PORTFOLIO
MANAGEMENT
RESEARCH
By With Intelligence

the journal of PORTFOLIO management

volume 40
number 1

FALL 2013

jpm.pm-research.com

Clairvoyant Discount Rates

Robert D. Arnott, Feifei Li, and Geoffrey J. Warren

research affiliates™

Clairvoyant Discount Rates

Robert D. Arnott, Feifei Li, and Geoffrey J. Warren

Robert D. Arnott
is chairman of Research Affiliates in Newport Beach, CA.
arnott@rallc.com

Feifei Li
is director of Research Affiliates in Newport Beach, CA.
li@rallc.com

Geoffrey J. Warren
is a senior lecturer at the College of Business and Economics at the Australian National University, Canberra, ACT, Australia.
geoff.warren@anu.edu.au

The debate over efficient markets has raged for some 50 years. In one elegant step, the neoclassical finance theory community has recently finessed this controversy by embracing the idea that, just as investors require a different risk premium from stocks than from bonds, a different risk premium may be required of companies with different risk attributes. This move away from a static discount rate—or from the CAPM's beta-adjusted risk premium—allows us to explain all sorts of market anomalies without abandoning the efficient markets hypothesis (EMH).

There's a seductive elegance to the idea that market anomalies are a manifestation of cross-sectional differences in discount rates. If the risk premium is lower for safe-haven companies that are expected to deliver reliable growth than for companies that are expected to face grave challenges, then of course the former are priced at a premium and the latter at a discount. And if the companies that trade at a premium are expected to deliver a lower equity-risk premium, because they're seen as less risky, then naturally they will deliver a lower return, on average, creating a value effect in an efficient market.

Some might argue that there's little difference between an inefficient market and an efficient market with discount rates that vary from company to company. We choose merely to acknowledge this view, not to explore it. Rather, our investigation proceeds by examining, through the prism of a clairvoyant investor, the actual historical discount rates that would have justified share prices in the past half century. A clairvoyant investor is one who knows all future distributions from an investment, over a specific span of time, as well as its terminal price at the end of the time span.

No one can know the future, but we do know past cash flows and prices. So we can assess the value that a clairvoyant investor would have assigned to any company at any time, if they knew what we know now about subsequent distributions for a reasonably long span (10 years or more), as well as the terminal share price, using a preselected discount rate.

Bill Sharpe labeled this the clairvoyant value, a whimsical name that captures the concept with elegant simplicity. In 2009, Arnott, Li, and Sherrerd [2009a, 2009b, hereafter referred to as ALS 2009] explored the concept of clairvoyant value (CV) and its surprising implications. The findings of ALS 2009 were broadly consistent with the well-known size and value effect: There is a tendency for the share prices of value and small-cap stocks to be below CV, while the reverse tends to hold true for growth and large-cap stocks.

In ALS 2009, we learned that the market does a remarkable job—far better than most observers realize—of identifying which companies are likely to show the best growth in their fundamental metrics of business success,[1] but that the market historically tends to overpay for the winners and underpay for the losers, arguably creating the value effect. We also learned that the spread between the valuation of growth and value stocks mean-reverts over time, creating a mean-reversion-based growth/value "cycle."

ALS 2009 did not directly address the issue of the extent to which cross-sectional differences in discount rates across companies might provide a plausible explanation for the pricing anomalies observed.[2] Instead of choosing a discount rate and then deriving the CV, suppose we assume that the clairvoyant's net present value matches the price and derive the discount rate that aligns the two, reversing the mathematics of ALS 2009. In other words, we compute the subsequent discount rate that matches the current price with the net present value of future cash flows over a reasonably long time span. In so doing, we derive the discount rate that a clairvoyant investor must have intended to earn, in order to justify the then-prevailing price. We call this the clairvoyant discount rate (CDR).

Shifting the focus from cash flows to discount rates has a well-established history in financial research. Indeed, most asset-pricing research is conducted on stock returns. Just as CV helps us understand an array of nuances in market pricing, so does CDR. In fact, it brings forth some new surprises. We find that the most popular growth stocks require a CDR that seems implausibly low: Companies in the lowest book/price quintile often offer a CDR that's below the yield on bonds or even cash. This seems inconsistent with market efficiency. Could an efficient market really choose to accept a risk premium for safe havens and growth stocks that is negative relative to Treasury bills?

Further, we find that the size effect, which historically delivers superior returns to portfolios of small-cap stocks, is not evident in our CDR. The average CDR for individual small-cap stocks is, in fact, modestly lower than for large-cap stocks. The observed decline in average CDR as stock size decreases is the opposite of what we might expect, if discount rates were the primary explanation for any size-related effect.

The seeming paradox between the CDR and the size effects is easily resolved: CDR is more widely dispersed for small-cap than for large-cap stocks, leading to a handful of large-CDR outliers. A few stellar performers, therefore, raise the portfolio return of a small-cap portfolio, even though the average individual small-cap stock produces a lower internal rate of return (IRR) than the average large-cap stock. So return differences between large and small stocks are related to a dispersion effect, not the average discount rate.[3]

LITERATURE REVIEW

John Burr Williams [1938] was perhaps the first to suggest the net present value of all future cash flows as the basis for gauging intrinsic value. Grossman and Shiller [1981] proposed a "perfect foresight price" concept for an investor whose exposure is the entire Standard and Poor's Composite Index. Shiller [1981, 1987] used this concept to compare the ex post rational price of the index and the actual price, showing that excess volatility in the aggregate market price cannot be explained by the change in dividends, real interest rates, and the intertemporal marginal rate of substitution.

ALS 2009 explored this "crystal ball" concept within the stock market: What is an individual stock ultimately worth for an investor who purchases and owns it over decades? Ex post realized cash flows are used to measure the investment's CV with accuracy and objectivity. Over the half-century span in the ALS 2009 studies, the market demonstrated a very impressive and overwhelmingly significant capacity to discern stocks with faster ex post growth rates. Indeed, valuation multiples had a lofty average correlation of roughly 50% with subsequent observed growth rates. However, the market pays far more for superior growth expectations than the clairvoyant investor should be willing to pay, again with overwhelming statistical significance.

Related to our work, Fama and French [1992, 1993] showed that the variation in cross-sectional expected returns can be largely captured by size and book-to-market effects. Cochrane [2011] provided an extensive study on how risk premiums vary over time and across assets and how this variation affects the movement of valuation ratios. Our study assumes that discount rates are fixed for 10- and 20-year horizons, as well as longer spans, from various start dates coterminous in 2011, of up to 56 years.[4] For each of these spans, and for each stock in the study, we derive the discount rate that aligns the net present value with the starting price.

In so doing, we do not account for variations in discount rates that are time dependent. In this sense, our findings are consistent with the existing literature on the value effect. This means that our results do not necessarily contradict market efficiency. However, our results are consistent with market efficiency only if the time and cross-sectional variation in discount rates is sufficiently large to allow our results.

In addition, our work is partly inspired by the implied cost of capital literature. This literature typically generates estimates of the expected return on equity through solving for the internal rate of return that equates current share prices with expectations for future cash flows, with the latter often based on analyst forecasts (e.g., Claus and Thomas [2001]; Botoson and Plumlee [2005]; Guay et al. [2011]). This method has been used to investigate the relation between expected returns and stock characteristics (e.g., Gebhardt et al. [2001]), as well as to test asset-pricing models (e.g., Pástor et al. [2008]; Lee et al. [2009]).

As with the implied cost of capital literature, we examine CDR to better understand return generation. However, while the implied cost of capital literature measures expected returns based on a proxy for cash-flow expectations, CDR estimates the actual return that must necessarily have been expected by any clairvoyant investor with perfect foresight of future cash flows under a buy-and-hold strategy.

An issue with the implied cost of capital research is that analyst forecasts are imperfect measures of market expectations, in part because they are known to be biased (Easton and Sommers [2007]) and in part because prices may reflect other information. CDR contains no such measurement errors. Nevertheless, it must be interpreted with some care, as perfect foresight is a rather demanding assumption.

OUR METHODOLOGY

How does CDR differ from the realized total rate of return over the measurement period? In essence, we view the CDR as an internal rate of return (IRR) with broader cash-flow reinvestment assumptions, notably including external reinvestment alternatives. An IRR is defined as the rate that equates the sum of the net present value of all cash flows (typically including the initial cash outlay for the purchase price as a negative cash flow) to zero. An IRR is essentially a money-weighted return, because cash contributions to the portfolio influence the return. Total return, on the other hand, is a time-weighted return, in that the timing of cash contributions to the portfolio is irrelevant; only the time period over which the return is calculated matters. If the stock never pays any dividends, the two concepts are identical. If the stock pays dividends, we can secure the cash portion by spending it or reinvesting into either a risk-free asset or the market index.

This seemingly trivial nuance matters. Consider General Motors (GM) for an investor in 1955, the first year of our study. A buy-and-hold investor who reinvests all dividends in GM stock earns the cumulative total return for GM, assuming it is held until its delisting in the wake of its 2009 bankruptcy at pennies a share. Our 1955 investor achieves a negative total return of –1.08% annualized, losing nearly 50% cumulatively, even after reinvesting decades of rich dividends.[5]

If the dividends are invested in T-bills, the stock investment itself was still a big loss. Nearly 100% of the initial investment is gone when the stock is delisted from NYSE at a price of $0.61. But our clairvoyant investor got the steady income from the safe-haven GM investment, until it finally failed. So the CDR for the clairvoyant investor was modestly positive, compounding at an annualized 5.82%. For the investor who spends the distributions, the 56-year return was even larger at 7.54%, because the company paid out more than $200 in dividends for every $100 invested in 1955. If our clairvoyant investor reinvests the dividends into the S&P 500 Index, the return improves to a very respectable 9.15%. (The differences in these calculations can be seen in the subsequent formulas.)

The average annualized geometric total return (TR) of a stock S (as it is traditionally studied and reported) is the geometric average of the single-month returns our database reports. There is no underlying assumption about its cash flow distribution and reinvestment.

$$TR = \sqrt[T]{(1 + R_{s,1})(1 + R_{s,2}) * \ldots (1 + R_{s,T})}$$

Here $R_{s,t}$ stands for the stock's total return for a single period t.

We use two ways to calculate CDRs. In method 1, we directly discount back the dividend payouts to t_0 and retrieve the IRR:

$$p_0 = \frac{CF_1}{(1+IRR)^1} + \frac{CF_2}{(1+IRR)^2} + \cdots + \frac{CF_{T-1}}{(1+IRR)^{T-1}} + \frac{CF_T + P_T}{(1+IRR)^T}$$

Method 1 differs from the conventional IRR calculation only in assuming that the proceeds of delisted stocks are reinvested in the S&P 500 or in T-bills for the remainder of the measurement period.

In method 2, we assume that dividends are reinvested elsewhere:

$$P_0 = \frac{CF_1(1+R_{a,2})(1+R_{a,3}) * \cdots (1+R_{a,T})}{(1+IRR)^T} + \frac{CF_2(1+R_{a,3})(1+R_{a,4}) * \cdots (1+R_{a,T})}{(1+IRR)^T} + \cdots + \frac{CF_{T-1}(1+R_{a,T})}{(1+IRR)^T} + \frac{CF_T + P_T}{(1+IRR)^T}$$

Here $R_{a,t}$ stands for the assumed return on reinvestment for period t. In method 2, we reinvest dividends into the S&P 500 or T-bills. If the stock is delisted, we reinvest the delisting proceeds into either the S&P 500 or T-bills.

Our base case is that investors spend all the dividends they receive.[6] After all, most stocks will eventually disappear due to bankruptcy or mergers and acquisitions, as Schumpeter's "creative destruction" changes our economic landscape. A clairvoyant investor won't worry about an end price of zero, many decades hence, as long as the IRR reflects lofty distributions to investors along the way. Thus, in this article, method 1 reflects the assumption that the clairvoyant investor will spend all the dividends.

Starting with all the publicly traded companies listed in the CRSP tapes, the sample period spans 672 months, from December 1955 to December 2011. At each year-end from 1955 to 2001, we track all the companies' cash flows and terminal price information for the next 10 years, 20 years, and until the end of 2011. When the company does not survive for the full CDR span, we compute CDRs using all available information.

Finally, our analysis ends with a starting date of December 31, 2001, so that we're always relying on at least a 10-year span, unless a bankruptcy or corporate action intervenes. We don't want very short CDRs, which will be very volatile and may be particularly sensitive to the extreme events of the global financial crisis,[7] clouding our results.

We retrieve financial data from the Compustat database. At the end of each year over our study span, all the stocks in the sample are ranked and partitioned into five quintiles by market capitalization. The independent sorting and ranking are also done for the book-to-market ratio, which is the stock's book value from the previous fiscal year divided by its market capitalization at the end of each calendar year.[8] We exclude stocks with missing or negative book values, earnings, or cash flows in the year immediately preceding a clairvoyance span. In order to minimize any risk of survivorship bias, we include all companies for which financial and price data are available from CRSP and Compustat.

Finally, we classify each stock in our sample into one of the 25 portfolios sorted on quintiles of market capitalization and book-to-market ratio. Unlike most anomaly studies, we focus on the average CDR for individual stocks, instead of portfolio returns. In so doing, we can avoid rebalancing issues and more accurately measure the break-even rates of return investors receive for holding individual stocks of a particular type over a long horizon. The tables showing portfolio results in this section are simply the equally weighted average or market-capitalization-weighted average of the CDRs for the individual stocks that belong to the particular size/style partition.

DATA AND EMPIRICAL RESULTS

Exhibit 1 illustrates the difference in discount rates measured by the two approaches, as well as the total-return concept for the top 10 companies in the U.S. stock market at the end of 1955. GM was the company with the largest market capitalization as of December 31, 1955, and stayed on top of the list for two more years. The annualized total return for GM over the next 56 years is quite depressing, at –1.08%. However, due to its large dividend payouts, its return is much higher: 7.54%, when we assume that the dividends are spent and the eventual delisting proceeds invested into the S&P 500.

These data show that a 10-year clairvoyant (able to see the next 10 years of distributions and the terminal price at the end of 1965) must have required a CDR of 12.71% in order to be willing to pay GM's year-end 1955 price of $46.25 per share. If this same investor

EXHIBIT 1

Illustration of CDR Method 1 vs. Method 2 vs. Total Return Using Top 10 Largest Companies as of 12/31/55

All the companies in CRSP tape as of 12/31/1955 are selected and their market capitalization is computed as the product of total shares outstanding and the last trading day's quoted price in 1955. The top 10 largest companies are listed here. The 1-month T-bill rates are taken from Kenneth French's data library (Ibbotson Associates). S&P 500 returns are retrieved from Global Financial Data.

Rank	Company Name	Delisting Proceeds Reinvested into S&P 500			Delisting Proceeds Reinvested into 1 Month t-bill		
		CDR-10yr	CDR-20yr	CDR-2011	CDR-10yr	CDR-20yr	CDR-2011
Method 1							
1	General Motors Corp	12.71%	7.62%	7.54%	12.71%	7.62%	7.54%
2	Du Pont E I De Nemours & Co	7.52%	2.82%	6.02%	7.52%	2.82%	6.02%
3	Standard Oil Co N J	8.73%	7.56%	10.64%	8.73%	7.56%	10.64%
4	American Telephone & Teleg Co	12.70%	8.94%	9.21%	12.70%	8.94%	9.20%
5	General Electric Co	10.16%	5.75%	9.46%	10.16%	5.75%	9.46%
6	Texas Co	14.99%	8.88%	9.93%	14.99%	8.88%	9.85%
7	Union Carbide & Carbon Corp	5.25%	3.89%	6.82%	5.25%	3.89%	6.81%
8	United States Steel Corp	3.85%	4.76%	4.88%	3.85%	4.76%	4.88%
9	Standard Oil Co California	10.86%	7.34%	10.54%	10.86%	7.34%	10.54%
10	Sears Roebuck & Co	16.48%	9.94%	9.02%	16.48%	9.94%	8.97%
		Dividends Reinvested into S&P 500			**Dividends Reinvested into 1 Month t-bill**		
		CDR-10yr	CDR-20yr	CDR-2011	CDR-10yr	CDR-20yr	CDR-2011
Method 2							
1	General Motors Corp	12.68%	6.80%	9.15%	11.43%	6.72%	5.82%
2	Du Pont E I De Nemours & Co	8.16%	3.70%	8.05%	7.06%	3.45%	5.51%
3	Standard Oil Co N J	9.39%	6.99%	10.12%	7.90%	6.83%	8.50%
4	American Telephone & Teleg Co	12.66%	7.83%	9.60%	11.22%	7.66%	6.54%
5	General Electric Co	10.47%	5.77%	9.40%	9.41%	5.57%	7.88%
6	Texas Co	14.66%	7.75%	9.82%	13.66%	7.71%	7.08%
7	Union Carbide & Carbon Corp	6.37%	4.44%	8.32%	4.96%	4.18%	6.09%
8	United States Steel Corp	5.98%	5.26%	8.26%	3.67%	4.80%	5.02%
9	Standard Oil Co California	11.11%	6.84%	10.08%	9.86%	6.68%	8.39%
10	Sears Roebuck & Co	16.09%	9.00%	9.43%	15.41%	8.96%	7.12%
		Total Return					
Total Return		10yr	20yr	2011			
1	General Motors Corp	13.80%	6.76%	−1.08%			
2	Du Pont E I De Nemours & Co	8.12%	2.58%	7.12%			
3	Standard Oil Co N J	9.22%	7.92%	12.43%			
4	American Telephone & Teleg Co	12.30%	8.04%	7.69%			
5	General Electric Co	10.40%	5.30%	9.57%			
6	Texas Co	15.18%	7.18%	11.20%			
7	Union Carbide & Carbon Corp	5.39%	4.43%	10.07%			
8	United States Steel Corp	3.05%	5.55%	6.42%			
9	Standard Oil Co California	11.31%	7.12%	11.98%			
10	Sears Roebuck & Co	17.03%	9.21%	9.61%			

**If the stock got delisted prior to the sample period end, the annualized return from 1955/12/31 to the delisting date is taken.*

Source: Research Affiliates, based on data from Compustat and CRSP.

could see 20 years into the future, that same price would imply a CDR of 7.62%. If the investor could see through to year-end 2011, after the price had gone near zero (the delisting price was $0.61, which we presume was reinvested into the S&P 500), the year-end 1955 price would have been fair, if our clairvoyant investor were happy with a 7.54% return. (This is a mirror image of the analysis in the ALS 2009 papers.)

By contrast, a buy-and-hold investor in 1955, choosing to buy Standard Oil Company of New Jersey (Exxon Mobil today), benefits from its overall good performance over the next 56 years, earning a compound annual gain of 12.43%. If the investor spends dividends from the oil giant instead of reinvesting in the stock, the earned return is a bit lower, at 10.64%. If the dividends are reinvested into the S&P 500 or T-bills, the annualized returns are a bit lower, at 10.12% and 8.50%, respectively. Investors in most of the top 10 companies on December 31, 1955, enjoy a better performance if all the interim cash flows are invested in the market index rather than the default-free instrument; this is a natural consequence of the S&P 500's higher return relative to T-bills over the next 10-year span, 20-year span, and the full span through 2011.

For an investor with 10 years of clairvoyance, the stock with the highest CDR was Sears at 16.48%; the stock with the lowest CDR was U.S. Steel, at 3.85%. For a wise seer with 56-year clairvoyance through year-end 2011, Standard Oil Company of New Jersey was priced to deliver an IRR of 10.64% over that very long period, while U.S. Steel was priced to deliver an IRR of 4.88%. Presumably Standard Oil Company of New Jersey drastically exceeded expectations over this half-century span, especially when compared with U.S. Steel. But because we're assuming our clairvoyant investor could see this future, we are assuming that he or she was genuinely happy to earn 6% less on U.S. Steel than on Standard Oil Company of New Jersey. For this to make sense, U.S. Steel must have been seen as a low-risk investment that amply deserved its anemic CDR.

A critic might say that this is silly: No clairvoyant will be willing to accept 6% less on one large-cap, blue-chip investment than another. Alternatively, our critic may point out that this outcome just reflects that our CDR calculations impose the unreasonable assumption of perfect foresight, and U.S. Steel might have generated low returns because it disappointed versus what were entirely reasonable expectations at the period's start. Fair enough.

But both critiques miss our key point. If we partition stocks based on size and style attributes that are self-evident up front, such as starting market cap and valuation multiples, and if the average CDR for these partitions differs widely, then investors have presumably been pricing these stock categories to reflect widely—even wildly—different demands for an equity-risk premium associated with those starting characteristics. Further, averaging across many stocks should cancel any idiosyncratic forecasting errors, leaving just the returns associated with those characteristics.

We now move on to observe the range of returns across portfolios formed along the size and value dimension. We find the average presumed ex ante risk premium of large-cap growth stocks (again for the clairvoyant investor) is remarkably anemic and often negative, relative to the prospective return for government bonds or even T-bills. We also find that small-cap stocks do not seem priced to generate greater returns on average, relative to large-cap stocks.

FROM SNAPSHOT TO A HALF-CENTURY HISTORY

We now shift attention from a snapshot of the top 10 by market cap at a particular time (1955), and consider all companies for all start dates from year-end 1955 until year-end 2001 by size and valuation quintiles.[9]

Panel A in Exhibit 2 shows that the discount rate required to justify current prices for the 10-year clairvoyant has averaged 3.98% for the average stock in the growth quintile and 10.75% for the average stock in the value quintile. (By assumption, 10-year clairvoyants can foresee 10 years of future distributions, plus the end price after 10 years.) This result assumes that dividends are withdrawn and delisting proceeds are reinvested in the S&P 500. For 20-year clairvoyants, the growth stocks required a 6.75% discount rate; value stocks are priced to deliver an average 10.98% return. For clairvoyant investors who could see all distributions for all U.S. companies through 2011, as well as the closing price at year-end 2011, growth stocks are fairly priced only if the clairvoyant uses an average discount rate of 5.01%. Value stocks required an average discount rate of 8.37%.

An obvious surprise is that most of these CDRs are lower than market returns over like spans. One of the main reasons is simple returns dispersion. Suppose that growth stocks are priced at a CDR averaging just 5%, but that some are priced for 0% (IRR is zero) and others are priced for 10%. Over a 20-year span, this two-stock portfolio does not deliver 5%. The stock with zero IRR remains fixed at $1, while the stock with a 10% IRR (assuming no dividends or corporate actions) grows from $1 to $6.73. Our two-stock portfolio grows four-fold, from $2 to $7.73, which is worth 7% per annum, not 5%. The wider the dispersion, the bigger the gap between the average company's CDR and the size or style portfolio returns. The dispersion effect can largely account

Exhibit 2

Average Clairvoyant Discount Rate, over 10-Year and 20-Year Spans and through 2011

Assumptions: No reinvestment of dividends, delisting proceeds reinvest into S&P 500 (Method 1).

	Average Return			Average Excess Return		
	10yr T-bond	1m T-bill	S&P 500	10yr T-bond	1m T-bill	S&P 500
10-yr	7.19%	5.53%	10.19%	--	–1.66%	3.00%
20-yr	7.76%	6.02%	10.96%	--	–1.75%	3.19%
Through 2011	6.30%	4.39%	8.70%	--	–1.90%	2.40%

Panel A: Equal-Weighted Averages

10-yr CDR

		Growth	2	3	4	Value
	Averages	***3.98%***	***6.82%***	***8.56%***	***10.06%***	***10.75%***
Small	***6.53%***	1.91%	4.73%	6.97%	8.78%	10.27%
2	***7.17%***	2.11%	5.81%	8.14%	9.42%	10.37%
3	***8.27%***	3.55%	7.01%	9.02%	10.62%	11.14%
4	***9.06%***	5.49%	8.24%	9.49%	10.92%	11.17%
Big	***9.14%***	6.83%	8.32%	9.16%	10.57%	10.81%

20-yr CDR

		Growth	2	3	4	Value
	Averages	***6.75%***	***8.59%***	***9.58%***	***10.59%***	***10.98%***
Small	***8.03%***	5.13%	7.16%	8.25%	9.54%	10.06%
2	***8.75%***	5.64%	8.04%	9.19%	10.23%	10.63%
3	***9.65%***	6.80%	8.83%	10.08%	10.99%	11.57%
4	***10.18%***	7.78%	9.59%	10.31%	11.39%	11.82%
Big	***9.88%***	8.41%	9.32%	10.05%	10.81%	10.80%

Through 2011 CDR

		Growth	2	3	4	Value
	Averages	***5.01%***	***6.57%***	***7.57%***	***8.29%***	***8.37%***
Small	***5.95%***	3.25%	4.96%	6.33%	7.44%	7.77%
2	***6.57%***	3.72%	5.89%	7.17%	7.88%	8.16%
3	***7.28%***	4.85%	6.64%	7.84%	8.49%	8.59%
4	***7.90%***	6.07%	7.50%	8.29%	8.84%	8.79%
Big	***8.11%***	7.15%	7.85%	8.24%	8.78%	8.54%

Panel B: Value-Weighted Averages

10-yr CDR

		Growth	2	3	4	Value
	Averages	***4.05%***	***6.91%***	***8.66%***	***10.16%***	***10.73%***
Small	***6.65%***	1.84%	4.89%	7.11%	9.04%	10.39%
2	***7.21%***	1.99%	5.85%	8.24%	9.61%	10.35%
3	***8.39%***	3.79%	7.04%	9.07%	10.73%	11.31%
4	***9.01%***	5.47%	8.24%	9.46%	10.87%	11.00%
Big	***9.25%***	7.16%	8.54%	9.44%	10.55%	10.58%

20-yr CDR

		Growth	2	3	4	Value
	Averages	***6.78%***	***8.69%***	***9.67%***	***10.67%***	***11.06%***
Small	***8.12%***	5.02%	7.32%	8.34%	9.69%	10.22%
2	***8.79%***	5.58%	8.08%	9.28%	10.37%	10.64%
3	***9.74%***	6.90%	8.82%	10.10%	11.09%	11.79%
4	***10.12%***	7.76%	9.61%	10.21%	11.32%	11.72%
Big	***10.10%***	8.64%	9.63%	10.41%	10.88%	10.92%

Through 2011 CDR

		Growth	2	3	4	Value
	Averages	***5.09%***	***6.61%***	***7.66%***	***8.36%***	***8.38%***
Small	***6.06%***	3.24%	4.98%	6.54%	7.62%	7.91%
2	***6.58%***	3.66%	5.92%	7.20%	7.99%	8.14%
3	***7.37%***	4.99%	6.67%	7.88%	8.56%	8.74%
4	***7.90%***	6.11%	7.53%	8.25%	8.86%	8.74%
Big	***8.20%***	7.47%	7.93%	8.44%	8.79%	8.39%

Source: Research Affiliates, based on data from Compustat and CRSP.

for the seemingly low CDRs we report here, relative to the traditional portfolio returns of the corresponding size/style portfolios.

It's striking to note that the growth stocks are not priced, on average, to deliver any risk premium whatsoever relative to 10-year Treasury bonds and are often priced to deliver less return than T-bills. For reference, the average yield over the average 10-year span for U.S. T-bills was 5.53%, while 10-year Treasury bonds carried an average starting yield[10] of 7.19% and a near-identical average 10-year return. (How quickly we forget these lovely yields of yesteryear!)

The average growth stock's CDR fails to match T-bills, let alone T-bonds, at only 3.98% over 10 years, and 5.01% through 2011. The small-growth corner portfolio is an even bigger surprise. The average stock in the small-growth portfolio is priced to deliver a 10-year return of 1.91%, a 20-year return of 5.13%, and a return through 2011 of just 3.25%.

It is well known that small-growth stocks have historically delivered a slender risk premium. It is less widely recognized that this modest portfolio-risk premium is delivered courtesy of a handful of extreme winners. If we owned a selection of small-growth companies and missed the extreme winners, we would have experienced a negative risk premium, because the average stock[11] in this sector delivered a negative risk premium.

Exhibit 3 graphically displays the average implied CDR through 2011 for each of the five quintiles, from growth to value, calculated year by year from 1955 until 2001.[12] This is plotted with the subsequent average 10-year T-bond and T-bill yields.[13] Exhibit 4 shows the average through-2011 CDR for the five quintiles, net of the subsequent average returns for the 10-year bonds. This vividly illustrates one of our key findings: Growth stocks appear to be priced, on average, to deliver a negative risk premium relative to ordinary T-bonds.

Clearly, past investors were sufficiently rattled by the troubled prospects for value stocks that they demanded a large premium IRR—a large discount rate—on average, rather than contentedly relying on the wide IRR dispersion compounding to deliver respectable portfolio returns. At the other end of the spectrum, can it be that the average small-cap growth stock is priced to deliver 1% less return than T-bills, and 3% less than ordinary T-bonds? So it would appear! These investors are satisfied to rely on extreme outliers to salvage their portfolio risk premium.

Panel B of Exhibit 2 suggests that the use of value-weighted portfolios, instead of equally weighted portfolios, for the size and value quintiles, makes surprisingly little difference. We get a few basis points, more or less—rarely as much as 20 bps—with value weighting. This may seem at odds with the ample literature suggesting that equally weighted portfolios generally outperform value-weighted portfolios. But recall the dispersion effect: The total weight of outliers in the equally weighted list is more than in the value-weighted list. Our equally weighted investor can accept a lower average IRR for portfolio stocks and still garner a higher average IRR for the portfolio in total.[14]

Because it makes comparatively little difference whether we look at 10 years, 20 years, or full-span results, and also little difference whether we use equal weight or value weight, the next table focuses only on the equally weighted results through 2011. Exhibit 5 shows the differences in the average CDR if dividends are reinvested into the S&P 500 or into T-bills (as described in method 2), not simply withdrawn from the portfolio. Reinvesting the dividends and delisting proceeds into the S&P 500 produces results averaging about 243 basis points per year better than reinvesting in the T-bills. This is not surprising, given that the S&P 500 earns a total return of 8.70% annually, or about 4.31% above the average T-bill rate, in our study span.

In this last case, it's striking to note that the average growth-stock investment, bought in each year over the past 56 years and held through 2011, delivers a CDR of just 3.19% with dividends and delisting proceeds reinvested into T-bills. In other words, a clairvoyant investor would have found the average growth stock fairly valued using an average discount rate of 3.19%. This is markedly less than placing those same investment dollars into T-bills, which delivered 4.39% per year through 2011 for the average of this half-century's start date.

Portfolio effects can save us from our negative risk premium, because upside departures from our average CDR of 3.19% deliver more benefit to us than symmetric downside outliers can hurt us. Nevertheless, these data suggest that it is entirely fair to expect a negative equity-risk premium on a typical, single growth-stock investment.

EXHIBIT 3

CDR vs. Returns on Cash and Bonds, Average Subsequent Annualized Returns through December 2011

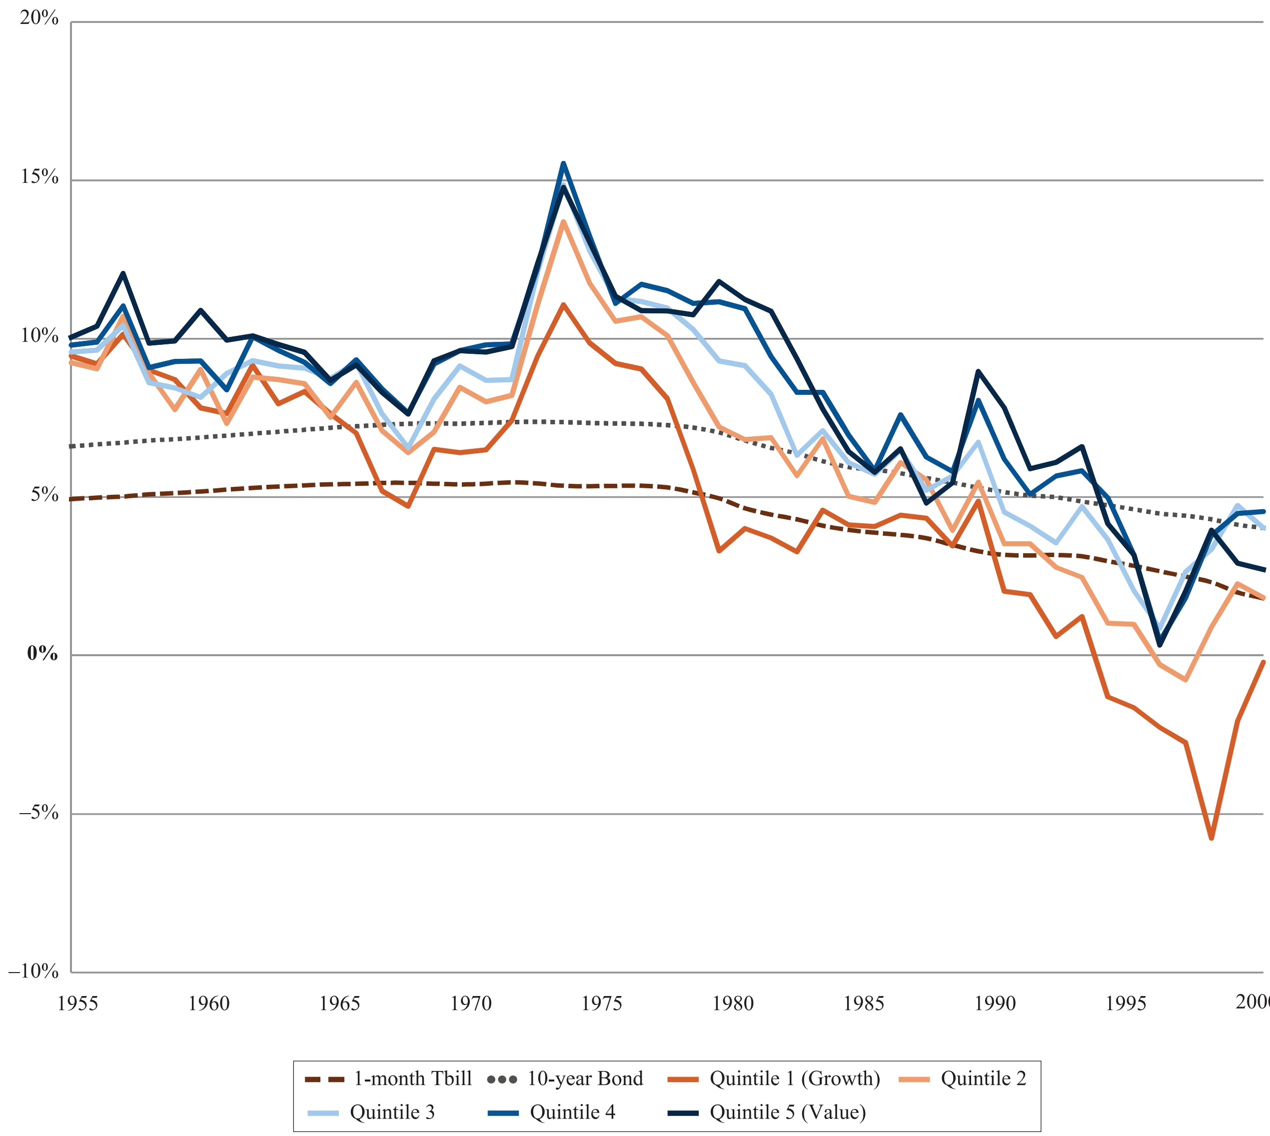

Source: Research Affiliates, based on data from Compustat and CRSP.

WHAT HAPPENED TO THE SIZE EFFECT?

A second surprise in our results is that the size effect seems to reverse for an average single-stock holding. Consider panel A of Exhibit 2, showing the 10-year clairvoyance tests. The quintile with the smallest market-cap companies has an average CDR of 6.53%, while the biggest market-cap quintile delivers an average CDR of 9.14%. It doesn't matter whether we equal weight or value weight our analysis. It doesn't matter if we use method 1 (spend the dividends) or method 2 (reinvest in T-bills or in the S&P 500). It would appear that the average small-cap stock has been priced to deliver somewhere between 2% and 3% less in future return than the average large-cap stock, over a 10-year window.

Here we can offer a simple explanation with supporting evidence: the dispersion effect. It should come as

Exhibit 4

Excess CDR over 10-Year Bond, Average Subsequent Annualized Returns through December 2011

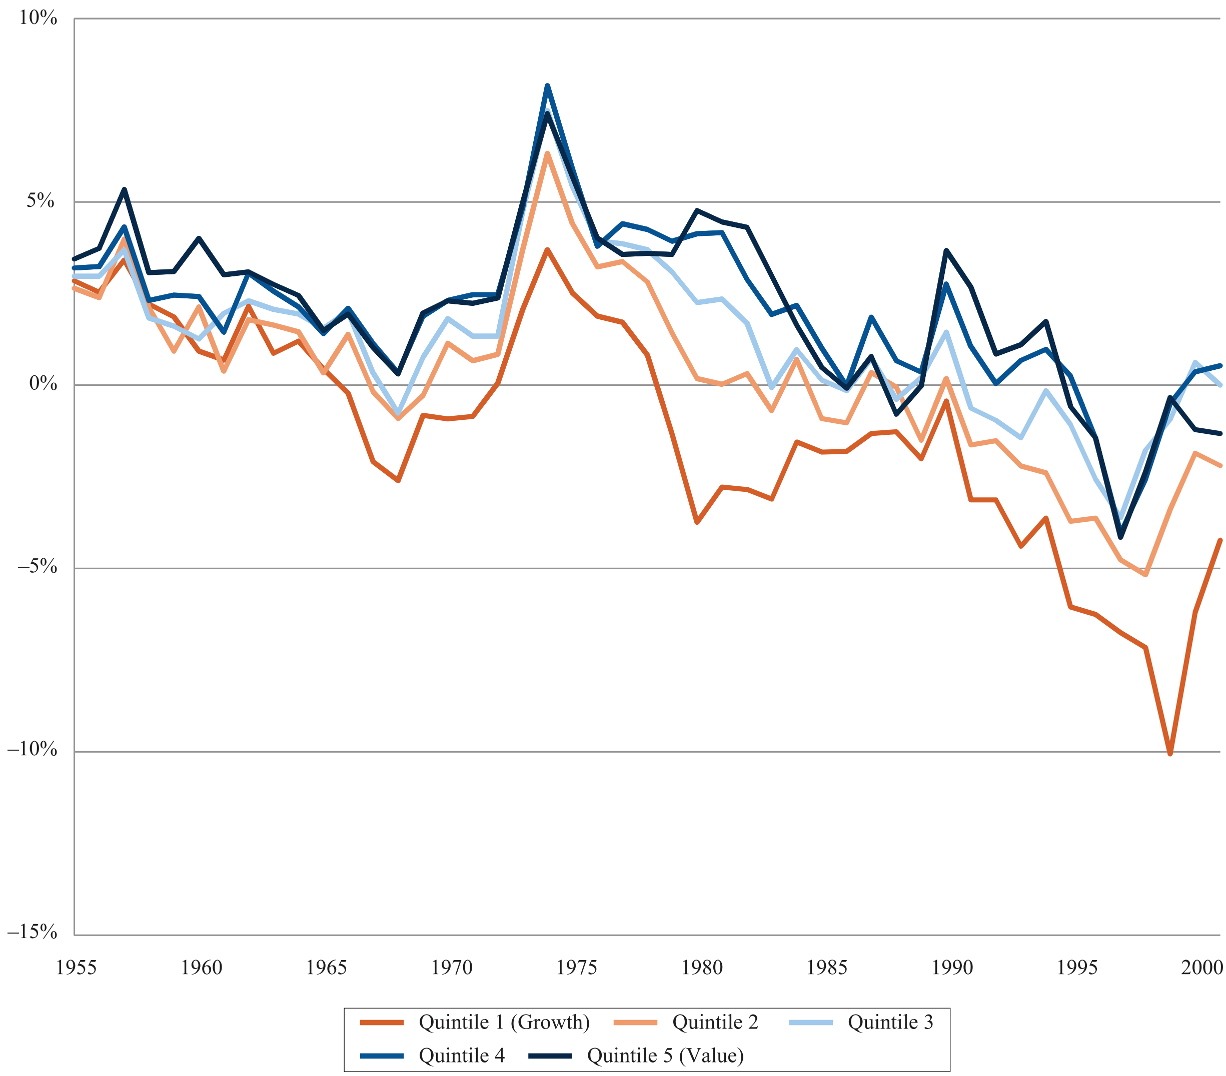

Source: Research Affiliates, based on data from Compustat and CRSP.

no surprise that the dispersion of CDR for small companies is far wider than for large companies, if only because the small companies often have a less diversified business, with far greater company-specific risk. To test this intuition, we compute the average standard deviations of CDRs across size and style (value/growth) portfolios. In Exhibit 6, the CDRs' standard deviations monotonically decrease across all the size groups, as we go from small to large companies, regardless of the growth/value quintile. The dispersion seems to shrink when we move into longer horizons, from 10-year to 20-year to through-2011. The average difference in CDRs' standard deviations in the smallest quintile, relative to those of the biggest quintile, is 5.61% for the 10-year span (13.61% for the small-cap quintile versus 8.00% for the large-cap quintile, a difference of 5.61%), 3.79% for the 20-year span, and 2.24% for the entire 56-year span.

We don't observe such a large difference in dispersion for the growth quintile, relative to value. Here we see something of a saddle, with wider outcome dispersion observed in the growth quintile and the value quintile than in the middle three quintiles. Again, this

EXHIBIT 5

Equal-Weighted Average Clairvoyant Discounts Rates Through 2011

Assumptions: Dividends and Delisting Proceeds Reinvest into S&P500		Growth	2	3	4	Value
	Averages	*5.40%*	*6.93%*	*7.80%*	*8.33%*	*8.43%*
Small	*6.08%*	3.48%	5.18%	6.46%	7.48%	7.82%
2	*6.80%*	4.08%	6.24%	7.44%	7.96%	8.27%
3	*7.53%*	5.36%	7.06%	8.02%	8.57%	8.62%
4	*8.13%*	6.53%	7.89%	8.53%	8.85%	8.86%
Big	*8.35%*	7.57%	8.27%	8.54%	8.79%	8.57%

Assumptions: Dividends and Delisting Proceeds Reinvest into T-bills		Growth	2	3	4	Value
	Averages	*3.19%*	*4.57%*	*5.36%*	*5.81%*	*5.81%*
Small	*3.51%*	1.08%	2.67%	3.90%	4.83%	5.08%
2	*4.28%*	1.72%	3.81%	4.92%	5.38%	5.59%
3	*5.06%*	3.02%	4.64%	5.59%	6.03%	6.01%
4	*5.71%*	4.37%	5.56%	6.10%	6.32%	6.22%
Big	*6.17%*	5.76%	6.16%	6.26%	6.50%	6.16%

Source: Research Affiliates, based on data from Compustat and CRSP.

is unsurprising. The out-of-favor value stocks either shake off their difficulties or they don't, leading to wide CDR dispersion. Meanwhile, the high-flying growth stocks either deliver on the market's lofty expectations or they don't, also leading to wide dispersion. The sensible, steady middle companies have notably less CDR dispersion, an outcome that any seasoned practitioner would expect.

The results in Exhibit 6 are based on the CDRs from method 1—dividends are spent, with no reinvestment, and the delisting proceeds are reinvested in the S&P 500—and are therefore best compared with Exhibit 2, which puts dividends and delistings through the same treatment.[15] The large CDR dispersion present within the small-size portfolios supports the explanation for the lack of a small-stock premium in our long-horizon buy-and-hold CDR study.

We illustrate this point in Exhibit 7, which constructs an example calibrated using data from Exhibits 2 and 6 for 10-year clairvoyance, with the companies in each bucket equally weighted. Consider the small-growth and large-value categories. The average CDR for small growth is 1.91% with a standard deviation of 14.69%; for large value, the average is 10.81% with a 7.88% standard deviation.

Let's assume that the small-growth bucket has two stocks: one brilliant, delivering a 16.60% CDR, and the other awful, delivering –12.78% CDR (this is 1.91% ± 14.69%, matching our empirical mean and standard deviation). The large-value bucket has two stocks: one delivering 18.69% and the other sagging with 2.92% CDR (equivalent to 10.81% ± 7.88%).[16]

EXHIBIT 6

Average Standard Deviation of Clairvoyant Discount Rates over 10-Year and 20-Year Spans and Through 2011

Assumptions: No reinvestment of dividends, delisting proceeds reinvest into S&P 500 (Method 1).

		Growth	2	3	4	Value
Standard Deviation of 10-year CDRs						
	Averages	*13.38%*	*11.50%*	*10.84%*	*10.41%*	*11.41%*
Small	*13.61%*	14.69%	14.11%	13.53%	12.43%	13.30%
2	*13.15%*	15.87%	12.88%	12.40%	12.15%	12.45%
3	*12.20%*	14.49%	12.00%	11.37%	11.04%	12.10%
4	*10.58%*	12.48%	10.51%	9.45%	9.18%	11.30%
Big	*8.00%*	9.36%	8.01%	7.47%	7.27%	7.88%
Standard Deviation of 20-year CDRs						
	Averages	*8.61%*	*7.41%*	*7.02%*	*6.97%*	*7.63%*
Small	*9.04%*	10.00%	9.07%	8.74%	8.37%	9.01%
2	*8.69%*	10.16%	8.60%	8.20%	8.08%	8.42%
3	*7.92%*	9.40%	7.78%	7.40%	7.24%	7.78%
4	*6.73%*	7.63%	6.54%	6.13%	6.11%	7.26%
Big	*5.25%*	5.84%	5.08%	4.62%	5.03%	5.70%
Standard Deviation of Through-2011 CDRs						
	Averages	*7.48%*	*6.73%*	*6.44%*	*6.42%*	*6.95%*
Small	*7.53%*	8.05%	7.72%	7.47%	7.02%	7.40%
2	*7.51%*	8.46%	7.48%	7.26%	7.01%	7.36%
3	*7.23%*	8.15%	7.16%	6.59%	6.72%	7.53%
4	*6.45%*	7.17%	6.22%	5.78%	6.12%	6.93%
Big	*5.29%*	5.55%	5.07%	5.09%	5.22%	5.53%

Source: Research Affiliates, based on data from Compustat and CRSP.

EXHIBIT 7
Illustration of the Dispersion Effect

	Small-Growth Companies			Large-Value Companies		
	Winner	**Loser**	**Portfolio**	**Winner**	**Loser**	**Portfolio**
Initial Price	$100.00	$100.00	$200.00	$100.00	$100.00	$200.00
Dividends	$3.00	$3.00	$6.00	$6.00	$6.00	$12.00
	$3.45	$2.55	$6.00	$6.60	$5.70	$12.30
	$3.97	$2.17	$6.14	$7.26	$5.42	$12.68
	$4.56	$1.84	$6.41	$7.99	$5.14	$13.13
	$5.25	$1.57	$6.81	$8.78	$4.89	$13.67
	$6.03	$1.33	$7.37	$9.66	$4.64	$14.31
	$6.94	$1.13	$8.07	$10.63	$4.41	$15.04
	$7.98	$0.96	$8.94	$11.69	$4.19	$15.88
	$9.18	$0.82	$9.99	$12.86	$3.98	$16.84
Final Dividend + Terminal Value	$362.50	$18.35	$380.85	$365.00	$81.50	$446.50
CDR	16.60%	−12.78%		18.69%	2.92%	
Mean CDR	1.91%			10.81%		
Portfolio Return			9.18%			13.06%

Source: Research Affiliates, based on data from Compustat and CRSP.

For our small-growth stocks, the winning stock rises more than 250% in the decade, while the loser sheds more than 80% of its value. This sounds rather like the big winners and losers in a typical small growth portfolio. Combine the two stocks—with an average CDR of 1.91%—and our portfolio produces a return of 9.18%, which is 7.27% better than the two components' average CDR. The broad return dispersion means that our big winner adds far more value than our big loser subtracts.

For our large-value category, the winner again rises more than 250%, while the loser pares its dividends year after year and winds up losing about 20% of its value after 10 years. Combine these two stocks into a portfolio, and it delivers 13.06%. As with the small-growth stocks, the winning outlier in our large-value portfolio helped us much more than the loser hurt us. Our portfolio outpaces our average CDR of 10.81% by 1.25%. In the small-growth portfolio, the winning outlier makes a whopping 7.27% difference. Our dispersion is narrower for large value than for small growth, leading to a smaller boost from the winning outlier.

Of course, our hypothetical example focuses specifically on the outliers. Two-thirds of our portfolio lies between these two poles. In our final test, instead of taking the average of individual stocks' CDRs in each quintile portfolio, we construct the actual portfolios that our quintile partitioning creates and examine the annualized aggregate portfolio returns over subsequent years.[17] At each year's end, starting from 1955 to 2001, we partition stocks into 25 size and style portfolios and hold them until the end of 2011. It is unsurprising that the inverted size effect we observe in Exhibit 5 vanishes, as shown in Exhibit 8.[18]

CONCLUSION

In this study, we extend the intuitive framework of clairvoyant value to its complement, the clairvoyant discount rate: the annualized discount rate over an entire holding period that justifies a security's initial purchase price. By looking at the distribution of CDRs, we gain a sharper insight into whether a discount rate explanation may be responsible for some of the apparent pricing inconsistencies observed in CVs and generate some new insights into value and size effects.

EXHIBIT 8
Buy-and-Hold Portfolio Total Returns through 2011

		Growth	**2**	**3**	**4**	**Value**
	Averages	***9.53%***	***10.31%***	***11.24%***	***11.46%***	***11.61%***
Small	***10.15%***	7.44%	9.32%	10.84%	11.77%	11.38%
2	***10.97%***	9.72%	10.16%	11.63%	11.37%	11.96%
3	***11.33%***	10.01%	10.80%	11.55%	11.89%	12.38%
4	***11.32%***	10.48%	11.00%	11.46%	11.62%	12.03%
Big	***10.39%***	10.00%	10.28%	10.72%	10.63%	10.31%

Source: Research Affiliates, based on data from Compustat and CRSP.

For our hypothetical world of long-term buy-and-hold investors, a vastly lower discount rate must be applied to justify the choice of most growth stocks with high valuation multiples over value stocks with low valuation multiples. The average CDR (or prospective required rate of return) for value stocks ranges from 50% to more than 200% higher than the corresponding rate of return earned from the average growth stock. More telling, the average growth stock is priced to generate returns below default-free proxies.

Further, the well-documented size effect is missing from our results. We provide an explanation based on dispersion, with small-cap portfolios winning due to a handful of huge outliers. Neither of these findings seem consistent with variations in discount rates as being chiefly responsible for the value and size effects, unless we're prepared to accept the notion of a negative risk premium for growth stocks. Instead, they respectively suggest a role for pricing errors and skewed payoffs.

The CV concept, measuring what stocks were truly worth given a presumed discount rate, and the CDR concept, measuring the discount rate that a clairvoyant must have assumed to justify past stock prices, provide powerful and intuitive tools to better understand the market-pricing mechanism. With the CDR, we can see that investors in growth stocks, especially small growth stocks, must have historically been content to earn anemic returns—even a negative risk premium—on their average holdings. They were (and perhaps still are) lottery investors, looking for a few big winners to offset a very large number of losing individual investments. CV and CDR offer us insights that we can only see with perfect hindsight.

ENDNOTES

We thank Jaynee Dudley, Li-Lan Kuo, Philip Lawton, Katy Sherrerd, Darren Wagner, and Lillian Wu for their excellent research and editorial assistance.

[1]We use the same measures used to define the Fundamental Index® concept in Arnott, Hsu, and Moore [2005]: sales, cash flow, book value, and dividends.

[2]ALS 2009 employed one of two estimates for the discount rate for each company: a common, market-clearing discount rate that established the S&P 500 Index as fairly valued, and a risk-adjusted approach that varied the discount rate across companies, based on the CAPM. This left open the issue of whether differences in discount rates might explain the pricing differences, given the doubts that exist over the CAPM as an adequate adjustment for risk. We explore this remaining puzzle element in this article.

[3]This has some important implications for investors who like to choose among trendy, high-flying small-growth stocks. They ought to diversify, to assure that they catch some of the spectacular winners, or they need to be awfully good stock pickers!

[4]The 56-year span is the longest span we tested, from 1956–2011, inclusive.

[5]GM was delisted from NYSE on June 2, 2009, at a price of $0.61. The cumulative loss in price, even adjusting for intervening stock splits, is actually 95.6%, if we don't include the dividend distributions.

[6]An alternative interpretation is an implicit assumption that the reinvestment rate for dividends is the IRR over the study span (Dudley [1972]). The two interpretations are equivalent if we consider IRR as an answer to the question, "At which discount rate am I indifferent about the investment opportunity—to spend all the cash I get, or to reinvest all the cash at the discount rate?"

[7]The documented findings here remain robust with the inclusion of the CDRs starting after 2001. Naturally, all of the through-2011 results include the effects of the global financial crisis, so we're not arbitrarily excluding these very important years from our study.

[8]This is the reason for excluding companies with negative financials.

[9]The CDR calculated for starting dates after year-end 2001 is dropped from the tables and graphs because of the noise in the calculation for the short spans.

[10]This starting yield for a 10-year T-bond will very nearly approximate its 10-year buy-and-hold IRR, which is a fair apples-with-apples comparison with the buy-and-hold stocks in our test.

[11]We might think of this as a single-stock portfolio.

[12]Again, we ignore more recent data, because the CDR calculation for any later starting year will include less than 10 years of data.

[13]Of course, our 10-year and 20-year clairvoyant investor should also know the 10-year and 20-year IRR on bonds and bills. We do not include this comparison, because it's not radically different from what's shown here.

[14]If the proceeds from delisting (acquisition or bankruptcy) are invested in T-bills instead of the S&P 500, the CDR is lower. unsurprisingly. The difference averages 130 basis points per year. This is a noteworthy difference, considering that the average CDR (when delisting proceeds are reinvested into T-bills) is only 6.9% per annum across all of our test spans, including all size and value quintiles. The difference between equally weighted and value-weighted port-

folios remains slight. These results are not reported, but are available upon request.

[15]The finding is robust if we change the reinvestment assumption and the difference between size groups' standard deviations is significant, based on the ANOVA test.

[16]This illustrative example was constructed as follows. Our two small-growth stocks both begin with a dividend yield of 3%; the large-value stocks begin with a yield of 6%. These are reasonably typical for the average growth or value stock over the past 50 years. The winning large-value stock sees its dividend rise 10% a year; the loser falls 5% a year. For the small-growth stocks, the growth rates exhibit far wider dispersion, up 15% and down 15% per year, respectively. We set the terminal prices so that winners match these four one-sigma outlier returns. The result is a simple—indeed, simplistic—illustration of how outliers can lead to our vanishing size effect paradox.

[17]This differs from the classic size and style quintile portfolios studied by Fama and French, as well as by many others. Most such studies involve reconstituting the portfolio annually, based on the latest financials and market cap. Our portfolios are buy-and-hold portfolios, held for the entire 10-year or 20-year span, or held for the full span until the end of 2011. For this reason, our size and value effects are markedly more muted than the classic Fama–French results, which incorporate a rebalancing return that has a profound affect on long-term results. It is interesting to observe that the entire size effect, and most of the value effect, is reliant on this rebalancing return (Chaves and Arnott [2012]).

[18]If we rebalance the portfolio every year, the size effect reported in Fama–French [1993] can be completely restored.

REFERENCES

Arnott, R., J. Hsu, and P. Moore. "Fundamental Indexation." *Financial Analysts Journal*, Vol. 61, No. 2 (March/April 2005), pp. 83-99.

Arnott, R., F. Li, and K. Sherrerd. "Clairvoyant Value and the Value Effect." *The Journal of Portfolio Management*, Vol. 35, No. 3 (Spring 2009a), pp. 12-26.

——. "Clairvoyant Value II: The Growth/Value Cycle." *The Journal of Portfolio Management*, Vol. 35, No. 4 (Summer 2009b), pp. 142-157.

Botosan, C., and M. Plumlee. "Assessing Alternative Proxies for the Expected Risk Premium." *Accounting Review*, Vol. 80, No. 1 (January 2005), pp. 21-53.

Chaves, D., and R. Arnott. "Rebalancing and the Value Effect." *The Journal of Portfolio Management*, Vol. 38, No. 2 (Fall 2012), pp. 59-74.

Claus, J., and J. Thomas. "Equity Premia as Low as Three Percent? Evidence from Analysts' Earnings Forecasts for Domestic and International Stock Markets." *The Journal of Finance*, Vol. 56, No. 5 (October 2001), pp. 1629-1666.

Cochrane, J. "Presidential Address: Discount Rates." *The Journal of Finance*, Vol. 66, No. 4 (August 2011), pp. 1047-1108.

Dudley, C. "A Note on Reinvestment Assumptions in Choosing between Net Present Value and Internal Rate of Return." *The Journal of Finance*, Vol. 27, No. 4 (September 1972), pp. 907-915.

Easton, P., and G. Sommers. "Effect of Analysts' Optimism on Estimates of the Expected Rate of Return Implied by Earnings Forecasts." *Journal of Accounting Research*, Vol. 45, No. 5 (December 2007), pp. 983-1015.

Fama, E., and K. French. "The Cross-Section of Expected Stock Returns." *The Journal of Finance*, Vol. 47, No. 2 (June 1992), pp. 427-465.

——. "Common Risk Factors in the Returns on Stocks and Bonds." *Journal of Financial Economics*, Vol. 33, No. 1 (February 1993), pp. 3-56.

Gebhardt, W., C. Lee, and B. Swaminathan. "Toward an Implied Cost of Capital." *Journal of Accounting Research*, Vol. 39, No. 1 (June 2001), pp. 135-176.

Grossman, S., and R. Shiller. "The Determinants of the Variability of Stock Market Prices." *American Economic Review*, Vol. 71, No. 2 (May 1981), pp. 222-227.

Guay, W., S. Kothari, and S. Shu. "Properties of Implied Cost of Capital Using Analysts' Forecasts." *Australian Journal of Management*, Vol. 36, No. 2 (August 2011), pp. 125-149.

Lee, C., D. Ng, and B. Swaminathan. "Testing International Asset Pricing Models Using Implied Costs of Capital." *Journal of Financial and Quantitative Analysis*, Vol. 44, No. 2 (April 2009), pp. 307-335.

Pástor, L'., M. Sinha, and B. Swaminathan. "Estimating the Intertemporal Risk-Return Tradeoff Using the Implied Cost of Capital." *The Journal of Finance*, Vol. 63, No. 6 (December 2008), pp. 2859-2897.

Shiller, R. "Do Stock Prices Move Too Much to Be Justified by Subsequent Changes in Dividends?" *American Economic Review*, Vol. 71, No. 3 (June 1981), pp. 421-436.

——. "The Volatility of Stock Market Prices." *Science*, Vol. 235, No. 4784 (January 2, 1987), pp. 33-37.

Williams, J.B. *The Theory of Investment Value*. Flint Hill, VA: Fraser Publishing, 1938.