that it is c dollars per year. And in the background, you can earn r as an interest rate on investments in the bank, and we would like to see: how much would you pay for this stream of income going out into the future? Now, what you could do instead of buying the million dollar prize is you could put your p dollars in the bank, and then you could earn $r \cdot p$ as interest every year. That is an alternative to receiving the winnings of c dollars every year. And the price where you should be willing to exchange your annual winnings for a unique payment that you could put in a bank, is one that equates these two streams of income, that is $r \cdot p = c$. You will be indifferent between getting the p dollars and putting them in the bank or receiving c every year. So, we can set up a very simple equation which is the basic discounting formula that we use in economics:

$$p = \frac{c}{r} \qquad (1)$$

If you have a stream of income that pays c dollars each year, the price that you would be willing to pay for it is c divided by the interest rate r. So if we take an example, suppose that interest rates are 7%, c is one dollar, then you might be willing to pay about 14 dollars to buy this lottery prize from a friend, or a friend might be willing to pay you 14 dollars to purchase it. This is one very simple illustration of how we can use a little bit of mathematics, a little bit of formal reasoning, to solve what at first seems like a somewhat complicated question.

Unfortunately, when we think about stock prices it is not just that stocks pay us income in the future, but they pay us a stream of income which grows over time. This

[4] cf. p. I, Index of Symbols.

complicates things slightly. So let us now imagine that there is a different lottery prize which grows at a rate that I will call g, and that might be for example 3%. So, if c is one dollar, this year I earn one dollar, then next year I might get a dollar and three cents, the following year I might get about a dollar and six cents. We now have to figure out the price for this stream of income which is increasing over time. So again, we will start with the price p. And you might be willing to accept that it is money in exchange for this stream of income which is growing over time. You put it in a bank, after one year you will have $(1+r)\cdot p$. The difference now is: As your annual income from the prize would rise every year, you have to adapt the stream of earnings you would get if you just left the money on the bank in a way that your "bank account" would grow at the same rate. Instead of taking out r, you only take out the portion that corresponds to the difference between r and g. So, if you put one dollar in the bank, you take out 4 cents as income, but you leave the other 3 cents. Thereby, you start again in the next year with $(1+r)\cdot p$, then minus the amount you take out, which is $(r-g)\cdot p$. What you are left with is $(1+g)\cdot p$ next year[5]. So, again, you take $(r-g)\cdot p$ out of your bank account every year, and the amount of money left on the bank account will grow with the rate g over time. And then you can have a stream of payments $(r-g)\cdot p$, a balance which is growing.

Now what I will compare to that is c in the first year, $(1+g)\cdot c$ in the second year, $(1+g)\cdot(1+g)\cdot c=(1+g)^2\cdot c$ in the third year, and so on. So, c is growing with the rate g. But the money you take out of the bank account is also growing at the rate g, so you take out $(r-g)\cdot p$ in the first year, you can take out that times $(1+g)$ in the second year, and so on in each subsequent year.[6]

Therefore we should equate the initial payment c with $(r-g)\cdot p$, rather than with $r\cdot p$. Then we can solve for a related formula to the last one, which is our discounting formula for streams of income that grow at a rate like g:

$$p=\frac{c}{(r-g)} \qquad (2)$$

---

[5] $[(1+r)-(r-g)]\cdot p=(1+g)\cdot p$

[6] cf. p. II, Annotation 1.

So to use another example, if r is still 7%, but g now is 3%, and c is a dollar, then the price will be a dollar divided by 0.04, or 25 dollars. So, already you can see the intuition that lies behind much of the thinking about growth: A stream of payments which is growing over time is more valuable than a stream of payments that is not - nearly twice as valuable, here.

Application to the Corporate Sector

So let us relate that to the corporate sector. Let us imagine that Y stands for all the outputs in an economy. Let earnings be a fraction α of the total output in an economy and new investments that firms undertake be a fraction s. So, a firm has a certain amount of net earnings, but it cannot pay all of these earnings out to the shareholders for them to go out and use for consumption purposes. It must take some of these earnings and reinvest them to keep the business growing.

The Free Cash Flow, the stream of payments which you might be able to take away from a corporation would be the earnings minus the new investment. So c, the thing that would stand for cash, the c that I used in the formula before, will mean earnings minus the net new investment that the firm does. And what we want to do is value this stream of income over time under the assumption that it grows. So, I will keep the two previous formulas. We need that the cash payment now is $c \equiv (\alpha - s) \cdot Y$ and - notice that we are going to assume that that quantity is positive - α must be bigger than s.

Let us stop for an instance. This is one of those places where you have to think, "Do I take the mathematics seriously?" What would it mean if this were true? What would it mean if α was less than s? If I owned a corporation, and every year it generated a certain amount of earnings - say, 7% of its output -, but I supposedly had to reinvest 10% of the output to keep the business growing, that would mean that all I do is feed money into this corporation and never take anything home. Because I own this corporation, I get to put money into it every year, and for the infinite future. It is a little bit like if you owned a boat, for example. Someone described a boat as a hole in the water that you throw money into. So that is what the corporation, the entire corporate sector, would be if s were bigger than α. That cannot be right. So, if the corporation is valuable, you get something from owning it. It must kick off some free cash that we can take and spend and use. So I will use the assumption that α is bigger than s, because anything else would not make any sense.

I will also use the fact that there is a proportionality between earnings and total output in an economy: $E = \alpha \cdot Y$. I will include that into my basic formula for discounting streams of income into the future, formula (2). And I will substitute for Y earnings divided by α[7]. Then I can divide the term by earnings and get the kind of expression that I want to use, a constant price earnings ratio for the entire corporate sector:

$$\frac{p}{E} = \frac{(\alpha - s)}{(r-g)\cdot E} \cdot \frac{E}{\alpha} = \frac{(\alpha - s)}{\alpha} \cdot \frac{1}{(r-g)} \qquad (3)$$

This is the basic valuation framework I want to use for thinking about valuing stocks.

For example - and this is the same kind of calculation I did before - if you set $\alpha = 0.3$, and $s = 0.15$, if r were 7% and g were 3%, then you calculate a price earnings ratio of about 12.5. That is not too far off from historical norms until the last couple of centuries or more in the United States, but significantly lower than the values that we saw in the late 1990s. Now, let me keep clear about what I am going to use the math for here. We do not know the numbers α and s - and r and g, for that matter - precisely enough for me to tell you what price earnings ratios should be, so we cannot work out what the golden number for price earnings ratios is. But I should be able to use this formula to ask the central question of this lecture: If g changes, how does that change the price earnings ratios? So, you will not leave this lecture with an idea about what quantity the true price earnings ratio has, but you should be able to get some idea about how it changes.

When we make comparisons in the cross section, we stop at a point in time and look across a number of firms. We can see some firms that have earnings growing rapidly into the future, and others, where the earnings will grow more slowly. And if you think about this formula applying per firm, that gives you a relationship that many people are familiar with: Firms with rapid growth in their earnings should have high price earnings ratios, and this is something you see very dramatically if you look at stock prices. Rapidly growing young firms have price earnings ratios that are much higher than those of traditional firms with few prospects for growth. We actually can carefully draw that

---

[7] $p = \frac{(\alpha - s)\cdot Y}{(r-g)} = \frac{(\alpha - s)}{(r-g)} \cdot \frac{E}{\alpha}$

inference, because we can assume that r is constant in the cross section. In a particular point of time, the rates of return on things like bank accounts or alternative investments are determined for the economy as a whole, while g will change between different firms. And so, as we vary g, that is the only thing in the formula that changes. As g gets bigger, the bottom part of the fraction gets smaller, so the fraction as a whole gets bigger. So, we get this cross section result which I just described as being familiar to many people[8].

But now you can appreciate where I am going. If we think about a change over time that affects the entire economy - suppose that g increases for the entire economy, i.e. all rates of growth are going to be faster on average. Then you have to think about what determines r. When comparing, for example, the United States in the 1970s and the United States in the 1990s, we might think g is higher, but then we have to ask whether r potentially undergoes changes in time as well. Growth might actually have an influence on the rate of return. And I will try to convince you that r should change, and the direction in which it should change is that r also goes up as g goes up. In the bottom line, I will try to shed a light on the implications of that for the price earnings ratio.

Effects on the Interest Rate

We can examine the effects on the interest rate with a thought experiment. This is a technique I have learned from physics, and it is one I think people in the social sciences should use more, and particularly in economics.

Think about an idealized, very unlikely situation that is very extreme. Because if you go to the extremes, it is easier to think about what is likely to happen. Small changes are hard to understand, but extremes are sometimes much easier to understand. Let me give you one example of this. Many people who have not taken economics think that a firm always earns more income, net revenue, or profits, if it charges a higher price. Higher prices lead to more profits - sounds plausible. But try the following on your child: Tell your child that you control the entire supply of diamonds, or whatever you want, and ask your child: “How can you get the most money from selling off diamonds?” Well, you

---

[8] $\frac{\partial P/E}{\partial g} = \frac{(\alpha - s)}{\alpha} \cdot \frac{1}{g^2} > 0$; Assumption: *r* constant in cross section; Result implicates that higher growth rates lead to higher price earnings ratios.

could charge a Euro per carat, or a thousand Euro per carat, or a million, or a billion, or a trillion per carat. If you go to the extremes, at some point (if, for instance, you charged a trillion trillion Euro per carat), even your child will recognize you would not sell any diamonds at all. So something is wrong.

We can do that same kind of thought experiment for the valuation model here. Imagine that a spaceship lands from another galaxy and brings us technological progress that makes Moore's Law[9] look like nothing. It is not doubling every two years, it is doubling every week, growth rates just go skyrocking and out of control. It is not something that you want a bank on happening, but it is conceivable. And, what should nevertheless be the case is that humans should still be able to make rational positions about what to pay for things. They should not just go crazy when that happens. So, think about an increasing g, in particular think about g getting up close to r. The spaceship has arrived, and we are going to be having faster growth for the infinite future. And, let us assume that r stays constant, but g is getting bigger and bigger. If g finally gets close to r and the bottom of equation (2) goes to zero, that would say that prices go to infinity. People would be willing to pay an infinite amount to own just one share of any firm you can consider. For one share of Coca-Cola, you would pay a trillion trillion Dollars because g is so high, and every single share is worth an infinite amount. That cannot be right, people would not pay a trillion trillion dollars for one share of Coca-Cola. So, r has to adjust, that is r must increase along with g.

By this point, I hope you have got some kind of idea of what can go wrong with the usual intuition for the basic formula. And then some idea of what can go wrong if you think about a change for the economy as a whole, that when g changes, r can change as well. If that is the case, we have to examine which effect is bigger. If we look at modest changes of g - a 1% increase in g, for example - what is the net effect on price earnings ratios, if both g and r go up?

The rest of the analysis requires a higher level of math, so I will cross most of the formal details, but anyway I will try to give you some idea about how economists can use

[9] Gordon Moore, Co-Founder, President, and CEO of Intel, formulated prospects on the future increase in numbers of transistors on computer chips in 1965, assuming them to double every second year. This prediction of exponential growth is referred to as Moore's Law.

math to reason about this question. If you have taken economics courses, you might be familiar with a diagram like the following:

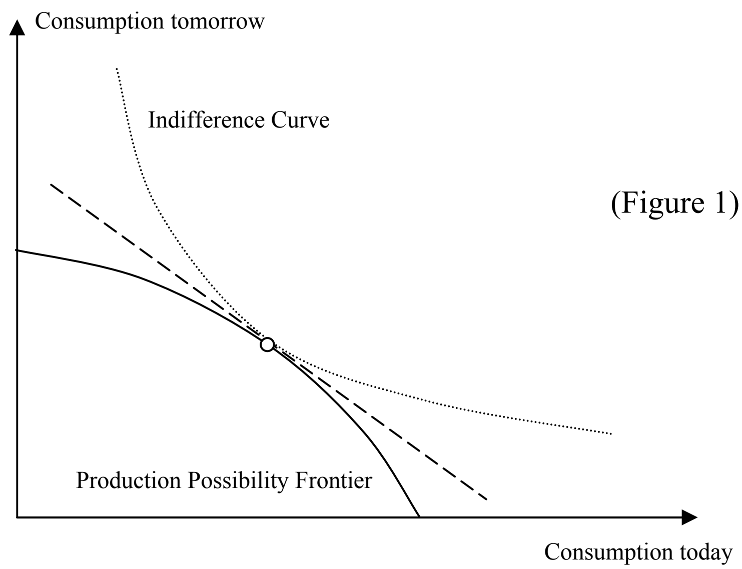

(Figure 1)

It tells you that a way to think about rates of return is to think about it as a kind of trade-off between what we have today and what we can have tomorrow. The first curve, the one that bounds the area towards the origin, describes the possibilities: The economy can have some more tomorrow if it is willing to accept less today. Invest more today, so you consume less today, but then you have more to consume tomorrow. That is what we call the production possibility frontier. The other curve which has the alternate curvature is named indifference curve and tells us something about how people would trade off more versus less tomorrow. Of course, they would like more of both, more goods today, and more tomorrow, but that is impossible. Their intertemporal preferences determine a rate at which they are willing to trade off between how much they have today, and how much they have tomorrow.

To figure out the interest rate that determines those trade-offs, we look at the slope of the flat line which divides the two regions. We take the point that determines our consumption today and our consumption tomorrow, and the slope at that point helps us understand what the rate of return would be. The usual tradition for modeling asset prices in economics is to look at the piece of the picture, where the flat line is tangent to the curves. Then, to measure the slope, one of those two curves is sufficient.

The Production Side - the Solow Growth Model[10]

First, we might employ a growth model. It is a model of what is possible and tells us what an economy can produce today versus tomorrow. We will try to find the slope by examining the corresponding point on that production possibility curve. I will skip the mathematical modeling of what this curve looks like, and try to make my point verbally. The slope I will find in that point will be our interest rate r. And the last step will be to see how the interest rate changes when people move along the curve.

This procedure is a standard exercise in what is referred to as the Solow Growth Model[11]. For those who are not familiar with this model, I will try to give a broad, rush overview over how we can use it. I will comment on various points, and, though not going into mathematical detail, try to explain the chain of formal reasoning.

People familiar with economics know the difference between net domestic product and gross domestic product. I am going to assume in my measure that for Y_t the effects of any depreciation are already removed, i.e. I will use the net domestic product. The basic equation of the Solow model that describes total output in an economy has three components. The first one, A_t, is the way we represent technological change, so A_t is a factor that grows over time (which is suggested by its subscript t). K_t represents the stock of capital and, finally, L represents the stock of labor. I will assume that L is constant, and I will not have anything more to say about it. I will just suppose that A_t grows over time, and that its growth speeds up.

By the way, that is exactly what New Growth Theory is all about: What might cause A_t to grow faster? - But in this lecture we will just assume it happens, and we will not ask where it came from.

The way an economy gets more capital stock K_t is that it does more investment. So I will assume that a fraction s of total output gets re-invested in the economy. We produce lots of things in our economy. Some of these things are DVD-players or TV-sets we can have fun with, others are things like blast furnaces and machine tools that are not any fun

[10] cf. p. II, Annotation 2.

[11] Named after Robert Solow, its main developer along with Curtis Swan.

today but that make things that are fun tomorrow. So a certain fraction of the total output of an economy is directed in these things that are inherently stocked.

Now, there are two results I will use here. The one apart from the expression for the price earnings ratio is that we can figure out the interest rate from what we call the marginal product of capital - and this is where it presumes some economics. We will just take the derivative of the production function, and that tells you what the interest rate should be in this economy. I can work out that the interest rate r in this economy should be a constant once things settle down - this is what we call the steady state. After a little bit of algebra[12], I can derive an operable expression for r:

$$r = g \cdot \frac{\alpha}{s} \qquad (4)$$

It includes α, the share of capital income in total output, which is actually something economists can measure. Furthermore, it involves the savings rate s, which we can also measure by seeing how much of the output is devoted to new investment, and the growth rate g.

Now we have to think back to what I said about α and s before: The whole model does not make any sense unless α is bigger than s, because $(\alpha - s)$ is the money that a corporation kicks off for me to go and do fun things with. With that background, we have an answer to the question of how much r will change. Before, we used the thought experiment to see that r would have to go up if g goes up, now we can actually put a number on it, or a magnitude. Not only would r go up, but r would go up by more than g goes up, and it will go up by this ratio α/s. And again, those numbers like α and s are economic figures we can actually measure.

If I combine equation (4) that tells you how interest rates change with the growth in the economy, and equation (3) which tells you what price earnings ratios are as a function of r and g, I get to my final equation. It shows that when the growth rate g increases,

[12] For the basic formal approach, cf. p. III, Annotation 3.

what should actually happen is that the price earnings ratios - according to this model - should go down[13]:

$$\frac{p}{E} = \frac{(\alpha - s)}{\alpha} \cdot \frac{1}{(r-g)} = \frac{(\alpha - s)}{\alpha} \cdot \frac{1}{\left(\alpha/s - 1\right) \cdot g} \qquad (5)$$

It is no question that a higher g is a good thing in an economy, because we get faster growth of consumption. But it turns out not to be good in the end, because it raises stock prices. It turns out that we will have a stream of dividends from the corporate sector, a stream of earnings that will grow faster, but we will also be more impatient. The r that we choose, the rate that we use to discount these streams of income, will be higher, and so price earnings ratios will actually fall.

So you get a fairly unambiguous answer to the question of what a rate of growth should do to price earnings ratios. It is one that goes exactly counter to the assumption that is implicit in much of the discussion that surrounded what was going on with stock markets in the late 1990s.

The Preference Side

Alternately, we might take the indifference curve. We model preferences, i.e. what people like, and how they feel about goods tomorrow versus goods today. We find a point on the curve, and we calculate a slope. We can do this in a quite elaborate form when involving uncertainty, and we can work out all the theory of asset prices on that basis.

Let me just explain this more traditional model of asset pricing very briefly. One disadvantage of using the production model I introduced before is that it is not easy to incorporate uncertainty in the mathematical framework - although you can do it. The expression for the interest rate we get from the analysis on the production side can be formally adapted. One can model preferences, utility people receive from consumption, in a way that consumption today is raised to some exponent $(1-\sigma)$, and so is

[13]Under the assumption that $\alpha > s$ (cf. p. 7): $\left(\frac{\alpha}{s} - 1\right) > 0 \quad \Rightarrow \frac{\partial\, p/E}{\partial g} < 0$ Result now implicates that higher growth rates lead to lower price earnings ratios.

consumption tomorrow.[14] And again that exponent is a number we definitely can measure.

What comes out of that kind of model is that if σ is between 0 and 1, the amount of savings in an economy will increase when growth does. In that case, it could be that r goes up by less than g. Now the question arises why that is. Until now I have been assuming that the savings s were about a constant fraction of our income. For an economy as a whole it is the same as saying that a constant fraction of total output in an economy goes into machine tools and blast furnaces. It could be that that assumption was wrong, it could be that when g goes up, s goes up as well, and then r is not increased by as much as g is. And that gives you the hope of a story where price earnings ratios truly rise.

The problem is that the evidence does not support either side of this implication. When we try and measure what σ is - and we have a variety of ways to do that - we usually conclude that σ is bigger than 1, not less than 1, and in that case, the implication actually does not follow. So, our understanding of preferences and σ suggests that this is not the way out of the puzzle. A similar impression arises if you look at what actually happened to savings rates. Did people save more in the United States or anywhere in the world because returns in asset markets were so much higher? There is just no evidence that savings rates and investment rates went up, particularly that savings rates from the preference side actually went up. So, even though there is a way to try and suggest that faster growth leads to higher price earnings ratios, the evidence argues pretty strongly against it. In that case, the best working assumption is that the parameter s is a constant. And therefore, finally, an increase in g will lead to a decrease in price earnings ratios.

Conclusion

So, where does that lead us? At the most narrow level of the question what may cause an increase in price earnings ratios as observed in the 1990s, I think that this should bias us strongly in the direction of taking seriously the hypothesis that it was a bubble. It was not driven by fundamentals, it was driven by a self-fulfilling set of expectations in

[14] cf. p. III, Annotation 4.

financial markets. So, if you take the theory seriously, it actually argues against the fundamentals explanation, and pushes it towards the bubbles explanation.

As an aside, there is also an interesting experimental on this. You take real people and let them train in a securities market that you control, using the kind of methodology that Vernon Smith, the 2002 co-recipient of the Nobel Prize, developed. Vernon pioneered experiments with real human subjects to see how they work in asset markets that we can control in the laboratory just the way a physicist could control a set of apparatus. What you find is that it is not that hard to generate bubble-type movements in asset prices. So, theory argues against the fundamentals explanation, and empiricism suggests that bubbles are quite plausible. I think it is an explanation which we should be taking more seriously, one that we should have been taking more seriously in the course of the late 1990s, I argue.

That is the narrow question, but the broader question is that if the theory is so relatively unambiguous and so clear, why did economists pay so little attention to that theory? Referring to the quote at the beginning, it was typical of discussions throughout the 1990s. In that time, even the skeptics who said that it was a bubble usually conceded the point that, if there were faster technological change, that might explain why price earnings ratios were higher. On the other hand, they attacked the assumption that we were indeed in a world with faster technological change. They said things like "well, there must be a bubble, because there really is not that much evidence for faster technological change". So, as we engaged in this debate, almost no-one took the logic in this model seriously, and this, I think, is a troubling sign. It is evidence that, also frequently, economists reasoned from their intuition and developed models as supporting arguments. But they did not sufficiently often go the opposite way and use the model to actually discipline their intuition, to check whether maybe their presentiment was wrong.

My conclusion about where we need to go as economists is that we should use simple models in the very first path, because simple models are the easiest to understand. And I think we should take those models very seriously.

Index of Symbols (in order of appearance)

p: price

c: cash flow (constant over time)

r: interest rate

g: growth rate (constant over time)

Y: total output of an economy

α: earnings rate ($\alpha Y = E$: total earnings in an economy)

E: total earnings in an economy

s: savings rate (sY: total savings/net new investments in an economy)

t: time index

Y_t: net domestic product NDP in year t (↔ GDP: gross domestic product)

δ: depreciation rate

A_t: technological change in year t

K_t: stock of capital in year t

L: stock of labor (constant over time)

Π: profit

w: wage

k: capital share in total income

σ: ???, $0 < \sigma < 1$

Annotation 1

Comparison of the two equivalent streams of income:

t	Annual Cash Flow	Unique Payment
0	c	(r-g) p
1	(1+g) c	(1+g) (r-g) p
2	$(1+g)^2$ c	$(1+g)^2$ (r-g) p

→ Equation of: c = (r-g) p

↔ p = c / (r-g)

Example: p = 1$ / (0.07-0.03) = (1 / 0.04)$ = 25$

Annotation 2

Net Domestic Product: $Y = NDP = GDP - \delta K$

Solow Growth Model: $Y_t = A_t K_t^{\alpha} L^{1-\alpha}$

Growth of capital stock: $K_{t+1} = K_t + sY_t$

Assumptions (in part congruent with those of the original Solow model of 1956)[15]:

- Continuous time
- All factors of production fully employed
- No government or international trade
- Constant labor force
- Endogenous technological change
- Technological change nonrival, yet (partially) excludable
- Equilibrium: monopolistic competition
- Interest rate constant in steady state

[15] cf. Paul M. Romer, *Endogenous Technological Change*; in: Journal of Political Economy, Vol.98, No.5, Pt.2; Chicago: Chicago University Press, 1990.

Annotation 3

Marginal Product of Capital: $MPC = \partial Y_t / \partial K_t = \alpha A_t K_t^{\alpha-1} L^{1-\alpha}$

Profit Maximization of Firms: $\max_K \Pi = A_t K_t^{\alpha} L^{1-\alpha} - wL - rK_t$

$\partial \Pi / \partial K_t = \alpha A_t K_t^{\alpha-1} L^{1-\alpha} - r$

FOC: $r = \alpha A_t K_t^{\alpha-1} L^{1-\alpha} = MPC$

Capital Share in Total Income: $k = rK_t / Y_t = \alpha A_t K_t^{\alpha-1} L^{1-\alpha} K_t / A_t K_t^{\alpha} L^{1-\alpha} = \alpha$

$\leftrightarrow \alpha = rK_t / Y_t \qquad \leftrightarrow \qquad r = \alpha Y_t / K_t$

Steady state (r constant): $\alpha/r = K_t/Y_t = K_{t+1}/Y_{t+1} = (K_t + sY_t) / Y_{t+1}$

$= (K_t + sY_t) / Y_t \cdot Y_t/Y_{t+1}$

$= (K_t/Y_t + s) \cdot Y_t/Y_{t+1} = (\alpha/r + s) / (1 + g)$

$\leftrightarrow (1 + g) \cdot \alpha/r = \alpha/r + s$

$\leftrightarrow g \cdot \alpha/r = s \qquad \leftrightarrow \qquad r = g \cdot \alpha/s$

Annotation 4

From production side: $r = g \cdot \alpha/s$

Model of Preferences: $U(C_t, C_{t+1}) = C_t^{1-\sigma} + \beta C_{t+1}^{1-\sigma}$

Ramsey consumers' preferences; imply a parallel relation between rate of growth of consumption and marginal rate of intertemporal substitution.[16]

Assumption: $1 > \sigma > 0$

→ If g rises then s rises, so r rises less than g

→ higher growth rates lead to higher price earnings ratios

Empiricism: No evidence for higher savings rates in periods of accelerating growth.

Empirical analysis implies that $\sigma > 1$.

Result that higher price earnings ratios follow from higher growth rates does not hold.

[16] cf. ibid, pp. S87f, S93, S97.