Growth and Value Investing: Understanding the Sources of Excess Returns

Thomas K. Philips
Chief Investment Officer
Paradigm Asset Management Company
New York City

Growth and value investing in theory (as described in the academic literature) is quite different from growth and value investing in practice (as carried out by investors). A theory, the Fundamental Theorem of Growth Equality, is proposed that states that over time, the earnings growth of value and growth stocks is about the same. The index construction methodology of growth and value indexes (notably the periodic rebalancing to equal capitalizations in each index) is largely responsible for this equality in long-term growth rates, which has implications for the value premium, the size premium, and the importance of dividend equivalents in growth and value investing.

Value stocks over the long run, and I stress the long run, have higher returns than growth stocks. But as I will show in this presentation, a number of very subtle issues are associated with value investing. I will start by addressing the history of value investing. Then, I will discuss value investing in theory (the *Journal of Finance* version of value investing) and value investing in practice (what portfolio managers do for a living, which by and large means buying an active product that looks like an index or an actual index product). I will then go on to address EPS growth versus economic growth for growth and value investing, which will lead me to the Fundamental Theorem of Growth Equality. To illustrate this theorem, I will provide several examples. As I will show, the process of index construction has enormous implications for the nature and the source of value returns and has enormous implications for how return differentials between two or more indexes are realized. Finally, I will touch on the implications for the size premium and index construction methodologies.

Evolution of Value Investing

A very rough definition of value investing is the business of buying businesses for less than they are worth, an idea initially associated with Benjamin Graham. A crude expansion of that definition is that value investors tend to be balance sheet investors and growth investors tend to be income statement investors. That statement captures the essence of the two philosophies, but the lines are becoming increasingly blurred.

The notion of what constitutes value investing, what constitutes buying a business for less than it is worth, has clearly evolved over time. Benjamin Graham was perhaps the first person to use the term "value investing," and his notion of value was inextricably linked with buying companies for less than their net current assets. Executing a value strategy in Graham's eyes meant that investors would buy stocks that were selling for less than two-thirds of their net current assets and sell them as soon as they rose to roughly the value of their net current assets. Investors played this game for many years until the bull market of the 1960s and early 1970s caused these opportunities to disappear. Graham said at that time that the net current assets game seemed to be over but that value could be thought of in other ways, such as price to book. Of course, this finding echoes in the work of Fama and French 20 years later in the 1990s.

So, in what other ways can value can be defined? Sanjoy Basu showed in the late 1970s that low-P/E stocks performed better than high-P/E stocks. Thus, by making this observation, he captured the notion of value via P/E. In recent years, the accounting folks—Jim Ohlson, Gerald Feltham, Charles M.C.

Lee, Bhaskaran Swaminathan, and so on—have combined book to price (B/P) and earnings to price (E/P) into a single value measure, *V*. They have shown that the ratio of *V* to price is a terrific value variable.

The literature of value investing is so rich that the mention of these few studies cannot do it justice. But at an aggregate level, at an index level, all of these definitions of value are roughly similar.

Value Investing: Theory vs. Practice

In the academic—the typical *Journal of Finance*—approach to value, a value variable (B/P, P/E, cash flow to price, dividend yield, etc.) is used to sort and rank a universe of stocks (maybe the 1,000 biggest stocks in the world or the entire Compustat universe) from high to low. The deep value stocks would be at the top, and the growth stocks would be at the bottom. The returns of the top decile (value) and the bottom decile (growth) are then subtracted to get a value spread. The result is the value premium in the *Journal of Finance* sense.

But that is not the way value investors invest in practice. Typically, value investors invest through indexes that are all constructed somewhat differently, depending on the index provider. Each index provider takes a value variable (such as B/P, which is what Standard & Poor's uses to construct its indexes, or a mixture of B/P and earnings growth, which is what Russell uses to construct its indexes, or cash flow to price or whatever the variable might be) and sorts the universe of stocks using this value variable. The index provider then breaks the universe of stocks into two halves (two indexes) with equal market caps but an unequal number of names. One half includes the more value-oriented stocks, the value index, and the other half includes the more growth-oriented stocks, the growth index.

Periodically—every three months, six months, or a year—depending on the index provider, the value and growth indexes are rebalanced to equal market caps. This rebalancing is fundamental to many of the observed interrelationships in the value index versus growth index area. For example, the expected value premium is the difference between the expected returns of the value index and the growth index, and the realized value premium is the return of the value index minus the return of the growth index.

Several properties of these indexes bear exploration, and I will discuss them in this presentation. One issue involves the earnings growth of each index. If the S&P 500 Index is broken into a growth index and a value index (S&P/Barra 500 Growth Index and S&P/Barra 500 Value Index), in aggregate, their earnings growth rate must equal that of the S&P 500. But does the earnings growth of one grow faster than the others? Is one riskier than the other? Are the returns of one more volatile than the others? Are the earnings of one more volatile than the others? How do the expected returns of the two indexes compare? How does rebalancing affect all of these properties? And finally, how is the value premium—the difference in return between the growth and value indexes—realized over the long term by investors? Is it realized in the form of capital gains? In the form of dividends? Some combination of the two?

EPS Growth vs. Economic Growth

Before I delve into the properties of earnings, earnings growth, and EPS growth, I need to define "aggregate earnings" and "per share earnings." Aggregate earnings are the earnings that are associated with every company in the economy at any point in time. When I talk of earnings growth, I am talking about the growth of aggregate corporate earnings, which include the earnings generated by new companies that did not exist in the prior time period. The growth in aggregate earnings is the growth in the earnings of the existing economy and the new economy, otherwise known as "entrepreneurial capitalism."

When I refer to "per share earnings" or "EPS growth," however, I mean the growth in the earnings of a *unitized* portfolio—a portfolio that does not receive inflows of capital (e.g., an S&P 500 index fund). The growth of per share earnings can be no larger than the growth of aggregate earnings because per share earnings do not have the benefit of the influx of new cash flows. If an investor has a buy-and-hold strategy, that investor cannot realize the true earnings growth rate of the economy. The investor can only realize the portfolio's per share earnings growth rate, which cannot be higher than the aggregate earnings growth rate. That distinction is subtle, but it is important.

The rate of growth of aggregate earnings has to be the same as the rate of growth of nominal GDP. Why? In an open economy with a competitive labor market, the fraction of national income that accrues to capital, as opposed to labor, is roughly constant. Although it has varied between 4 percent and 8 percent over time, it is now about 6 percent and does not seem to be increasing. Therefore, the theory that a company's earnings can grow faster than the economy because of rapidly growing foreign subsidiaries, lower interest costs, and increased productivity is rubbish. The benefits associated with such growth typically accrue to employees, not to shareholders.

Shareholders tend to get a fixed slice of the national income pie. Warren Buffet wrote two articles on this topic that appeared in *Fortune* magazine, one in late 1999 and another in late 2001, both of which are well worth reading.[1]

And the growth of per share earnings must be the same over the long term as the growth of aggregate earnings. The reason is very simple. Suppose assets are always deployed efficiently in an economy. So, when a company earns more than it can reinvest in its own business, it pays out the excess earnings to investors in the form of dividends. Using those dividends, investors can finance new investment. And if the dividend yield—the free cash flow yield, to be more precise—is high enough, investors can finance all the new investment in the economy and can ensure that a unitized portfolio contains the entire economy. In other words, this assumption is equivalent to having existing firms finance all of the entrepreneurial activity in an economy.

This scenario does require certain elements, such as a bankruptcy court that redeploys in an effective manner the assets in businesses that go bust. But if these functions are efficiently executed, and by and large they are, then the growth of per share earnings will equal the growth of earnings. My claim has not been true historically, in part because the bankruptcy court has been erratic in its operations and in part because the capital inflows into most developing countries have been extremely large. Rob Arnott and Peter Bernstein show that over the long term in both the United States and the United Kingdom, per share earnings have grown roughly in line with GDP per capita, not GDP.[2] But in the future, I believe the growth rate of the two quantities will be similar.

The empirical evidence can help clarify my assertion. **Figure 1** shows nominal GDP growth as well as normalized S&P 500 earnings growth and EPS growth for the period December 1976 to December 2001. The earnings of the S&P 500 grew faster than the economy over this time period, which is not surprising. When the S&P 500 removes a company from the index, it is typically a small company in the throes of bankruptcy or one that has shrunk to a shadow of its former self. When the company is replaced in the index, it is replaced with a relatively large company. Therefore, it is not surprising that S&P 500 earnings grew faster than the economy. Per share earnings, in contrast, grew roughly in line with the economy except in 2001. This earnings implosion in 2001 should be ignored, however, because it was caused in large part by accounting issues and the true economic earnings of the S&P 500 in 2001 are as yet unknown. Thus, the evidence over this period shows that per share earnings grew roughly in line with nominal GDP.

[1]Carol Loomis and Warren Buffett, "Mr. Buffett on the Stock Market," *Fortune* (22 November 1999):212–220. Warren Buffet, "Warren Buffet on the Stock Market," *Fortune* (10 December 2001):80–94.

[2]Robert Arnott and Peter Bernstein, "What Risk Premium is 'Normal'?" *Financial Analysts Journal* (March/April 2002):64–85.

Figure 1. Normalized S&P 500 Earnings and EPS Growth and GDP Growth: December 1976–December 2001

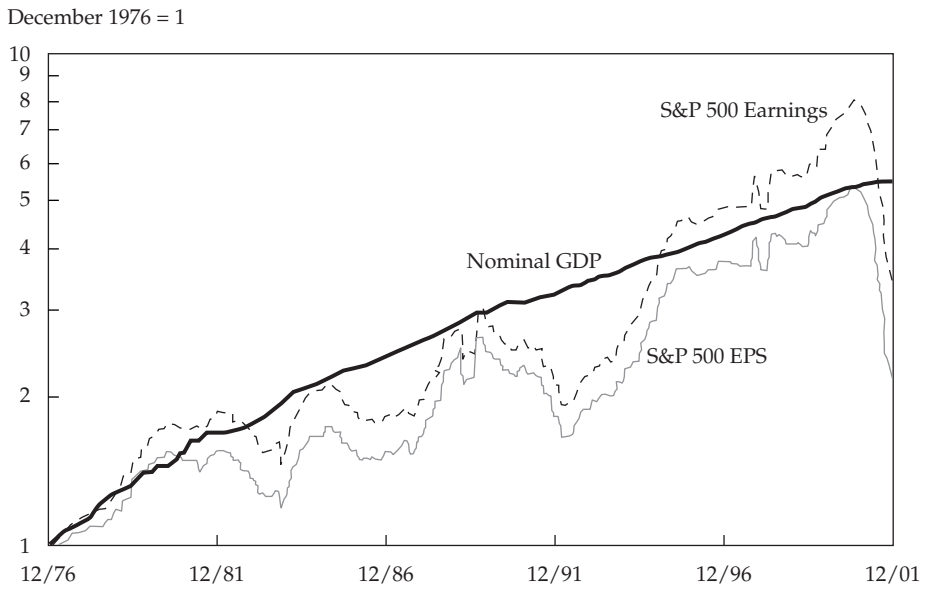

Earlier, I posed several questions that I would like to address now: How do the earnings of the growth and value indexes relate to the earnings of the broad market? In particular, does the S&P 500 growth index experience faster earnings growth and the S&P 500 value index experience slower earnings growth than the market? Investors instinctively feel that S&P 500 value stocks have slower earnings growth than S&P 500 growth stocks, but as it turns out, the reality of the relationship is much more nebulous. I have constructed a time series of EPS and GDP data from 1974 to 2000—a mixture of data from Bloomberg (index values), Barra and Morgan Stanley's Quantitative Strategies Group (P/Es), and the Federal Reserve's FRED database (GDP). **Figure 2** shows once again the growth of nominal GDP, which is compared with the growth of per share earnings for the S&P 500 growth and value indexes. Notice that the EPS of both the growth and value indexes are not growing faster than the economy but roughly at the same speed. Thus, growth earnings and value earnings appear to grow at the same rate over the long term, although they have very different patterns of behavior. The earnings of the value index in particular are extremely volatile. In the 1991–92 recession, nominal per share earnings for the value index retreated to levels last seen in 1976. And in 2001, they again dropped to 1976 levels. So, value earnings implode in recessions and improve sharply in recoveries, but over the long run, they grow at about the same rate as the earnings of the growth index. Notice that the earnings of the growth index are surprisingly stable and track GDP quite well.

A good way to measure the difference between the growth and value growth rates is to do a *t*-test for the difference in means. **Table 1** shows data from three starting points and does not include 2001 because 2001 was such a horrid year for earnings that it cannot be used to draw any reasonable conclusions. The S&P growth index saw its earnings grow annually at an arithmetic average rate of 8.29 percent and a geometric rate of 7.75 percent from 1974 to 2000. The value index had growth of 9.25 percent (arithmetic)

Table 1. EPS Growth: S&P 500 Growth versus S&P 500 Value

	Annualized Growth			Standard
Period/Index	Arithmetic	Geometric	Volatility	Error
December 1974 to December 2000				
S&P growth	8.29%	7.75%	11.22%	
S&P value	9.25	6.45	25.26	
Difference	–0.96	1.31	20.82	4.08%
December 1975 to December 2000				
S&P growth	8.62	8.43	11.32	
S&P value	10.30	7.83	25.20	
Difference	–1.67	0.60	20.92	4.27%
December 1976 to December 2000				
S&P growth	7.75	7.60	10.67	
S&P value	9.84	7.28	25.63	
Difference	–2.09	0.32	21.27	4.43%

Figure 2. Normalized S&P 500 Value and Growth Index EPS Growth and GDP Growth, December 1976–December 2001

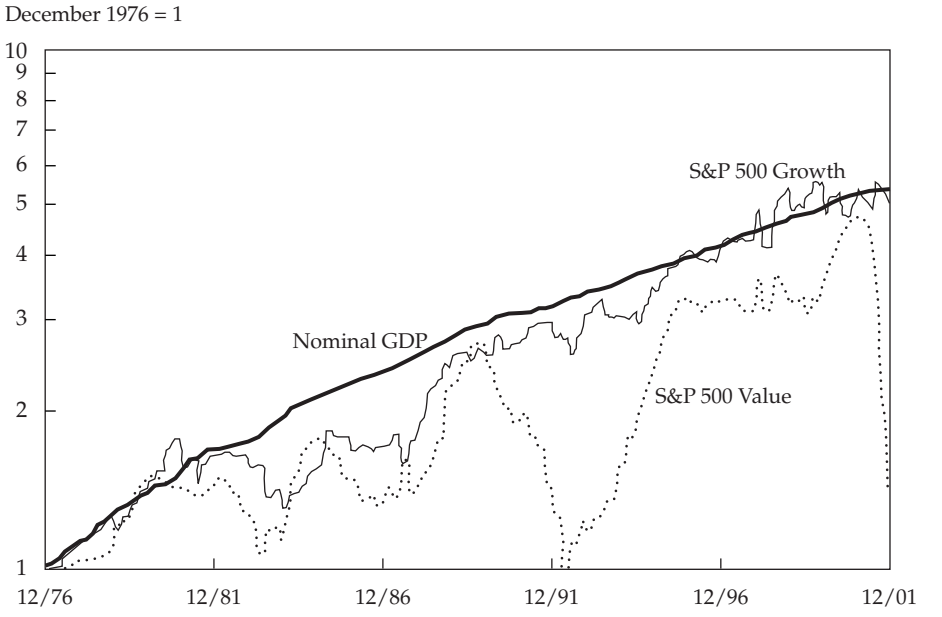

and 6.45 percent (geometric) for the same period. The difference between the growth and value indexes is about 1 percentage point (pp) expressed in both geometric and arithmetic terms. But how significant is 1 pp? The standard error of the difference is about 4 percent. So, the difference is one-fourth of a standard error away from zero (i.e., not really distinguishable from zero).

For the 1975–2000 period, the arithmetic difference in growth rates is 1.67 pps and the geometric difference is 0.60 pps; the standard error is still about 4 percent (or one-third of a standard error). And from 1976 when the recession was clearly over, with both the growth and value indexes starting from roughly the same base, the difference (geometric) in their earnings growth rate is 0.32 pps a year. The difference in growth rates is one-tenth or one-fifteenth of a standard error for the 1976–2000 period (i.e., completely indistinguishable from zero). Therefore, the empirical evidence supporting the equality of growth rates is very strong.

Fundamental Theorem of Growth Equality

Investors might reasonably wonder why the earnings growth of value and growth stocks is about the same. I have a theory that explains it, which is succinctly expressed as the Fundamental Theorem of Growth Equality:

> If all stocks have time-invariant growth rates, then all indexes that are rebalanced back to a fixed fraction of the market's capitalization, regardless of their construction methodology, must experience the same long-run rate of earnings and EPS growth as the market.

Following is an informal proof of this theorem.

Suppose I have a growth index and a value index. If the growth index has a faster earnings growth rate than the value index, then between rebalancings, I would expect the capital gains of the growth index to be higher than those of the value index, assuming no P/E bubble. But what happens when the index provider rebalances? Remember the index provider has to rebalance to equal market caps in each of the two indexes, so in the next rebalancing, a few growth companies will be kicked over into the value index. What happens to the value index? Its growth rate will be pushed up a little bit as a result. If the growth rate continues at the same fast pace, once again, at the next rebalancing some companies will be transferred into the value index, pushing up its growth rate even further. The net result is that after a while, the two growth rates will be perfectly equalized. Although huge discrepancies can exist for short periods of time, as I showed earlier, over the long term, the two rates must converge.

An immediate implication can be drawn from this theorem. If both indexes have the same long-run rate of earnings growth and if a P/E bubble does not occur, the capital gains component of return for these two indexes must be identical. Furthermore, any difference in return between the two indexes is entirely the result of a difference in their dividends or dividend equivalents. By "dividend equivalents" I mean all the possible uses of free cash flow: dividends, share buy backs, and takeovers.

Examples

Each of the following four examples reaches the same conclusion about the growth rate of earnings for growth and value stocks, but each starts with a different premise. I have called the examples "Only Cisco," the "Last Days of Disco," "Living with Risco," and the "Slow Death of Misco." **Table 2** lists information on the five companies used in the examples.

Cisco is a growth company with $100,000 a year in earnings, a growth rate of 10 percent a year, and a P/E of 100. It has 10,000 shares outstanding at a price of $1,000 each, which translates into an initial market cap of $10 million. Cisco grows at 10 percent every year; it never misses.

Disco, on the other hand, has $1 million in earnings, but those earnings will never grow in perpetuity. The company has a constant growth rate of zero and an initial market cap of $10 million.

Table 2. Company Information for Example Companies

Company	Earnings	*g*	P/E	Shares	Price	Growth Pattern
Cisco	$100,000	10%	100	10,000	$1,000	Constant 10%
Disco	$1,000,000	0	10	10,000	$1,000	Constant 0%
Risco	$500,000	10	20	10,000	$1,000	Alternating 0%, 21%
Misco	$1,000,000	–10	10	10,000	$1,000	Constant –10%
Mini C	$10,000	–10	100	1,000	$1,000	Constant –10%

Notes: Earnings are annual earnings; *g* = annual growth rate.

Risco is a risky growth stock. It has $500,000 in earnings and a P/E of 20. It too has an initial $10 million market cap, and although its long-term growth rate is 10 percent, the growth rate is not realized as a constant 10 percent every year. The growth rate alternates; for example, in some years the growth rate is zero, and in some years it is 21 percent, which averages out to 10 percent a year compounded.

Misco is a company that constantly misses its earnings forecast. It has $1 million in earnings today, but the earnings are declining at 10 percent a year. It has a P/E of 10 and also a $10 million market cap.

Mini C is a company that spun out of Cisco. Mini C is a small company with only $10,000 in earnings and a $1 million market cap. It has a growth rate of –10 percent, but it has a P/E of 100 because it has a new business model (as did most dot-coms during the bubble!).

In the examples, value is defined by a low P/E. So, Cisco is a growth stock, and Disco and Risco are value stocks. I will use the classic definition of value and growth indexes. That is, the indexes have equal market caps after rebalancing. In addition, I assume a stock can be partially allocated to each of the two indexes, which allows me to get around the granularity issue. Finally, indexes are rebalanced periodically.

Only Cisco. In this first example, the entire universe is only one stock, Cisco. Therefore, the growth index and the value index must each hold 5,000 shares of Cisco. This case is admittedly trivial. The growth index and the value index are identical, so naturally they have exactly the same long-run rate of earnings growth and EPS growth. Everything works out very nicely; there is perfect equilibrium, and no transfers are needed at rebalancing.

The Last Days of Disco. Disco is a company that has no earnings growth. It has an initial $10 million market cap, but its earnings never grow. Initially, the growth index is pure Cisco (10,000 shares of Cisco with a market cap of $10 million), and the value index is pure Disco (10,000 shares of Disco with a market cap of $10 million). But as time goes buy, Cisco's market cap keeps rising because its earnings are growing at 10 percent a year, while Disco's market cap stays the same.

Because the indexes are always rebalanced back to an equal market cap, a little bit of Cisco will begin to enter the value index. **Figure 3** shows the allocations to Cisco and Disco for the entire market. Initially, the market is half Cisco and half Disco. At infinity, the market is pure Cisco, no Disco. But between Time 0 and infinity it is a mixture of Cisco and Disco. **Figure 4** illustrates how the makeup of the growth and value indexes changes from Time 0 to Time 50. The fraction of Cisco in the value index is zero at Time 0 and rises to nearly 100 percent at Time 50. Similarly, the fraction of Disco in the value index is 100 percent at Time 0 and drops almost to zero at Time 50. By the same token, Cisco is always 100 percent of the growth index, and Disco is always 0 percent of the value index.

Figure 5 shows the per share earnings of these two indexes from Time 0 to Time 50. Notice that the growth index is always growing at a 10 percent annual rate. The growth index is always pure Cisco, whose earnings grow steadily at 10 percent a year. And because

Figure 3. Entire Market: Allocations to Cisco and Disco

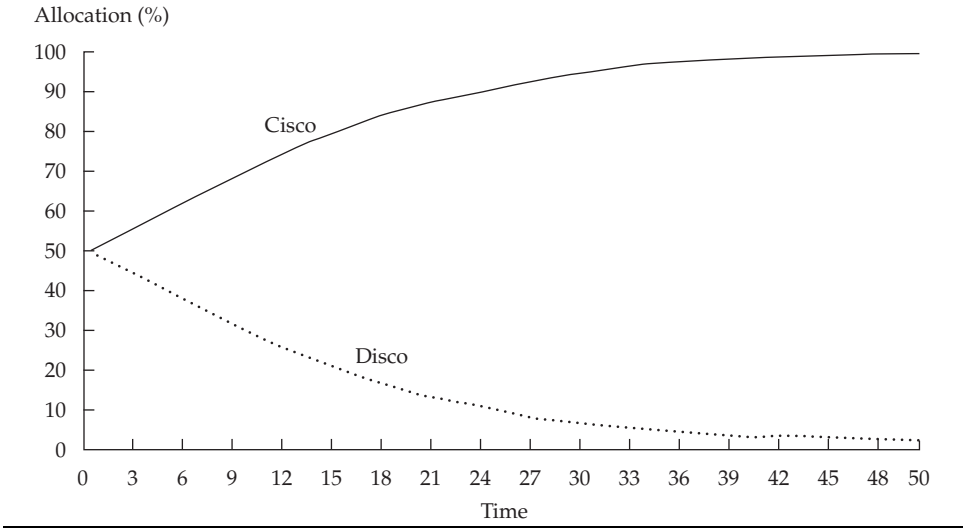

Figure 4. Value and Growth Indexes: Allocations to Cisco and Disco

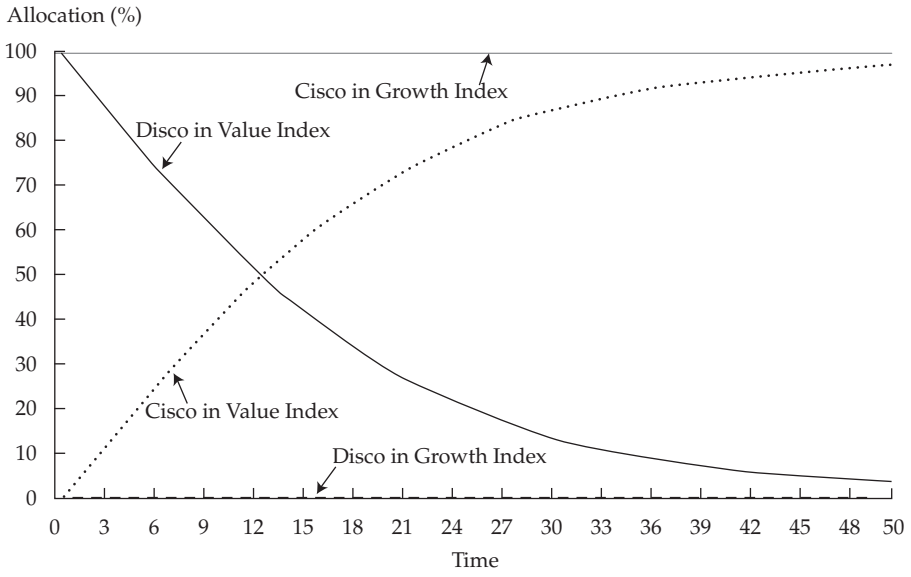

Figure 5. EPS Growth Rates: Value, Growth, and Market

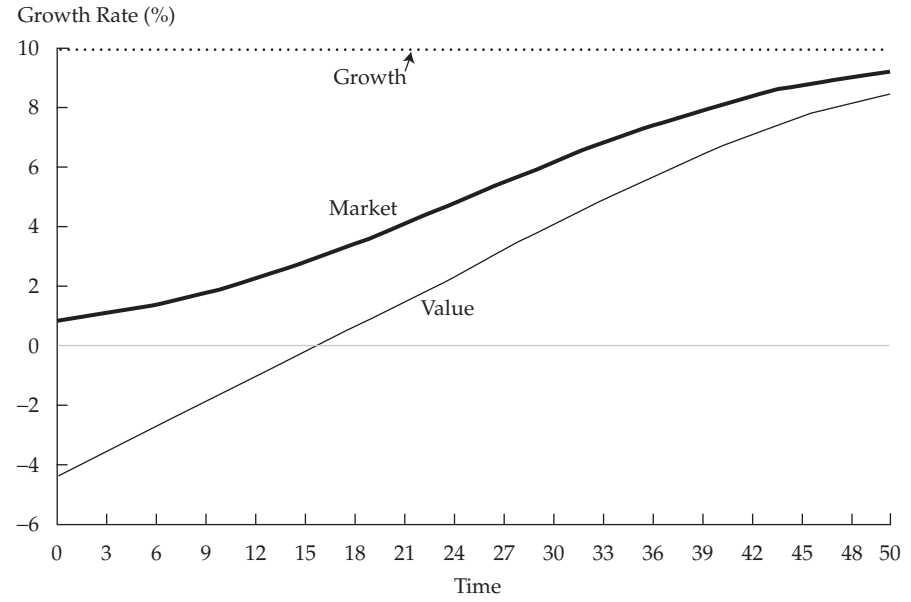

the growth index never holds any Disco, its growth rate must also always be 10 percent a year. The value index, however, starts with a negative EPS growth rate, which then climbs almost to 10 percent a year.

The reason for the beginning negative growth rate for the value index, considering that Disco has a growth rate of zero (not a negative rate), is rather interesting. It turns out that the process of rebalancing has an impact on per share earnings growth. As some of the high-P/E Cisco is moved into the value index with the low-P/E Disco, the P/E of the value index increases. That increase in P/E is equivalent to depressing per share earnings growth. This point is very important: Rebalancing can affect EPS growth. To calculate the long-run rate of EPS growth, one would have to take the rate of EPS growth between rebalancings and multiply that number by the ratio of the P/Es before and after each rebalancing. That product gives the true rate of EPS growth for this rebalanced strategy. The fact that rebalancing reduces EPS growth by reducing P/E changes cannot be overemphasized. So, an index can have a negative EPS growth rate even though the smallest growth rate of any constituent stock is only zero.

In the end, the value index is pure Cisco, and it reduces to my first example, Only Cisco. So, yet again, I have growth and value indexes at the same long-run rate of EPS growth.

Living with Risco. The third example is much more realistic than the first two. The market is composed of three stocks, each with 10,000 shares outstanding—Cisco, Disco, and Risco. Initially, the growth index is all of the Cisco shares plus half the Risco shares, and the value index is all of the Disco shares plus half the Risco shares. As time progresses, Risco's earnings keep growing at 10 percent a year, albeit a noisy 10 percent—alternating between being flat one year and up 21 percent the next—while Disco's earnings remain flat. So, slowly but surely, Risco starts to displace Disco in the value index. As time approaches infinity, Disco's representation in the market becomes completely immaterial, and Cisco and Risco in equal amounts constitute all of the capital in the market. By Time 50, the growth index becomes almost pure Cisco, and the value index becomes almost pure Risco—although a little bit of Cisco pops in and out as Risco's growth rate alternates between 0 percent and 21 percent. And once again, rebalancing equalizes the long-term growth rate of the two indexes.

The EPS growth rates for the growth index, the value index, and the market are very volatile. They oscillate wildly around their long-term average. These enormous oscillations are the result of Disco exhibiting constant growth of zero, Cisco exhibiting constant 10 percent growth, and Risco exhibiting a growth rate alternating between zero and 21 percent. Because the growth index becomes pure Cisco toward Time 50, its EPS growth rate stabilizes. But Cisco's entering and exiting the value index has a huge distorting effect on the EPS growth rate of the value index, which is entirely caused by the change in the P/E of the index. Remember, Cisco has a P/E of 100; Risco has a P/E of 20. The long-term EPS growth rate for the value index is 10 percent, but in the short term, the growth rate is extremely noisy. This tendency for rebalancing to inject noise into the EPS growth rates of the value and growth indexes significantly complicates the proof that the long-term equalization in the subindex growth rates will hold regardless of the pattern of EPS growth rates experienced by the universe of stocks. If it were not for this noise, the proof would go cleanly.

The Slow Death of Misco. The last example illustrates the possibility of constructing a model in which the growth rate of one index is higher than that of the market. In this example, the market consists only of Misco, with a growth rate of –10 percent a year and a P/E of 10, and Cisco, with a constant growth rate of 10 percent a year. So, if the market contains only Cisco and Misco, the market exhibits zero growth. In addition, every year Cisco spins out an underperforming division called Mini C. Mini C has a –10 percent growth rate, but thanks to its new business model, Mini C has a P/E of 100.

From Time 0 to Time 100, the market allocations to each stock in the universe move from an initial equal proportion of Cisco and Misco to a final allocation of equal proportions of Cisco and Mini C. The value index, originally composed solely of Misco, ends with a 100 percent allocation to Mini C, and the growth index begins and ends the period with a 100 percent allocation to Cisco.

In this universe, the growth rate of earnings for the market and the subindexes is zero. And in terms of EPS growth, although the market shows zero growth, the growth index grows at 10 percent and the value index at –10 percent. Thus, in spite of the index rebalancing, the constant flow of poorly performing companies from the growth index to the value index creates a difference in their long-term EPS growth rates. So, in this example, the EPS growth rate of the growth index is greater than that of the value index and of the market over the long term.

Lessons. What these examples indicate is that the necessary and sufficient conditions for equal EPS growth may differ substantially and that the model of Cisco, Disco, and Risco is the right model for representing actual market behavior. Furthermore, because of P/E effects from rebalancing, short-term EPS growth patterns can be surprisingly volatile.

Formal Proof

The formal proof of this theorem is quite involved and can be found in "The Source of Value."[3] I will skip the proof here, but I want to point out that it contains a very important condition, the "no-pumping" condition, that ensures that the growth rates are equal. Remember, I said that the long-term EPS growth rate can be found by taking the rate of EPS growth between rebalancings and multiplying it by the ratio of the P/Es just before and just after each rebalancing. The product of the P/E ratios has to be bounded. It is equivalent to being able to constantly buy low and sell high, which should not be possible in an efficient market. Thus, the no-pumping condition is necessary because without it, someone will always be able to construct examples in which the EPS growth rate of one index is higher than that of the market.

[3]Thomas K. Philips, "The Source of Value," *Journal of Portfolio Management* (Summer 2002):36–44.

Counterexamples. That said, the results of the theorem are much more general than the proof allows and certainly hold true in practice. It is, nevertheless, tempting to construct counterexamples to definitely disprove the theorem, but I found that whenever I constructed a counterexample, I had made three or four classic mistakes.

The first mistake is constructing a strategy that is not investable—one that requires cash inflows at rebalancing. This mistake is not unexpected because confusing EPS growth and earnings growth is very easy.

The second mistake is neglecting to adjust for the impact of the P/E change at each rebalancing, resulting in either too high or too low a rate of EPS growth.

The third and most subtle (and the easiest) mistake to make is taking the view of a single investor: "I can buy this hot stock, and then I can move into the next hot stock, and then I can move into the third hot stock, no problem." An investor can do that for a while but not in perpetuity. The model has to address the market in the aggregate, not the portfolio of an individual investor that represents a very small fraction of the market.

The fourth mistake is allowing the model to create portfolios that grow faster than the market and that allow continual reinvestment without any limits.

Observations. First of all, rebalancing is critical to the argument. Without the phenomenon of rebalancing, the theorem would not hold. A big gap exists between the value investing found in the *Journal of Finance* and the value investing done by investors, and this gap is caused entirely by this rebalancing effect. Rebalancing equalizes EPS growth by transferring securities between indexes; the fastest growing subset becomes the market and is eventually distributed evenly between the two halves of the market.

Second, I did not use any particular definition of value in the formal proof. I did not predicate my comments on P/E or price to book or any other measure. So, the proof must hold for any definition, any pair of indexes, just as long as the market is split into equal parts. Therefore, it must hold for small-cap versus large-cap as well as for growth versus value.

Implications

Several implications arise from this theorem. First, if the EPS values of the two indexes grow at the same rate, then they must generate the same rate of capital gains over the long run, assuming no P/E bubble.

Second, because total return equals capital gains plus dividend equivalents (i.e., dividends, share buy backs, takeovers, etc.), if the capital gains components of two indexes are equal, then the only difference in their returns must come from the difference in their free cash flow yields. Therefore, dividends and dividend equivalents matter, and they matter much more than most people think they do.

Third, the value premium is in some sense a fundamental premium, because any split between two indexes must induce a difference in return based only on a difference in the free cash flow yields; the value premium is entirely caused by excess free cash flow in the value index. Following is a test of this implication. **Figure 6** shows the price component of return of the S&P 500 value index versus the S&P 500 growth index from 1974 to 2001. The capital gains components of the two are essentially identical. Although value stocks dominated the market in the 1980s, growth stocks dominated in the late 1990s, and the Internet bubble dominated in the period from 1999 to 2000, the annualized price return over the entire period is essentially identical. **Table 3** shows that the differences between the price returns of the growth and value indexes for each of the periods are within a small fraction of a standard error of zero. Thus, it is not too far fetched to see that the true mean is zero and the data simply show some minor noise around the mean. I use three starting dates to ensure that I have not simply picked a particularly favorable starting point to validate my theory.

Risk vs. Return. Two schools of thought dominate the explanation for the value premium. The efficient markets school says that value stocks are riskier than growth stocks and this added risk explains their higher return. The behavioral, or the inefficient markets, school says that investors exhibit biases in their investment decision making—in particular, they consistently overestimate the growth potential of growth stocks—and that these biases explain the higher return of value stocks.

The evidence that I have presented from 1974 to 2001 shows that the earnings of value stocks are much more volatile than the earnings of growth stocks and that the returns of value stocks are much less volatile than the returns of growth stocks. So, if an investor thinks that risk is measured by the risk of earnings, then the value premium is consistent with the risk-based explanation. If an investor thinks that risk is measured by the risk of return, then the value premium seems more consistent with the behavioral explanation. One possibility for the lack of clarity is that the time series I used is not long enough to reach a conclusion. Another possibility is that a "peso problem" exists. That is, for a lengthy period of time, value stocks can be largely unaffected by the vagaries of the market and then a particular catalyst causes the value market to fall apart, which is exactly what happened to value stocks during the depression. The peso

Figure 6. Price Returns: S&P 500 Growth and S&P 500 Value, December 1974–December 2001

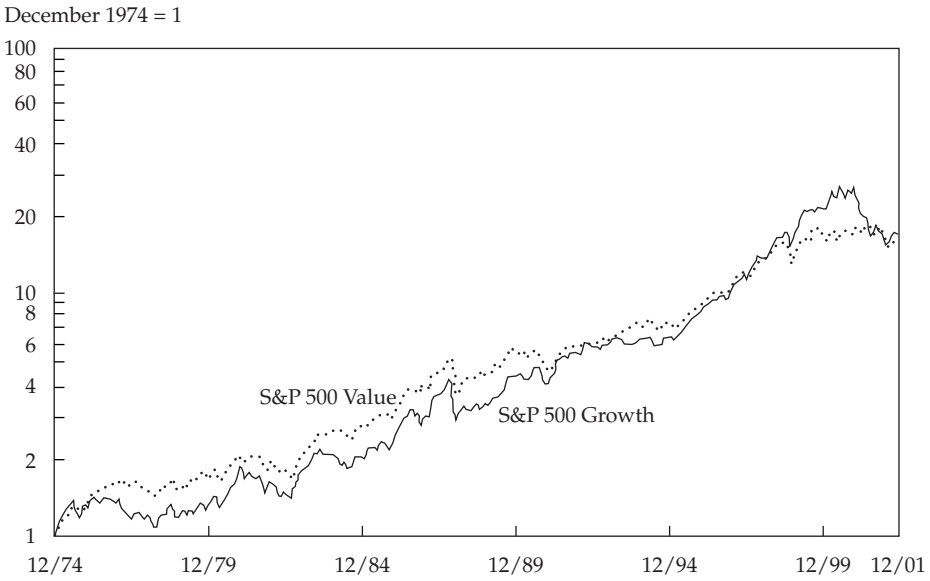

Table 3. Price Returns: S&P 500 Growth versus S&P 500 Value

Period/Index	Annualized Growth		Volatility	Standard Error
	Arithmetic	Geometric		
December 1974 to December 2001				
S&P growth	12.53%	11.06%	18.24%	
S&P value	11.54	10.76	13.27	
Difference	0.99	0.31	12.03	2.31%
December 1975 to December 2001				
S&P growth	11.93	10.89	18.33	
S&P value	10.64	10.34	12.67	
Difference	1.29	0.55	12.03	2.41%
December 1976 to December 2001				
S&P growth	11.99	10.91	18.70	
S&P value	9.96	9.66	12.43	
Difference	2.04	1.25	11.65	2.38%

problem could be the explanation for the unobserved risk of value stocks. Therefore, my results are consistent with either view, and the question of what drives the returns of value stocks will elicit spirited debate for a long time to come.

Size Premium. I am very skeptical of the size premium. If my argument is correct that over the long term all indexes that are rebalanced back to a fixed fraction of the market's capitalization must have exactly the same rate of EPS growth, then the only way small-cap stocks can have a higher relative return is if they have a higher free cash flow yield. But think about two companies, a big company and a small company, in the same line of business. They have comparable fixed costs, so the free cash flow yield as a fraction of revenue must be smaller (not bigger) for the small company than for the big company.

I believe that the small-cap premium is a classic case of data mining, and a lot of evidence supports this claim. Peter Knez and Mark Ready in their notable *Journal of Finance* paper trim 1 percent of the outliers in the Compustat database and find that the small-cap premium disappears and a large-cap premium appears.[4] And Jeremy Siegel has said that by simply taking a few years (1975–1983) out of the study period, the small-cap premium disappears and a large-cap premium appears.[5] So, I believe that size is not a priced factor. The size premium will be sporadic; investors will see it every now and then, but it is not something that can be reliably priced.

Expected Premiums. Assuming the expected return of the market is about 8 percent, then the difference between the expected return of value stocks and growth stocks is about 25 bps, and the small-cap premium is actually negative. I expect that small stocks will underperform large stocks, although I am not particularly confident of this claim for a very simple reason: The small cap indexes (Russell 2000, S&P 600, and so on) are not defined in terms of a fixed fraction of market cap; they are defined in terms of the bottom 600 stocks, the bottom 2,000

[4]Peter Knez and Mark Ready, "On the Robustness of Size and Book-to-Market in Cross-Sectional Regressions," *Journal of Finance* (September 1997):1355–82.

[5]Jeremy Siegel, *Stocks for the Long Run* (New York: McGraw Hill, 1998).

stocks, and so on. Thus, more slack and more give exist in the earnings growth rates of small-cap stocks, but over the long run, I believe the per share earnings growth rates of small-cap and large-cap stocks will be equal.

Summary

All rebalanced indexes must have exactly the same rate of earnings growth. As a consequence, the capital gains components of their returns must be identical. Any difference in their returns over the long run must be explained 100 percent by their free cash flow yields (i.e., dividends and dividend equivalents). Thus, dividends and dividend equivalents are much more important than most people believe.

Finally, the size premium is in all likelihood the result of data mining. Furthermore, the expected value premium is about one quarter of a percent, and the expected return of the market as a whole is about 8 percent.

Question and Answer Session

Thomas K. Philips

Question: What are the implications for your theorem if the typical addition to the S&P 500 is a growth stock?

Philips: Additions have a huge impact, but it is not the impact that you would expect. Assume S&P adds a growth-oriented stock to the S&P 500 growth index. After a quarter or so, S&P rebalances the index. So, for a month or two, these stocks boost the growth index's earnings a little bit, but the effect is not long lasting.

The effect that most people do not expect to see is that because S&P is always adding big stocks, growth stocks, and high-P/E stocks to the S&P 500, it tends to depress the EPS growth rate of the S&P 500 itself. Figure 1 shows the earnings of the S&P 500 growing faster than the earnings of the economy, which results from the large-company stocks being added to the index. But the per share earnings of the index grew just about in line with the economy because of the P/E effect. Jeremy Siegel estimates that this P/E effect has caused EPS to be depressed by about 1–2 percent a year for the past 40–50 years. So, there is an important effect from adding growth stocks to the S&P 500; it is just not the effect that most people expect to see.

Question: You mentioned that the earnings growth rate of large-cap companies should be about the same as that of small-cap companies, but small companies seem to have faster earnings growth than large companies. How can that be?

Philips: You don't see the small companies that stumble. If GE has a hiccup, it makes the newspapers. But when a small company goes out of business, you don't see it. It might make the local newspaper, but it doesn't make the front page of the *New York Times*. A lot of those small companies go out of business, which cuts the earnings growth rate of the small-cap indexes. So, the weighted average of the two is about the same. It is the mystery of the missing observations.

Question: Which value benchmarks do you run your money against?

Philips: I think we could use any index without a problem. Over the very long term, it is a complete nonissue. Every index provider has something that sounds proprietary, but at the end of the day, the indexes are all about the same. In the short term, however, you have to make sure that your process reflects the biases of the index that you'll be measured against.

Question: You showed earnings for value stocks dropping much more than earnings for growth stocks in 2001. That doesn't make sense given the drop in technology earnings.

Philips: Keep in mind that I was showing index level data. So, for every Cisco that saw an implosion of earnings there was an Armstrong Holdings or JDS Uniphase that also saw an implosion of earnings. Much of the drop in earnings isn't real; it's just accounting write-offs. It is not economic earnings disappearing; it is accounting earnings disappearing.

Question: How often are the Barra growth and value indexes rebalanced?

Philips: Once every six months, in June and December. The Russell indexes are rebalanced once a year, at the end of June.

Question: Rob Arnott has said that the theoretical expected return of stocks is zero. How did you get an expected return of 8 percent?

Philips: Rob is not arguing that the expected return of stocks is zero; he is arguing that the expected risk premium is zero. A 5.5 pp difference exists in those two points of view. Rob and I discussed this issue at a TIAA-CREF/AIMR-sponsored equity risk premium forum.[1] My estimate of the equity risk premium is a little higher than Rob's but not hugely different. I would say there is a 1–2 percent risk premium, maybe 1.5 percent; he says it is about zero, maybe 0.5 percent.

But Rob also says that free cash flow is a very important component of return, and I agree completely with that statement because capital gains can only grow in line with the economy. Any excess return above those capital gains has to come from dividends or dividend equivalents.

Question: Do the value and growth indexes in your analysis exhibit survival bias?

Philips: There is a peculiar form of bias in the data, but it is not biased in that sense. The caveat is that the data are all backfilled. S&P created its indexes in the 1990s and, with Barra's help, used the Compustat tapes to value all the stocks that existed at any point in time. So, S&P was able to determine what the S&P 500 looked like at any point in time. Then, using that live data, it split the S&P 500 into two halves based on price-to-book ratios or the particular measure the Compustat tapes allowed it to calculate. So, the indexes are affected by the backfilling effect, but it is not a bias. The numbers are real.

[1]To view the presentations made at the Equity Risk Premium Forum, go to the Additional Publications section at www.aimr**pubs**.org.