Measuring the Link between Asset Returns and Economic Growth

by

Zachary Grossman[*]

David Love, Advisor

A thesis submitted in partial fulfillment
of the requirements for the
Degree of Bachelor of Arts with Honors
in Economics

WILLIAMS COLLEGE

Williamstown, Massachussetts

7 May 2007

[*] This paper could not have been completed without the help and support of many people. Special thanks goes to my advisor Professor David Love, as well as Professors Peter Pedroni and William Gentry. I'd also like to thank Stephen Abbott, Doug Hammond, Rohan Mehra, and Elizabeth Bond for editing help, as well as all of my friends and family that had to listen to me give long-winded answers about the status of my project.

Abstract

The United States is projected to enter a period of slowing economic growth in the near future due to an aging workforce. These retiring workers will also increase the demand for resources from government insurance programs. Many reform proposals intended to restore the solvency of these programs rely on the exploitation of high asset rates of return. In this paper we examine how asset returns have moved in relation to economic growth. After developing theoretical predictions based on the major macroeconomic models, we specify two models to test empirically for the expected relationships. We find that there is a strong relationship between asset returns and output growth. These results contradict the assumptions used to calculate the impact of Social Security reforms and have major implications for the future of government insurance programs.

Table of Contents

I. Introduction

The United States has seen robust economic growth since the end of World War Two partially fueled by the labor of the baby-boom population. A predicted decrease in productivity along with the imminent retirement of the aging workers will contribute to the slowdown of economic growth of the U.S.[1] Additionally, the retiring baby-boomers will put strains on the major social insurance programs in the U.S. as they draw on Medicare and Social Security, as well as private (but government guaranteed) pension plans. Medicare is projected to cost at least 10% of GDP per year at the end of the next half-century. Social Security will also have a significant influence on the budget in the next seventy-five years, with an expected future cost of 6% of GDP per year (Gale et al., 2004). For reference, the government only collects approximately 20% of GDP from tax revenues. The math is not hard—in order to avoid serious shortfalls, it is likely that the government will need to undertake serious reforms.

The strain we expect these insurance programs to encounter depends critically on our assumptions about growth and the return to capital. If, contrary to expectations, the economy continues growing rapidly, a program like Social Security will remain reasonably affordable because the rising dependency ratio (number of beneficiaries to tax payers) will be at least partially offset by increased per-worker tax revenues. High returns to capital will also help keep social insurance programs solvent. To see this, briefly consider how much needs to be invested today to yield a $200,000 retirement nest egg in 25 years—this is called present value discounting. If the portfolio return is 7%, we

[1] Baker et al (2005) cite Social Security Administration estimates of 0.3% per annum growth in hours worked after 2015 compared to 1.6% growth from 1958-2004, and productivity growth of 1.6% per year after 2001 compared to 2.2% between 1990-2004.

require an initial investment of approximately $36,850. If the average portfolio return is 6%, though, the necessary initial investment rises to $46,600. The rate of return on investments, then, is critical in the evaluation the present value of future obligations—in our context, the present discounted value of the expected shortfall shrinks as the rate of return increases.

This paper examines how the returns on assets and output growth have covaried over the past sixty years. Specifically, we would like to better understand the implications of an impending slowdown in U.S. output growth on asset returns. After examining the previous empirical and financial research on the topic, we specify two empirical models in order to identify and quantify this relationship. The first model is a growth rate based time-trend regression, and the second is a log-level bivariate Fully Modified Ordinary Least Squares (FMOLS).

For each model we test the relationship over multiple time segments and make use of different stock index values, with and without dividends, as our measure of assets. We do find a relationship between asset returns and output growth that is of the hypothesized magnitude. While the relationship has not remained constant over the past sixty years, it is unclear whether asset returns have recently become more sensitive to changes in economic growth.

In Section II we examine the policy and decision making implications of this slowdown in the background section, followed by a treatment of the major macroeconomic models and their theoretical predictions in Section III. Section III also reviews the financial and empirical literature that is relevant to our examination and then

Section IV explains our empirical models. Section V reports and discusses our results, Section VI concludes, and Section VII contains tables and charts of the results.

II. Background

Social Security

Currently, Social Security collects enough money through payroll taxes to not only cover benefits but also run a surplus. As the baby boom generation retires, though, the benefits drawn by the new retirees will quickly surpass the revenues collected through taxes—we mentioned earlier that Social Security is estimated to cost up to 6% of GDP, compared to approximately 4% now, in the upcoming decades (40% of all government tax revenues). In the short term the government will draw on the Social Security surplus— a trust fund that is invested in government bonds, on average returning a rate below the return on equities. For a long-term solution, though, we must examine both economic growth and asset returns to understand the tradeoffs inherent in prospective reforms.

Most proposals for Social Security reform can be divided into three categories: pre (or fully)-funded, pay-as-you-go, or a hybrid of the two. The former automatically taxes workers and places these funds into private portfolios that individual workers can access once they retire. The existing system is mainly pay-as-you-go—the current generation of workers are taxed to pay for the benefits of retired workers. It also incorporates a redistributive component, as the benefits of the highest earners are not nearly as large as the present value of the taxes they pay in. It should be noted that the current system has a large implicit debt built into it—current workers are promised benefits, but the government has deferred raising the money necessary to pay them until the next generation enters the workforce. This means that if reforms occur that make Social Security a fully-funded program, the government will no longer have a source of taxes for the current retired generation and will need to borrow to pay their benefits.

Let us assume, for simplicity, that returns and growth can each either be high or low. This means there are four possible combinations of asset returns and economic growth that have major policy implications. First suppose that both rates are high. As mentioned earlier, the high economic growth would bring about an increase in the revenue from payroll taxes, partially offsetting the rising cost of the program. If the system has been switched to fully-funded privatized accounts where savings exploited market returns, we would see that the high market returns would decrease the necessary initial investment. The system that performs better in this state of the world depends on the spread between wage growth and asset returns, and so it is not immediately clear whether reform is desirable. Next, consider high economic growth combined with decreased asset returns. Once again we find that the current system avoids disaster despite the rising dependency ratio, although a fully-funded system with market investments suffers—personal portfolios require large investments in order to satisfy retirement obligations.

When economic growth remains high, a pay-as-you-go system performs at least as well as a fully-funded one. Next consider what happens if economic growth slows, as it is expected to do. If asset returns remain high we see the opposite result of what we saw before—the current system suffers as payroll tax revenue will not rise enough to cover the increasing obligations. A privatized system manages to exploit the high returns on assets, so meeting future obligations is much easier. Finally, if both returns and growth decrease the spread between asset returns and wage growth is once again important. Because both rates are falling, though, the size of future obligations in present value terms are increasing, and so we find that neither system performs well—tax revenues do not rise, while the required initial investments are large.

In the context of the current debate, these possibilities mean that there is only one outcome in which moving to a pre-funded system and investing in assets definitely outperforms a refined version of the current system: high asset returns coupled with slowing economic growth. In this scenario, we might even be able to use some of the high returns to pay for the increased government debt, incurred by the change to privatized accounts, and still maintain the benefits promised to the retiring generation. This would be ideal, as one major issue with switching to a privatized account system is that the government would have to shift the implicit debt of social security into a large explicit one.[2] One critique of this course of action is that it is questionable whether the U.S. should be relying on more debt with its already large obligations.

Economists are expecting a slowdown in economic growth in the near future.[3] If this comes to pass, we narrow down our possible states of the world: either asset returns remain high and privatizing Social Security is the best course of action or neither system performs well. In either case, a fully-funded system is at least as solvent as the current one. So if there are possible benefits and no anticipated extra costs of switching to private retirement accounts, why has Social Security not already been reformed? We have simplified the debate somewhat, but there are many unresolved issues slowing

---

[2] See CBO (2004), "The Outlook for Social Security."
[3] See Diamond (2000) and Baker et al. (2005).

political adoption: administration costs and redistributive tendencies are but two examples.[4]

Even more importantly, we must consider the likelihood of future asset returns that are significantly higher than the risk-free rate. If this state of the world is unlikely to occur there is little advantage in switching to a fully-funded system. Over the past sixty years, assets have traditionally returned an average of 7.0% per year, much larger than the 3.0% average returned by government bonds (Diamond, 2000). This sustained difference in returns cannot be explained by economic theory. If two investments have equal risk, we expect investment to pour into the opportunity with the higher return until each vehicle has the same yield. Since the difference between bond and asset returns has been stable, the higher return on assets is considered the premium for holding risky assets. This historical discrepancy between asset and bond returns is called the equity premium; according to Mehra and Prescott (1985), such a large risk-premium implies implausibly high risk-aversion, and therefore the phenomenon is known as the equity premium puzzle.

Despite having a large equity premium in the past, it is not so clear that it should be assumed into the future. Diamond (2000) points out that The Office of the Chief Actuary (OCACT), the department in charge of projecting Social Security fiscal policy, has been mandated to use 7.0% asset returns. This implicitly assumes the continuation of a large equity premium due to high asset returns. With this assumption the resulting projections of switching to a more fully-funded Social Security system may severely overstate the increased solvency any privatizing reform provides.

---

[4] For a more in depth discussion on Social Security reform, see Samwick (2007) and Feldstein (1996).

Pension Plans

Similar to Social Security, the solvency of pension plans hinges on the future of asset returns. Pension plans take two major forms: defined benefit and defined contribution. Defined benefit plans stipulate a salary for retired workers once they leave the workforce, while defined contribution plans pre-fund portfolios for workers.[5] Funding for defined benefit plans relies on proper planning by firms; they must accurately calculate their pension rate of return in order to meet their future obligations. Unfortunately, many firms fail to correctly judge their pension requirements, and so the government has helped remedy this problem by covering shortfalls through the Pension Benefit Guaranty Corporation (PBGC). There is an interesting question as to whether these shortfalls occur because of poor planning or moral hazard caused by the government, but we will not deal with that here.

Among economists, there is great debate over whether the correct method of funding is to use stocks or risk-free bonds to fulfill future benefit needs. Zvi Bodie (2006) believes that pensions should use bonds to cover their future obligations. The main argument goes as follows. Assume a firm invests in the stock market an amount that they believe will fully fund their pension. The major risk they face is that the market does not perform as well as expected, and so they do not have enough money to pay their obligations at the end date. To insure against this, they buy put-options. They also note that it does not matter if the stock market performs better than expected—they only need the fund to have enough in it to pay for their obligations, so they sell off the upside potential by writing a call-option. By put-call parity, however, this combination is

[5] DC plans are similar to a fully-funded Social Security, while DB plans are like the current Social Security system.

equivalent to simply buying a bond, so firms should simply purchase bonds.[6] This system works adequately for the PBGC, the government program that assumes responsibility for bad pensions.

Deborah Lucas and Stephen Zeldes (2006) argue that an optimal pension plan must make use of the stock market as well as bonds to meet future obligations. The argument rests on the fact that firms cannot simply plan out their future obligations—new and current workers require different pension treatment than retired and separated workers. Because wages are related to economic growth and stock market returns, a portfolio entirely composed of bonds will be insensitive to the dynamic components of the obligation—firms need a portfolio that can cope with obligations that vary based on the prevailing economic environment. Lucas and Zeldes go on to refute the argument that firms should mimic the PBGC, as their commitments are only for retired and separated workers—most firms have enough current workers that their obligations are significantly different than that of the PBGC.

If the return on assets relative to bonds changes in the near future, then optimal provisioning for large pensions will change. Slowing output growth in general should decrease the rate of return on bonds, so to meet future obligations more money will be required initially. Now consider what happens if asset returns also decrease with the output slowdown. For a pension using Lucas and Zeldes' model, the optimal balance of risky to riskless assets will most likely shift. Ideally firms will be able to make the necessary adjustments, although with the PBGC guaranteeing their pensions it is very possible that they make no alterations and the government ends up assuming the debt of

[6] Put-call parity states that for a single strike price:

stock = call - put ➔ stock + put – call = 0.

more poorly planned pensions, adding extra strain to already stressed insurance programs.

Decision Making

In addition to having major policy implications, linked economic growth and asset returns affect households decisions as well. Consider a representative economic household attempting to maximize lifetime utility. According to the standard lifecycle model, these households will attempt to smooth the marginal utility of consumption over time, which typically means saving when young in order to maintain consumption levels in retirement. Whether or not people actually achieve this level of saving is debatable—John Scholz, Ananth Seshadri, and Surachai Khitatrakun (2006) made use of a dynamic life cycle model and found that most households may in fact be saving too much for retirement. This opposes the more common negative outlook on saving, which points to a negative savings rate as the neglect we are paying to our future.[7] Whatever is actually optimal, it is worth noting that there is no complete consensus on the matter.

Economic models of consumer optimization, which we will revisit later, make use of two variables that are very closely correlated with asset returns and output growth. A representative household has two sources of capital inflows: wages and returns on assets. It has been empirically documented, and therefore used when solving these models, that wage growth is approximately the same as economic growth. Economists also sometimes make the simplifying assumption that wages and asset returns are uncorrelated. Here we may have a problem, as it is our contention that the two are correlated over the long run, and so changes in one will occur with changes in the other.

[7] See Darlin (2007).

Consider what this means for a representative household. If economic growth slows in the future, therefore decreasing their expected wage growth, they will increase their saving in order to smooth consumption. However, as hypothesized and tested in this paper, we also expect that asset returns will decrease, and so by increasing savings the household will lose current and future consumption both because of higher saving and lower returns on those investments. The net result is that consumption now and in the future suffers because of slowing economic growth in the optimal smoothed path. Now consider what happens if a household takes the Office of the Chief Actuary's projections of asset returns and economic growth as given, or in the case of the model leaves the two growth rates uncorrelated. The household will only react to the slowing wage growth, increasing saving and decreasing consumption. They will not make the secondary adjustment of increasing consumption that the imminent lower rate of return dictates because they incorrectly assume it will not happen. This will result in suboptimal amounts of current and future consumption. If people are already allocating their funds incorrectly, and Social Security benefits decrease to deal with the impending shortfall, an entire generation could be left without sufficient saving for retirement.

III. Previous Research

We will first examine theoretical and empirical work related to the return of assets and economic growth.

Theory

Many economists have made note of the problem with OCACT's assumption of continued high asset returns despite slowing economic growth. Baker, DeLong, and

Krugman (2005) use the major macroeconomic models to argue against the likelihood of continued high asset returns. Specifically, they assume that if a slowdown in output growth is imminent, there will in turn be a reduction in the rate of return on assets. Working through the Solow growth model with a Cobb-Douglas production function, they derive that the steady state change in the real rate of return, r, will be:

$$\Delta r = \tfrac{\alpha}{s}(\Delta n + \Delta g),$$

where n is the population growth rate, g is labor augmenting technological growth, α is the fraction of income paid to capital, and s is the saving rate. The slowdown in economic growth is attributable to the aging of the population (decreasing n), and forecasted decrease in productivity growth (decreasing g). With these projections, the Solow model predicts a decrease in the rate of return on assets because of the decline in the marginal productivity of capital. The fraction (α/s) is historically around three (although recently it has been higher), so we would expect a greater than one for one drop in asset returns as the baby boomers retire.[8] Baker et al. note that a decrease in saving rates could result in a constant return to assets, but this requires endogenizing savings decisions, and therefore a different model.

The authors extend their analysis to the Ramsey and Diamond OLG models in order to examine the possible effects of the economic slowdown on savings decisions. Working with the Ramsey model, they assume households have utility:

$$\sum_{t=0}^{\infty}(1+\beta)^{-t}(U(C_t))N_t^{1-\lambda},$$

[8] We assume an approximate historical savings rate of $s \approx 0.1$ and a share of income paid to capital of $\alpha \approx 0.3$.

where β represents time preference, C_t is household member consumption, and N_t is the number of representative household members, which grows at a rate of n. The standard setup assumes $\lambda = 0$ (Romer, 2000), which leads to the result that there is no relationship between steady state asset returns and labor force changes.

The assumption of $\lambda = 0$ means that each person cares about their own utility and the utility of future people equally. This seems slightly unrealistic, and so Baker et al. (2005) propose the idea of "imperfect familial altruism," where each person cares more about their own utility than the utility of future people. Mechanically this means that $0 < \lambda < 1$. Assuming log-utility and consumption growth equal to labor augmenting technology growth, g, they solve the model to reveal that in a steady state:

$$r = (1+\beta)(1+n)^{\lambda}(1+g) - 1,$$

which for the distant future becomes:

$$r = \beta + g + \lambda n.$$

The Ramsey model predicts that the return on assets moves one for one with productivity changes, and less than one for one with population changes.

The assumption that consumption growth equals labor augmenting technology growth may not be realistic, especially as a population is aging—as people get older they consume a higher proportion of wealth and income than when young, and so in an aging population consumption growth will outpace technology growth. Baker et al. (2005) move to the Diamond OLG model in order to examine different consumption and savings decisions among agents. Assuming a Cobb-Douglas production function, they show that the real rate of return r_t, income on capital over total capital, is:

$$r_t = \frac{\alpha y_t}{k_t} = \frac{\alpha E_t^{1-\alpha} k_t^{1-\alpha}}{(1+n)^{\alpha-1}},$$

where α is the share of income paid to per capita capital, k_t. E_t is labor efficiency which grows at g, and n is the population growth rate. Assuming log-utility, they solve the model to show that:

$$r = \left(\frac{\alpha(1+g)(1+n)(2+\beta)}{(1-\alpha)} \right),$$

where β is once again a time preference. Taking partial derivatives shows us how r reacts individually to changing rates of population and efficiency growth.

$$\frac{\partial r}{\partial g} = \left(\frac{\alpha(1+n)(2+\beta)}{(1-\alpha)} \right),$$

where since we are assuming Cobb-Douglas, using $\alpha = 1/3$, this reduces to:

$$\frac{\partial r}{\partial g} = (1+n)(1+\tfrac{\beta}{2}) \approx (1+n+\tfrac{\beta}{2}).$$

The relationship between returns and population growth works out similarly:

$$\frac{\partial r}{\partial n} = (1+g)(1+\tfrac{\beta}{2}) \approx (1+g+\tfrac{\beta}{2}).$$

The Diamond OLG model therefore predicts, in the steady state, a slightly larger than one for one change in the rate of return in response to a change in either population or efficiency growth.[9]

It will also be helpful to examine how economists model stock pricing. A stock is a claim of ownership on a productive piece of equity. Companies can distribute earnings to shareholders through dividend payments, or by repurchasing shares from investors. Alternatively, firms can retain earnings and reinvest in the company. The return on a

[9] We can use the Solow model to gain some intuition as to why population growth rates act as a multiplier for the effect of productivity growth changes on asset returns, and vice versa. In the steady state, an economy with high population growth will have a high marginal product of capital, and so a positive shift in productivity growth will result in a larger shift in returns (graphically, the depreciation line is steeper).

stock is therefore dependent on the price changes of the stock and all of the dividends paid out to owners. The Gordon formula calculates the price, P, of a stock as:

$$P = \frac{D}{r_e - g},$$

where D represents dividends and buybacks (net cash returned to investors), r_e is the real return on equity, and g is the growth rate of dividends and buybacks. This can be rearranged to show us that the average expected return on equities is:

$$r_e = \frac{D}{P} + g.$$

Baker et al. (2005) define g as the difference between the expected real rate of output growth and the discrepancy between the CPI and GDP deflator. Holding all else constant, as output growth slows, then, the Gordon pricing formula also predicts a decline in the expected return on assets.

Financial and Empirical

Baker et al. (2005) have shown that the major macroeconomic models predict that slowing growth in the labor force coupled with a possible decrease in productivity growth will result in a reduction in asset returns. Depending on which model we believe best describes the relationship between output and assets, the output growth slowdown will occur along with a decrease in asset returns that is anything from equal to or greater than the slowdown in economic growth. While the models give a good simplified picture of how the economy works, the next step is to turn to the data and see whether their predictions hold empirically. There is plentiful data on aggregate asset returns and economic growth over the past century, so we can easily pursue our goal. The following is a review of the empirical and financial literature that will establish the context for our empirical tests.

We ideally would like to examine how the predicted slowdown in future economic growth will affect asset returns today. One major issue that we should acknowledge is that business cycles and countercyclical monetary policy will affect the relationship between asset returns and output growth. Ideally long-term projections of output and asset returns are unbiased by the short run influences of the business cycle. This long-term approach has been used to examine the relationship between deficits and interest rates.[10] While we do have Congressional Budget Office (CBO) data on projected output, we were unable to find a reliable measure of forward asset returns. Our models will therefore make use of current data that is not immune to distortions from business cycles or monetary policy. Even though we have no way to fix this issue, it is worth noting that these aspects of the economy may influence our results.

Reacting to OCACT's decision to assume high asset returns despite an aging population and slowing productivity growth, Diamond (2000) presents a detailed argument against the assumption of 7.0% long run asset returns coupled with a 4.0% equity premium. Changes in capital markets, like decreasing charges on mutual funds and increased ownership of stock by individual investors, help to lower the equity premium. Diamond also points to the high valuation of stocks we have seen recently. According to many established metrics, like earnings and book value, stocks are more highly priced now than they have ever been, and so a consistently high rate of return, without a major correction, is unlikely.

---

[10] For a more detailed treatment of this model, see Laubach (2005).

Diamond uses the Gordon formula to estimate the magnitude of the price correction necessary for long-run high asset returns. Taking future returns and output growth as given leaves dividends as the only free variable. Diamond examines multiple reasonable adjusted dividend (buybacks and dividends) payout ratios, and finds that, for returns to remain high, stock prices would need to decrease anywhere from 21% to 55%, a significant change. Finally, he briefly discusses the theoretical implications of the Solow model, which we have already examined in more depth. Diamond concludes that the likelihood of a sustained large equity premium, along with slowing economic growth, is extremely unlikely.

One of the trickier aspects of dealing with rates of return is that the relationship between asset returns and the risk-free rate is not fully understood. Mehra (2003), the original co-discoverer of the equity premium puzzle, surveys the attempts to explain the discrepancy between asset returns and bond returns. Because the puzzle rests on the inability of major economic models to explain observed trends, many economists have altered the preference structure of the models to better conform to reality. Mehra argues that generalized expected utility models, which unlink the elasticity of intertemporal substitution from risk aversion, as well as habit formation models, which incorporate previous levels of consumption into consumption choices, only manage to explain part of the puzzle.[11] Using the Gordon pricing formula, as in Diamond (2000), Baker et al. (2005) perform quick calculations to show that a stable and large equity premium is not anticipated if the economy slows. While this may imply that asset returns will be dropping, we must remember that the equity premium is still a puzzle—models do not

[11] Mehra (2006) elaborates in detail on these explanations.

adequately explain it, and so the fact that a model forecasts a small premium is a hard claim to fully believe.

While most of the figures and examples used so far have been taken from U.S. data, it is worth noting that the U.S. is not unique: a large equity premium and high stock volatilities are present in many other developed countries as well. John Campbell (1999) examines stylized facts about the United States and other nations in an attempt to gain insight into the equity premium puzzle and unusually high volatilities across the world. Campbell also tests the ability of the stock market to forecast various macroeconomic variables; the independent variable he uses is a log price-dividend ratio. The only regression that yields significant coefficients is the model that uses excess returns of the stock market over that of treasury bills—this holds true both internationally and in the U.S. While we will not be dealing with treasury bills, the ease with which his exploration generalized to international data seems to validate extending our own analysis to other countries besides the U.S.

Focusing on the United States, it would be useful to know what actually causes output or asset fluctuations. While economic growth and stock market returns are most likely co-determined, the most common relationship studied between the two variables has been the predictive power of stocks on output and investment. To see why this makes some sense, first consider the task of predicting stock market returns. If we believe the efficient markets hypothesis, stock markets incorporate all publicly available data—actually predicting their fluctuation should be out of the question. The actual efficiency of markets has greatly increased in the past twenty years as the speed of information has increased due to technological advances. We will return to the possibility of predicting asset returns with output growth later.

Asset returns, though, may foreshadow GDP changes. Over seventy years ago Burns and Mitchell (1938) noted the predictive power of the stock market for GDP growth. This makes some sense—stock markets, as a forward looking aggregation of information, are continually updating the value of the corporations that make up the economy, and therefore should be related to output. Given that the stock market continuously reports value while GDP estimates, although updated continuously, are released every three months, stocks can incorporate information about GDP levels prior to their official release.

Using an institutional framework, Ross Levine and Sara Zervos (1996) examined the relationship between banks, stock market liquidity, and GDP growth. They point out that much literature has detailed how the banking system can spur economic growth by funding productive investments. They also highlight the sparse, although growing, research on the relationship between well-functioning stock markets and output growth. Pointing to theoretical models developed recently they explain that a relationship should exist, as more liquid markets reduce disincentives to longer-term projects that help maintain long run output growth. Indeed, their empirical tests find significant and strong positive correlations for stock market liquidity and banking infrastructure on long-term output growth, productivity growth, and capital accumulation.

Many other economists, like Fama (1981), Harvey (1989), and those discussed below, have examined the claim that asset returns can predict output growth. Fischer and Merton (1984) explore the relationship between output, investment, stock prices, and alternative measures of asset returns. They find that the stock market is a good predictor of Gross National Product (GNP) because stocks are closely tracked by consumption and investment.

Robert Barro (1990) examines this link in more detail. Using vector autoregression with quarterly and monthly stock index values, Barro shows that up to fifteen-month lagged stock index values significantly explain variation in current investment and GNP growth. He finds that stock prices adjusted for inflation explain variation slightly better than nominal values do, so we will also be using real prices. Projections based on this data predict future output and investment fairly well except in times of stock market crashes (a time when the valuation of the stock market is severely corrected). Furthermore, these predictions appear to do better than projections using Tobin's q, the ratio of the value of capital divided by the cost of acquiring new capital, which had previously been the standard.[12]

Not everyone is convinced of the actual predictive ability of stock prices. Stock and Watson (2003) argue that "this link is murky" (p. 797). A combination of poor in-sample projections coupled with decreasing significance of stock prices when including lagged output values severely limits the robustness of previous results. Indeed, their results show that for certain periods of time, asset values can be good predictors for output growth, but at other times there is no significant predictive ability. More positively, they do show that combinations of measures of assets better forecast output growth than simple autoregressive models. Unfortunately, asset returns, and more

[12] One interesting note about Barro's study is that Canadian investment and GNP are better explained by the U.S. stock market than by Canadian indices. This fact may be important when extending this paper's analysis to other countries.

specifically stock market returns, seem to have significant predictive power only in particular circumstances.

The complexity of stock markets makes them incredibly hard to explain and predict. The Gordon formula is one example of how we can think about stock prices. Another well-known model is the Capital Assets Pricing Model (CAPM), where some combination of the market portfolio and a risk free asset create the "efficient frontier" that contains the optimal portfolio. While it is a relatively simple method of dealing with markets, the CAPM has major shortcomings in that it does not adequately explain average U.S. stock returns. In a famous effort to fix this Fama and French (1992) added extra metrics, book-to-market and size factors, that help to better explain U.S. stock returns.

Exploring the advantages of the Fama-French model over CAPM, Maria Vassalou (2003) adds in a variable for GDP growth-related news as a possible explanatory factor in stock market returns. She is careful to point out that in reality this news is unobservable and is distinct from expectations of GDP growth. She models the news by creating a mimicking portfolio that can predict output growth.[13] The discovery is that when she adds this portfolio into the CAPM, the faulty model correctly deals with average stock returns. It even turns out that when the Fama-French fixes are added into Vassalou's CAPM, there is no incremental increase in the explanatory power of the model (the Fama-French fixes only explain a small amount of variance depending on the mimicking portfolio). Despite the difficulty of actually verifying this model, the conclusion that

[13] See Vassalou (2003) for a more detailed explanation of these portfolios.

news concerning GDP growth helps explain stock market returns validates the methods we will use on our search for the relationship between output growth and asset returns.

IV. Estimation Methods and Data

Methods

While most economists have focused on using the stock market as a predictor of output, our concern is slightly different. For policy as well as theoretical reasons we are more interested in how slowing future output growth, and the expectation of it, may be related to asset returns. We should first briefly examine why it makes sense that output might have an effect on assets. High GDP will signal to companies and their owners that income is high, so firms can expect a good business year. In this manner strong GDP growth will create high stock prices and good returns. As argued before, it is also possible that higher stock values correspond to companies doing well, and so this will be reflected in GDP. We can see this in the Gordon formula, as the return is directly proportional to output growth.

We will use two models to examine the relationship between stock market returns and output growth: a time trend model that uses growth rates and a Fully Modified OLS model that makes use of log-levels. The simplest way to examine the co-depending nature of output and the stock market is to look at trends in the data. Although we may not be able to explain in detail how stock markets fluctuate and how their movement relates to real GDP changes because of endogeneity concerns, it is not hard to check if both variables move in the same direction most of the time. Keeping in mind that we are trying to examine the long-run consequences of slowing output growth, it will be helpful to know how the stock market and GDP moved in the past. A preliminary approach will check if there is any significant trend in the growth rates of either output or equities. This trend may not necessarily exist, but it is worthwhile to check if either variable has grown in a particular way over the past sixty years. To examine a more direct correlation between asset returns and output growth we will specify the model

$$r_t = \alpha + \beta_0 t + \beta_1 y_t + \varepsilon_t,$$

where r_t is the quarterly growth rate of real stock index values, t is time, and y_t is the quarterly growth rate of output growth. While it appears that β_1 is the coefficient of interest in this model, it actually is not. This coefficient measures the short run relationship between asset returns and output growth, not the long run steady state relationship that we are interested in. Additionally, the standard errors from a simple OLS regression will not be robust to problems arising from omitted variables.

The coefficient β_0 is the estimate with which we are primarily concerned. If, in the steady state, there is a trend in the growth rates of either assets or output our initial trend regressions will register it. When regressing output growth rates on asset returns, we will be examining β_0 to implicitly check if there has been a change in the growth rate trend relationship. If, as our examination of the Solow model highlighted, similar changes in economic growth have a larger affect on asset returns more recently, then β_0 should be non-zero. This coefficient is robust to omitted variables.

A few notes about the estimation of the model follow. Both stock market and output time series are non-stationary—the mean and variance of the series change depending on when we examine them. Since these series are changing over time, we may infer correlation when we are in fact noticing each series moving by itself. The nonstationarity that is present in these series is called a unit root. Because we are dealing directly with growth rates, we are avoiding the unit root that is embedded in asset returns and output growth. We will estimate our coefficients using Newey-West robust standard errors to correct for any serial correlations that are in the data. Barro (1990) found that coefficients lost significance after a fifteen-month lag in stock prices, and so, given that our data is quarterly, we will use up to five lags for Newey-West standard error estimation.

Another important issue is the timing of our series. GDP growth is a discretely released variable that is announced every quarter while stock index prices are flow variables that we have averaged into quarterly data-points. We have established the general notion that asset returns and output growth should be related, although the actual speed and channels of adjustment are unclear. As economic growth gradually converges towards its expected value over the upcoming decades, we assume the accompanying change in asset returns follows a pattern similar to that of the past sixty years. We will be testing the model with the relevant quarter's GDP announcement corresponding to the quarterly average of the three months that make up that quarter, so the stock index values are slightly lagged.[14]

This model is also very sensitive to the choice of time period; we will examine both the entire period and smaller segments of time with this model. Regressions using the entire period will give us some idea of the historical relationship between asset and output fluctuations. Dividing our time period into shorter segments will allow us to more closely examine the trends in the relationship—the Solow model predicts that changes in

[14] The stock index averages have components that are measured prior to the release of the GDP data. It is not a lag in the traditional sense.

output growth could affect asset returns more now than they have in the past.[15] Based on the analysis from Baker et al. (2005) we expect to see anything from a one for one to three (or higher) for one relationship, although there is no coefficient in this regression that will quantify the actual relationship.

Finally, it should be noted that this regression is susceptible to false negatives. It is unlikely that there is any noticeable trend in the growth rates of either assets or output. In fact, if we are assuming efficient markets, a direct time trend should not be present in asset returns. We should be a little concerned, then, as the trend coefficient is what we are expecting to show the relationship. There are three different reasons the trend coefficient may tell us nothing. If there is no actual trend then this coefficient will be zero. We should also find a zero point estimate for this coefficient if the relationship between output and assets has remained constant. Additionally, even if the response of asset returns to output growth has changed over the past few decades, this change may be entirely absorbed by the coefficient on output growth. This would be unfortunate, as we cannot use this output coefficient for inference. Since we do not really have a way to control for which coefficient picks up this change, we will specify the following alternative model as well.

Our next model allows us to quantify the relationship between asset returns and output growth. Instead of bypassing the unit root present in the levels of our output and stock market series, we can examine how our data has covaried over the past half-century with a log-levels bivariate model. We said earlier that using growth rates helped us avoid the unit roots that were part of time series involving output and stock market prices. This

[15] As mentioned on page 10, this is due to a decreasing savings rate.

levels regression will work by first checking for and acknowledging the unit root using an augmented Dickey-Fuller test (ADF) on each series. This test checks the null hypothesis of a unit root present against the alternative that the unit root does not exist. It should be noted that the distribution of an ADF test is not normal, and so we will report the critical values along with the t-statistic in our tables. Once we have established the unit root we also need to check if the series is cointegrated—that is, have the unit roots in the series moved together. We will use the Engle-Granger cointegration test by running the following bivariate regression and then using an ADF test on the residuals to check for cointegration:

$$r_t = \alpha + \beta_1 y_t + \varepsilon_t .$$

For this second ADF test the null hypothesis is that there is no cointegration, with the alternative being that cointegration is present. The EG test works by checking to see if the unit root has been eliminated by the linear regression.

The point estimate of β_1 will be robust to omitted variables and the unit root, although the standard errors will not. In order to fix this last problem we will use a Fully Modified OLS regression, which will make our standard errors consistent. We should note that the ADF test used for cointegration is only valid if we find a unit root in the first test, and both tests are required to use FMOLS. One advantage of the FMOLS model for us is that the procedure shows little distortion when using smaller samples, which is good for us because we examine relatively small segments in some regressions.[16] We are in effect measuring the long run steady state relationship between asset returns and output growth. By splitting our data at either an economically motivated or arbitrary date we

[16] The virtues and disadvantages of FMOLS and other panel cointegration methods are discussed at length in Harris and Sollis (2003).

can explore if there has been a one-time shift in the relationship between asset returns and output growth. Once again timing will be an issue, although the biggest decision we face is where to split the time series.

Data

We make use of quarterly U.S. data from 1947 through 2006. In addition to regressing our models over this entire period, we also split it up into arbitrary and economically significant time periods. We will examine each half of the time period (1947-1976, 1977-2006), three evenly divided twenty-year segments (1947-1966, 1967-1986, 1987-2006), as well as segments divided according to the Tax Reform Act of 1986 (TRA-86), a major tax code change that affected capital gains taxation. As previously discussed, the goal of segmenting our time period is to discover if the magnitude of the relationship between output growth and asset returns has changed. We use this economically motivated break point in order to see if a shift in how investors value equity and investments altered the relationship between output and assets. The timing of where we actually divide the periods deserves some consideration, as the macroeconomic adjustments to a tax code change will occur at varying times, and so the value of assets adjusts at different speeds. We choose to break the segments after the fourth quarter of 1986, as the reform was signed at the end of 1986.[17]

To test our models we make use of quarterly data from various sources. Seasonally adjusted real GDP figures as well as any inflation controls are taken from the National Income and Product Account (NIPA) tables, while data on asset values is the

[17] There are two effects we may expect with a tax code change. If taxes are decreased (as in 1986), investments become more valuable and their price will appreciate. If firms do not actually have higher earnings, though, then the rate of return their stock pays could decrease.

simple average of monthly S&P 500, Dow Jones Industrial Average (DJIA), and New York Stock Exchange (NYSE) index values. We also have data on monthly dividends for the S&P 500. Index values for the NYSE are only available back to 1965, while the S&P and DJIA data goes back to 1947. We have adjusted the nominal stock index values by the implied GDP deflator (calculated by dividing nominal values by real GDP). Projections of real GDP growth are taken from the Congressional Budget Office (CBO), and other measures of market-indexed assets are taken from Yahoo! Finance. The stock index and output values are transformed logarithmically for the bivariate model and then differenced to reveal the equivalent of a continuously compounded growth rate for our trends model. Our final measure of assets, which we only test with the bivariate model, uses the real S&P 500 index values with real dividends added on. Dividend values are taken from Shiller (2006). This sum is then transformed logarithmically to satisfy linearity. It should be noted that we only have data until the second quarter of 2004 instead of the third quarter of 2006 for this series.

There are two issues of which we should be aware. First, each stock index is calculated differently. The S&P 500 is a float-weighted index, meaning only publicly traded shares of a stock count towards its value in the index, while the DJIA is a price index that uses a divisor that is much less than one to value the average. Despite the differences in the ways that value changes are incorporated into the index prices, each index is calculated using a consistent method, and so the returns should be comparable—we do not foresee any reason these differences will bias any of our results.

We must also take note of the fact that three of the four metrics we use to proxy for asset returns are solely the capital gains of holding equity. While our fourth metric does include dividends, we never incorporate any measure of debt. We are concerned this may prevent us from having a complete picture of asset returns. Corporate debt can have a large effect on the return on assets, as debt financing reduces the amount of equity used for a given project. For two companies with the same earnings, the one with more leverage will have a higher rate of return on equity and on assets, although will also be considered riskier. There is evidence for a changing trend in corporate debt structure over the past sixty years, so any correlation between the corporate financing and economic growth may bias our regressions.[18]

V. Results and Discussion

Table 1 presents the general trend regressions for the past sixty years for different measures of asset returns and output growth, while Figures 1-4 are histograms of the growth rates for the S&P 500 index, the Dow Jones Industrial Average, the New York Stock Exchange, and real GDP, respectively. Figure 5 displays the growth rates of both real GDP and the stock indexes over time. Note that there does not appear to be a significant trend over time in the growth rates, and the growth rates appear more distributed with fat tails than normally. Despite the fact that these trends are insignificant, we retain them in our direct regressions because we hope that the long-run trend between our variables will display any fluctuation or permanent change in the relationship between asset returns and output growth. Another interesting point is that while real GDP, the S&P 500, and the DJIA have virtually no time-trend, the NYSE has a trend that is actually significant at the 10% level.

[18] Bernanke and Campbell (1988).

There are two possible explanations for this. There is only data on the NYSE starting in 1965, while the other three variables go back to 1947, so the measurement period may be responsible for the trend. Unfortunately, shortening the time segments for the other three variables does not make the time-trend any more significant, so this is not quite an adequate explanation. Another possibility is the way we have treated dividends. Dividend payments have slowly been decreasing over the past century.[19] For the NYSE, the large number of established companies that are seeing more stock price appreciation as they pay fewer dividends combined with smaller companies that are less likely to pay dividends could explain this result. We will examine this more with our levels bivariate model.

Tables 2, 3, and 4 show the results of the time-trend growth rate regressions. For all three asset metrics, the S&P 500, DJIA, and NYSE, the time-trend model yields mixed results. Table 2 documents the regressions that use the S&P 500 as the dependent variable. While the time-trend over the entire period is not significant, the coefficient on output deserves some attention. This coefficient does not actually tell us anything—it represents the short-run relationship between economic growth and asset returns and we cannot infer long-term significance from this coefficient. Additionally, it is susceptible to omitted variable bias. Despite the uninformative nature of this coefficient it is interesting to note that these point estimates are close to what we would expect to see if the predicted relationship exists—for the past sixty years we would expect a one-percentage point increase in GDP growth to be associated with almost a 1.7 percent increase in S&P 500 returns, *if* this coefficient were meaningful. Remember that the various theoretical

---

[19] See Fama and French (2000).

treatments of the expected relationship predict that the multiplier on changes in economic growth is somewhere around one or slightly higher. The bivariate regressions will give us a better idea of the actual relationship.

Table 2 also displays the regressions where the whole time segment has been arbitrarily split up into twenty and thirty year segments. The only segments that have significant time-trends are between 1947-1977 and 1967-1987—every other trend coefficient is insignificant. Since we do not have any way of interpreting what this marginally significant relationship means, we hope that the FMOLS model will shed some light on the issue. It is also interesting to note that the point estimates on GDP growth are of the magnitude that we would expect to see if the hypothesized relationship exists. Unfortunately, they are in reality uninformative.

Next we turn our attention to the results of the DJIA measure of asset returns. We find that the coefficients on output growth are for the most part meaningless, much as we actually expected to see in all of the trend regressions. For the entire time period the trend coefficient is once again insignificant. As before, though, the trend coefficients for the segments between 1947-1976 and 1967-1986 are significant, while the remaining segments have insignificant trend estimates. So some trend seems to have been present during this time period, although contrary to what we have expected, there seems to have been no change in the relationship between asset returns and output growth in the latest periods.

Table 4 displays the results of the regressions using the NYSE as our definition of assets. We do not have quite as complete of a picture for this stock index as we did for the other two indices as there is only data starting in 1965. The coefficients on GDP growth show the same pattern as the coefficients in the S&P 500 regressions. Part of this may be because there is a significant overlap among the stocks that make up the NYSE and S&P 500. The coefficients on output for the NYSE are also slightly smaller than their counterparts in the S&P regressions, although once again these estimates are not meaningful. Unlike our other two asset metrics, the NYSE does show a significant time-trend for the entire period, although in this case that is only the period between 1967-2006. Similar to the other metrics, the trend between 1967-1986 is significant while the estimate for most recent twenty-year segment is not. Because we only have forty years of data, though, it is hard to directly compare these results to our other two metrics.

It is difficult to conclude much from the trend regressions. We consistently see a significant trend between 1967-1986, which indicates some changing relationship between asset returns and output growth. We do not, however, find this significance in the point estimates for the entire period or the majority of segments. As mentioned briefly in the methods section, there are a few possible explanations for this result. If the steady state trend between asset returns and output growth does not exist then the trend model will have point estimates indistinguishable from zero. This would also be the case if the relationship were constant. Finally, it is possible that a changing relationship could be accounted for by the coefficient on output growth, which varies considerably across our regressions. However, the fact that we find two significant trends for each regression, combined with the theoretically meaningless yet interesting variation on the output point estimates seems to indicate some unrecognized relationship between output growth and asset returns.

Next we turn our attention to the log-levels bivariate model. Tables 5 and 6 contain the results of the augmented Dickey-Fuller tests for unit roots in the series and cointegration on the residuals, respectively. We do not reject the null of a unit root in any of the series for any time period. Results from the EG test for cointegration, however, are less decisive. For the Dow Jones and NYSE measure of asset returns we reject the null hypothesis of no cointegration, and come quite close to rejecting the null for the S&P 500 (with and without dividends), between the years 1977-2006. The same cannot be said of the other segments. We are close to legitimately rejecting the null for the other two segments for the NYSE, in addition to the first twenty years for the S&P 500 and DJIA. One noted problem with the ADF test is that it has relatively low power for small sample sizes, and so our inability to reject the null may be a result of this.[20] Additionally, cointegration between U.S. stock markets and output has been documented by Yin-Wong Cheung and Lilian K. Ng (1998),[21] so we proceed with our models.

Since we have satisfied the prerequisites of a unit root and examined the existence of cointegration, we are justified in using Fully Modified OLS to estimate the relationship between assets and output. We ran this regression for seven different time periods—the whole period, three equal twenty year splits, two thirty year splits, and from 1947-1986 along with the final twenty years. Remember that this regression is designed to pick up the long run steady state relationship between assets and output, and so by dividing it up into multiple periods we are directly asking if the relationship has ever been different. Table 7 displays these results.

[20] See Harris and Sollis (2003).
[21] According to Benzoni et al. (2005) it is not unusual for economists to assume cointegration if it is "economically plausible" despite "weak statistical evidence" of it.

First we will focus on the S&P 500. Over the entire period, we expect a 1% change in real GDP to be associated with a 1.2% change in the S&P 500—the coefficient is also significant with a p-value of .014. The next two columns in the table display the results of regressions between 1947-1976 and 1977-2006. For the first period, the point estimate is small and the p-value is large, and so we cannot reject the null that there is in fact no relationship between assets and output. For the most recent thirty years, though, we find that a 1% increase in output would be associated with a 2.45% appreciation in assets. This coefficient is highly significant.

The final two columns display a different division of the time period, corresponding roughly to the major tax code changes that went into effect in 1987. Among other things, these changes reduced the tax rate levied on capital gains by virtue of reducing the rate of the highest bracket. Like the thirty-year segments, the first forty years, from 1947-1986, reveal a relationship of only .65, with slight significance. The final twenty years are similar to the final thirty years in the previous segments, where a 1% change in real GDP between 1987-2006 is significantly related to a 2.38% change in assets.

Finally, we also divide the past sixty years into three twenty-year segments to see if we can pick up a more gradual change in the relationship between assets and output. For the years 1947-1966 we see that a 1% increase in output is significantly correlated with a 2.04% increase in the stock market. Most interestingly, between 1967-1986 the relationship between assets and output seems to have reversed itself. In this time period, for a 1% increase in output we would expect a 1.26% decrease in the S&P 500, although this point estimate is not significant. The final segment is reported in the previous paragraph as it is exactly same as the twenty year segment of the TRA-86 motivated regression. Rather than a gradual shift, the relationship seems to have disappeared and then reappeared.

The FMOLS regressions that use the Dow Jones Industrial Average as the definition of assets are similar to regressions that used the S&P 500. Over the entire time period, a 1% increase in output is associated with a 1.55% increase in assets, although this coefficient is less significant with a p-value above .06. When split into two even time periods, we also see the same results. From 1947-1976, the insignificant point estimate is 1.1, while the very significant coefficient for 1977-2006 is 2.62. The economically-motivated segmented regressions are also alike. We see the same pattern with the DJIA as we did with the S&P 500 when dividing the period into three segments.

We ran the FMOLS regressions using the New York Stock Exchange composite as our measure of assets. While the S&P 500 is made up of many stocks that are listed on the NYSE, there are many smaller stocks also listed on this composite, making it possibly a better measure of economy-wide assets. Because we only have data starting in 1965, we run our regressions from 1967 so that we can compare this asset metric to the others. Over the whole forty-year period, we find that a 1% increase in real GDP should be associated with a 1.38% increase in the real value of the NYSE. Once again, for the segment between 1967-1986 the estimate of this relationship is insignificant and negative, while the over the final twenty years we believe that a 1% increase in real output is significantly related to a 2.33% increase in assets.

Finally, we also used the S&P 500 with dividends included as a dependent variable in the bivariate model. This series only runs through the second quarter of 2004, so we ran our regressions with a truncated final period. Over the fifty-eight year period we find that a 1% change in output is significantly correlated with a 1.15% change in assets. This is very similar to the results obtained with the simple S&P 500 asset metric. In fact, there is almost no difference between the S&P 500 regressions with and without dividends added in.

For all four definitions of assets we see a significant relationship between assets and output over the entire period. When we break the period into smaller segments, we estimate that the strongest relationship occurs in the most recent segments, exactly as predicted. There is a question, though, of whether our results are influenced by the middle twenty years and the estimated negative relationship between assets and output. When we remove these years, we still notice that the largest relationship occurs in the most recent twenty years. The difference between the point estimates for 1947-1966 and 1986-2006 (2.04 vs. 2.38 for the S&P 500 and 1.84 vs. 2.64 for the DJIA, 2.0 vs. 2.7 for the S&P 500 with dividends) may be small enough that there are two ways to reconcile the findings of both models. If in fact there is no difference between the relationships in the early and late years, then the constant relationship would not show up at all in the trend model. It is also possible that even if this shift were significant it is small enough that the variation was completely accounted for in the real GDP estimate of the trend model, which therefore returned a false negative.

It is also interesting to note that the trend model did pick up an unexpected changing relationship, although when we take into account the bivariate model it seems most likely that from 1967-1986 the relationship between assets and output became dampened for some time before reemerging in the recent past. The series are graphed in

Figure 6—while real output grew consistently during this period, the value of the stock market declined. Despite the fact that variables in our data are real, high inflation may still be part of the explanation for this observation. Between 1970 and 1985 the United States, in addition to experiencing incredibly high inflation, saw: the end of the Bretton-Woods system of exchange rates, multiple oil shocks, a new monetary policy regime under Volcker, and a shifting fiscal policy. The fight against inflation that occurred brought the U.S. into multiple recessions. If the stock market suffered more relative to output, which Figure 6 indicates, we would see this negative correlation. As if this were not enough, the United States also experienced a slowdown in productivity at the same time as the mass entry of women into workplace. The dynamic and influential nature of all of these phenomena most likely disrupted the macroeconomic environment in a way that has not occurred since.

Our result of a lack of a consistent and significant relationship across our time period agrees with both Stock and Watson (2003) and Barro (1990). Stock and Watson found that the stock market was a good predictor of GNP only some of the time, while Barro (1990) found that market corrections inhibited the stock market's predictive ability. The recessions and economic turbulence that are responsible for the results of our middle segment agree with these other findings.

The results of the FMOLS regressions lead us to believe that there is a correlation between output growth and assets that is very close in magnitude to what we expected. Based on the point estimates we find when splitting the time period in half and into before and after TRA-86 segments, it appears that the multiplier on changes in output may have increased, as we might have predicted. This result, though, is questionable because of the strange results we encounter between 1967-1986. If we exclude the middle twenty-year segment, it is unclear whether we can conclude that the relationship is different in the first and last segments—the largest difference resulted when using the Dow Jones Industrial Average as the measure of assets: we estimate that the effect is 79 basis points larger in the most recent twenty years.

Indeed, the combination of the significant time trends found for the segment between 1967-1986 and the negative and insignificant point estimates found in the bivariate model are consistent with a changed relationship between assets and output. In comparison, the lack of a significant trend for the first and last periods along with similar FMOLS point estimates provides evidence that the relationship remained constant during those years.

VI. Conclusions

In this paper we explored the relationship between asset returns and output growth. We first examined the policy and decision making implications of a changing relationship between assets and output. This was followed by an inspection of the theoretical predictions of the major macroeconomic models, where the work of Baker et al. (2005) points to a positive, yet indeterminate in size, relationship between asset returns and GDP growth. Building on the framework of economists who have used stock prices to predict output, we specified a time-trend OLS model in addition to a log-levels bivariate model in order to examine in detail the relationship between the two variables.

Dividing the past sixty years into twenty and thirty year segments, we found that the stock market and GDP related to each other as the theoretical models predict, regardless of our definition of asset returns. This relationship tested significantly under FMOLS, although it is unclear whether the magnitude of the relationship has changed. We estimate that a 1% change in output is correlated with approximately a 2% change in assets. It appears that our theoretical models are correct, and in fact the return on assets behaves similarly to what the Solow growth model predicts.

Our results also lead us to ask other questions about the relationship between economic growth and the return on capital. The period between 1967-1986 deserves more examination. The unhinging of the relationship between asset returns and output growth that our results indicate does not fit well with our models. A more comprehensive reason for why this deviation occurred would also be useful when considering future scenarios where this dissociation might return.

Additionally, it would be interesting to extend this analysis to other countries. Barro (1990) discovered a relationship that may help us focus this investigation: Canada's macroeconomic variables were better predicted by the U.S. stock market than its own stock market. Rather than examine many countries, we can focus our analysis on a handful of economic regions to explore how asset returns and economic growth are related—although the EU is relatively new, it provides a very natural starting point. This region analysis may require creating a sphere of influence map that detailed what countries are the centers of their economic region. Campbell's (1999) examination of the many similarities between the United States and other nations sets up this extension as a good check on how well our models generalize. Most importantly we can examine whether the multiplier effect from the Solow model corresponds to each country's or region's saving rate—it may be possible to distinguish which macroeconomic model's predictions are most valid because of the variation of labor supply, capital investment, and savings across many nations.

Our analysis could also become more robust with the specification of a third model. We mentioned very early on that our data might be subject to biases because of business cycles and countercyclical monetary policy. In order to work around this we mentioned a framework proposed by Laubach (2005) where the independent and dependent variables are forward projections. In order to use this framework to test output growth and asset returns we would require projections of future output growth along with future asset price movements. The CBO has five-year forward quarterly projections for GDP dating back to 1991 and ten-year forward projections beginning in 1996, and so this type of analysis will be more feasible as time progresses. Forward projections for asset returns are not as easy find. A possible proxy for future asset returns is the inverse of a price to earnings ratio, or the yield, on an index. Since we are assuming all known information is accounted for in a stock's price, the yield on an index should be the expected future return—the one issue is that since this is a current measure it may be biased by the business cycle and the Fed.

This paper has been partially motivated by the questionable accuracy of government projections of future asset returns and output growth. Given that we believe the assumption of continuing high asset returns to be unreasonable, it may be interesting to examine how often the OCACT has made similar seemingly bad assumptions in the past. In this case the bad assumption makes the proposal for Social Security privatization, a major reform, appear much better than it actually is. This analysis simply requires additional research and will greatly strengthen the background and motivation for exploring this topic. There are ramifications for policy and decision makers if government projections have consistently been wrong.

If it is truly the case that slowing output will occur along with lower returns on assets, the OCACT must recalculate the costs of privatized accounts for Social Security. The lower return on capital means that without raising taxes it will be much harder to maintain benefits and pay off the debt incurred by switching regimes. The free lunch that proponents of private accounts are expecting may actually be quite costly.

References

Baker, D., J. B. Delong, and P. R. Krugman (2005): "Asset Returns and Economic Growth," *Brookings Papers on Economic Activity,* 1.

Barro, R. J. (1990): "The Stock Market and Investment," *The Review of Financial Studies,* 3(1), 115-131.

Benzoni, L., P. Collin-Dufresne, and R. S. Goldstein (2006): "Portfolio Choice Over the Life-Cycle when the Stock and Labor Markets are Cointegrated," *SSRN*. Online: [http://ssrn.com/abstract=875984].

Bernanke, B., J. Y. Campbell, B. M. Friedman, and L. H. Summers (1988): "Is There a Corporate Debt Crisis?" *Brookings Papers on Economic Activity*, 1988(1), 83-139.

Bodie, Z. (1995): "On the Risk of Stocks in the Long Run," *Financial Analysts Journal.*

Bureau of Economic Advisors (2006): National Income and Product Account Tables: Years 1947-2006).

Burns, A. F. and W. C. Mitchell (1938): *Statistical Indicators of Cyclical Revivals*, NBER Bulletin 69, NY. Reprinted in *Business Cycle Indicators*. G.H. Moore, ed. 1961. Princeton U. Press.

Campbell, J. Y. (1999): "Asset Prices, Consumption, and the Business Cycle," in *The Handbook of Macroeconomics*, 1, J. B. Taylor and M. Woodford, eds. Amsterdam: Elsevier, 1231-1303.

Cheung, Y.-W. and L. K. Ng (1998): “International evidence on the stock market and aggregate economic activity,” *Journal of Empirical Finance*, 5(3) 281-296.

Congressional Budget Office (2004): “The Outlook for Social Security.” Online: [http://www.cbo.gov/showdoc.cfm?index=5530&sequence=0].

Darlin, D. (2007): “A Contrarian View: Save Less And Still Retire With Enough,” *The New York Times*, 27 January 2007.

Diamond, P. A. (2000): “What Stock Market Returns to Expect for the Futrure?” *Social Security Bulletin*, 63(2), 38-52.

Congressional Budget Office (2006): “2006-Historical Qstar.”

Fama, E. F. (1981): “Stock Returns, Real Activity, Inflation and Money,” *American Economic Review*, 71, 545-565.

Fama, E. F. and French, K. R. (1992): “The cross-section of expected stock returns,” *Journal of Finance*, 47, 427-465.

Fama, E. F. and French, K. R. (2000): "Disappearing Dividends: Changing Firm Characteristics Or Lower Propensity To Pay?" *Center for Research in Security Prices*, 509.

Feldstein, M. (1996): “The Missing Piece in Policy Analysis: Social Security Reform,” *American Economic Review,* 86, 1-14.

Fischer, S. and R. C. Merton (1984): “Macroeconomics and Finance: The Role of the Stock Market,” *Carnegie-Rochester Conference Series on Public Policy*, 21, 57-108.

Gale, W. G., P. R. Orszag, E. M. Engen, and C. L. House (2004): “Budget Deficits, National Saving, and Interest Rates [and Comments and Discussion],” *Brookings Papers on Economic Activity,* 2004(2), 101-210.

Greene, W. H. (2002): *Econometric Analysis*, Prentice Hall.

Harris, R. and R. Sollis (2003): *Applied Time Series Econometrics*, New Jersey: Wiley.

Harvey, C. R. (1989): “Forecasts of Economic Growth from the Bond and Stock Markets,” *Financial Analysts Journal*, 45(5), 38-45.

Khitatrakun, S., J. K. Scholz, and A. Seshadri, (2006): “Are Americans Saving “Optimally” for Retirement?” *Journal of Political Economy*, 114.

Laubach, T. (2005): “New Evidence on the Interest Rate Effects of Budget Deficits and Debt,” Washington: Board of Governors of the Federal Reserve System.

Levine, R. and S. Zervos (1996): “Stock Markets, Banks, and Economic Growth,” *Journal of Economic Literature*.

Lucas, D. and S. Zeldes (2006): “Valuing and Hedging Defined Benefit Pension Obligations—

The Role of Stocks Revisited," Online: [http://finance.wharton.upenn.edu/department/Seminar/2006Fall/macro/Zeldes-macro110706.pdf].

Mehra, R. and E. C. Prescott (1985): "The Equity Premium: A Puzzle." *Journal of Monetary Economics*, 15, 145-161.

Mehra, R. (2003): "The Equity Premium: Why is it a Puzzle?" *Financial Analysts Journal,* 54-69.

Romer, D (2000): *Advanced Macroeconomics,* New York: McGraw-Hill, 2nd edn.

Samwick, A. (2007): "Optimal Saving?" Online: [http://voxbaby.blogspot.com/search/label/saving].

Shiller, R. (2006): Stock Market Dataset, *Irrational Exuberance*, Online: [http://www.econ.yale.edu/~shiller/data.htm].

Stock, J. H. and M. W. Watson (2003): "Forecasting Output and Inflation: The Role of Asset Prices," *Journal of Economic Literature*, 41, 788-829.

Vassalou, M. (2003): "News related to future GDP growth as a risk factor in equity returns," *Journal of Financial Economics,* 68, 47-73.

Wikipedia (2006): "Dow Jones Industrial Average," Online: [http://en.wikipedia.org/wiki/Dow_Jones_Industrial_Average].

Wikipedia (2006): "S&P 500," Online: [http://en.wikipedia.org/wiki/S%26P_500].

Wilcox, D. W. (2006): "Reforming the Defined-Benefit Pension System in the United States," *Brookings Panel on Economic Activity*.

VII. Figures and Tables

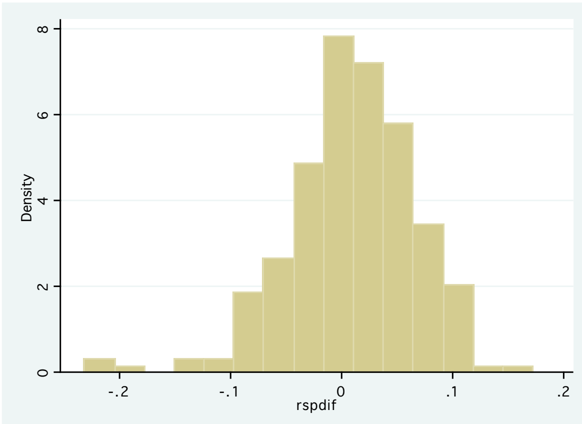

Figure 1: S&P500 growth rates

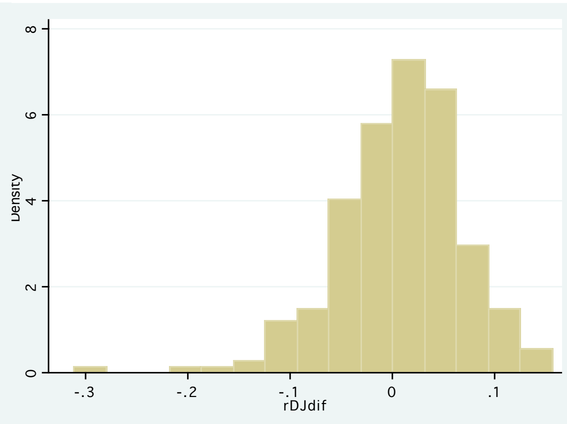

Figure 2: DJIA growth rates

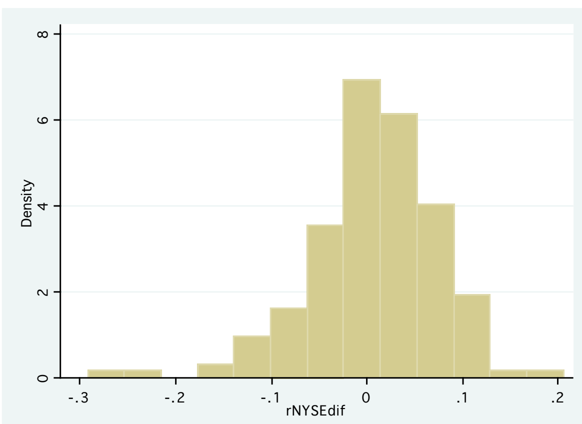

Figure 3: NYSE growth rates

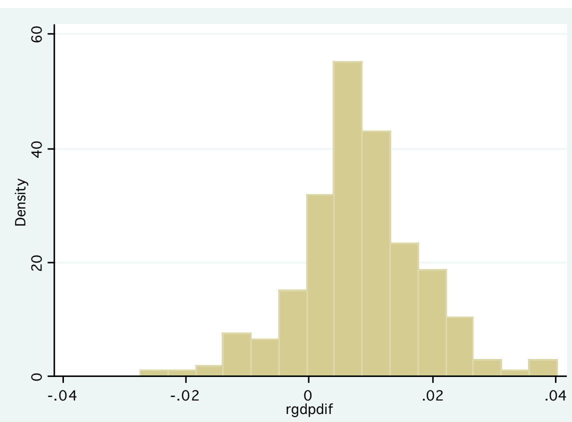

Figure 4: real GDP growth rates

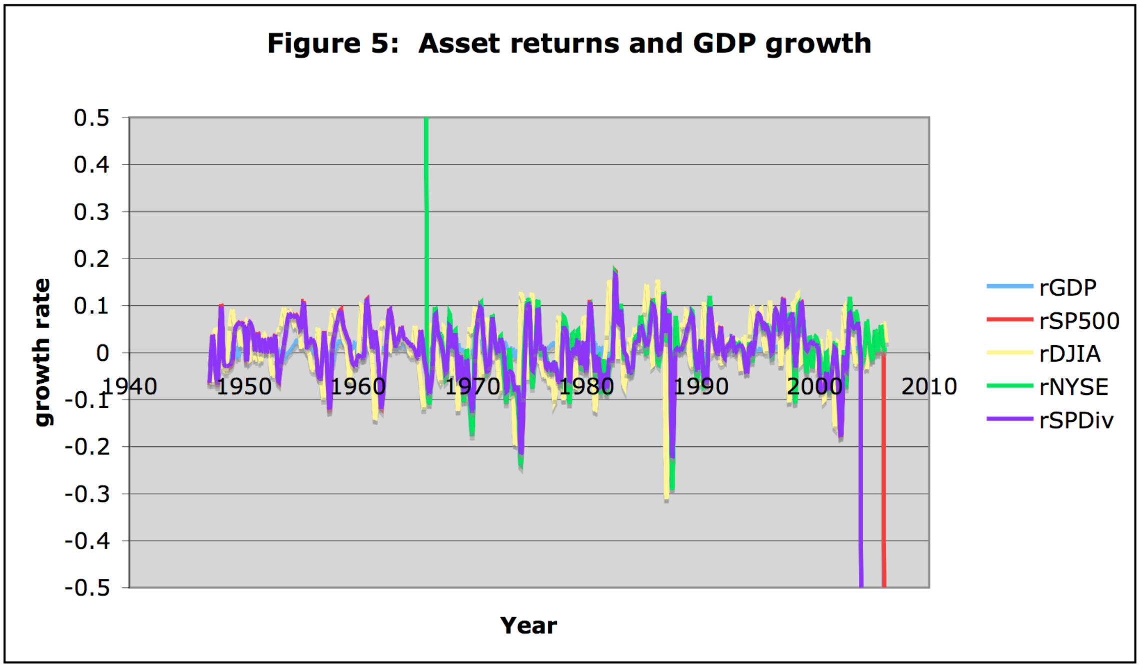
Figure 5: Asset returns and GDP growth
growth rate
0.5
0.4
0.3
0.2
0.1
0
-0.1
-0.2
-0.3
-0.4
-0.5
1940
1950
1960
1970
1980
1990
2000
2010
Year
rGDP
rSP500
rDJIA
rNYSE
rSPDiv

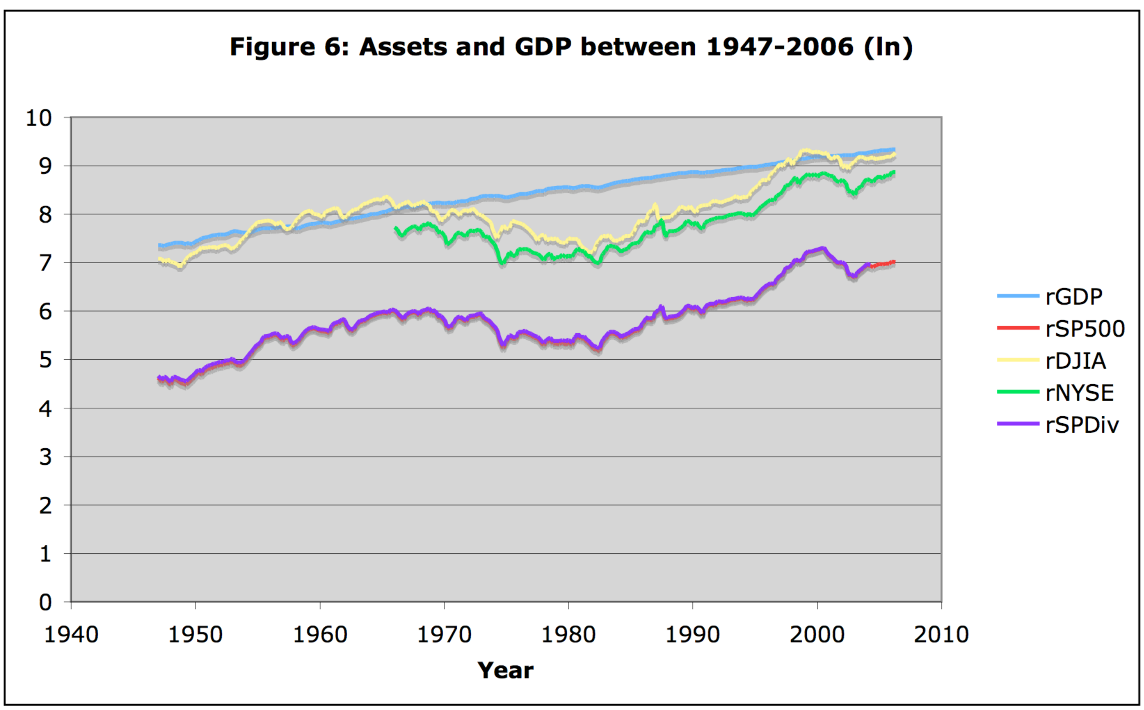
Figure 6: Assets and GDP between 1947-2006 (ln)
10
9
8
7
6
5
4
3
2
1
0
1940
1950
1960
1970
1980
1990
2000
2010
Year
rGDP
rSP500
rDJIA
rNYSE
rSPDiv

Table 1

	Growth rate time-trends real GDP	S&P 500	Dow Jones	NYSE
trend	-0.000011 [.000012]	0 [.000063]	0.000025 [.000059]	0.000218* [.000124]

Table 2 S&P 500 Regressions

	1947-2006	1947-1966	1967-1986	1987-2006	1947-1976	1977-2006
trend	0.000019 [.000057]	-0.00012 [.00026]	0.00059* [.00039]	-0.00018 [.00033]	-0.00031* [.00017]	0.000073 [.00018]
real GDP	1.678*** [.3722]	1.092*** [.3476]	1.8113** [.6969]	4.2245*** [1.05]	1.4*** [.4365]	2.312*** [.6240]

Table 3 Dow Jones Industrial Average Regressions

	1947-2006	1947-1966	1967-1986	1987-2006	1947-1976	1977-2006
trend	0.000018 [.00006]	-0.00009 [.00028]	0.0009** [.00036]	0.000015 [.00039]	-0.00038** [.00019]	0.00011 [.00018]
real GDP	-0.60133 [.4486]	-0.7497 [.538]	-0.9603 [.7619]	0.9723 [1.159]	-0.4343 [.5081]	-1.054 [.7208]

Table 4 NYSE Regressions

	1966-2006	1947-1966	1967-1986	1987-2006
trend	0.00022* [.00012]		0.00061* [.00033]	0.00003 [.0003]
real GDP	1.832*** [.6747]		1.513** [.7594]	2.844** [1.4103]

[Newey-West standard errors]
*- Significant at 10%
**- Significant at 5%
***- Significant at 1%

Table 5 ADF test

Series	1947-2006	1947-1976	1977-2006	1947-1966	1967-1986	1987-2006	1947-1986
rgdp	-1.22	-0.91	-0.05	-0.12	-0.32	0.21	-0.93
rsp500	-0.98	-1.85	-0.65	-1.04	-1.91	-0.8	-1.97
rdjia	-0.78	-2.22	-0.43	-1.48	-1.51	-1.41	-2.1
	1967-2006						
rnyse	0.05	---	-0.4	---	-1.92	-0.44	---
SP+Div	-1.02	-1.88	-0.61	-1.06	-1.9	-0.77	-2

t-stat shown Critical Values: 10% = -2.58, 5% = -2.9, 1% = -3.5

Table 6 ADF test on residuals

Series	1947-2006	1947-1976	1977-2006	1947-1966	1967-1986	1987-2006	1947-1986
rsp500	-1.5	-0.73	-2.39	-1.51	-1.62	-1.35	-1.69
rdjia	-1.28	-1.23	-2.56	-2.52	-0.46	-1.17	-1.95
	1967-2006						
rnyse	-2.2	---	-3.07*	---	-1.56	-2.02	---
SP+Div	-1.49	-0.72	-2.57	-1.51	-1.6	-1.46	-1.7

t-stat shown Critical Values: 10% = -2.58, 5% = -2.9, 1% = -3.5

Table 7 FMOLS Regressions

Series	1947-2006	1947-1976	1977-2006	1947-1966	1967-1986	1987-2006	1947-1986
rsp500							
rgdp	1.2**	1.23	2.45***	2.04***	-1.26	2.38**	0.65*
	[2.2]	[.78]	[7.61]	[3.5]	[-.93]	[2.08]	[1.32]
rdjia							
rgdp	1.16*	1.1	2.62***	1.85***	-1.88	2.64***	0.37
	[1.55]	[1.04]	[11.46]	[6.47]	[-.54]	[3.05]	[.68]
rnyse	1967-2006						
rgdp	1.38**	-	2.32***	-	-1.04	2.33***	-
	[2.95]		[13.61]		[-.84]	[5.47]	
SP+Div	1.15**	1.2	2.52***	2***	-1.24	2.7***	0.64*
rgdp	[2.06]	[.77]	[8.59]	[3.62]	[-.91]	[2.46]	[1.35]

[t-statistics] *- significant at 10%, **- significant at 5%, ***- significant at 1%