PORTFOLIO MANAGEMENT RESEARCH
By With Intelligence

the journal of PORTFOLIO management

volume 41
number 4

SUMMER 2015

jpm.pm-research.com

Rip Van Winkle Indexing

Robert Arnott, Noah Beck, and Vitali Kalesnik

research affiliates


research
affiliates

Rip Van Winkle Indexing

Robert Arnott, Noah Beck, and Vitali Kalesnik

Robert Arnott
is the chairman at Research Affiliates, LLC, in Newport Beach, CA.
arnott@rallc.com

Noah Beck
is a senior researcher for equity research at Research Affiliates, LLC, in Newport Beach, CA.
beck@rallc.com

Vitali Kalesnik
is a partner and head of equity research at Research Affiliates, LLC, in Newport Beach, CA.
kalesnik@rallc.com

If market prices efficiently reflect all known information about a company's future growth and stability, a capitalization-weighted portfolio that uses the most current information should deliver better risk-adjusted performance than alternative choices do. Sharpe [1964] demonstrated that, if we begin with an array of simplifying assumptions, including efficient markets, we can prove that cap-weighted indexing is mean–variance optimal. His theoretical revolution is gaining real-world traction. In 2013, more than one-sixth of global equity mutual fund assets were managed in accordance with a passive index approach, up from less than 9% in 1998 (Investment Company Institute [2013]).

If markets are less than perfectly efficient, opportunities exist to outperform the cap-weighted market portfolio; conventional wisdom is that money managers must act quickly on new information to realize profits. Our research suggests there is another, very simple way to make money. We call it the Rip Van Winkle approach. According to Washington Irving's story, Rip Van Winkle was an idler who could not support his family. Irving wrote, "The great error in Rip's composition was an insuperable aversion to all kinds of profitable labor." After a drinking bout, he fell asleep for 20 years. While he slept, the American Revolutionary War began, ran its course, and ended.[1] Rip Van Winkle awoke to a very different world.

In our narrative, Rip Van Winkle's lassitude causes him to reuse the 20-year-old market capitalization weights that he knew before his nap, leaving out stocks that no longer exist in their original form, as the basis for his current investment. So intrigued is he by the long interlude that he decides to maintain a 20-year gap in his subsequent portfolios and thus rebalances annually back to the next 20-year-old stale-weight portfolio. In testing over the past 67 years, this approach produces a risk-adjusted alpha of about 180 basis points per annum, which would place Rip Van Winkle in the very top cohort of active managers.[2]

To be entirely clear, we do not think Rip Van Winkle indexing is an attractive real-world strategy. Any investors who implemented the Rip Van Winkle strategy would see bizarre weightings as ghosts of the long-forgotten 1981 energy bubble arrive in Rip's 2001 allocations, and the tech bubble arises again in the 2019 and 2020 portfolio weights, which then plummet in 2021 and 2022. Rip's investors would have to wait until May 2032 for Facebook to show up in their portfolios. In addition, trading on price movements that took place in the distant past would introduce needless turnover and high tracking error.

We do not recommend Rip Van Winkle indexing. Nonetheless, the concept advances our understanding of index investing, in theory and in practice. Rip Van Winkle's approach to indexing represents an excellent controlled experiment to quantify the effect that past price movements, reflected in portfolio weights, would have on the performance of a cap-weighted index.

PREVIOUS LITERATURE

Rip Van Winkle's successful performance is at odds with common sense and arguably at odds with the capital asset pricing model (CAPM) assertion that cap-weighted indices are mean–variance optimal. Discarding 20 years of information should not be a recipe for investment success. The superior return is less surprising, however, once we accept the empirical evidence that stock returns historically exhibit some predictability, notably in the amply documented value and long-horizon mean reversion effects.

Recent short-term winners (stocks with high returns over the past 6 to 12 months) tend to outperform, while recent short-term losers tend to underperform. Known as momentum, this effect was documented by Jegadeesh and Titman [1993]. Conversely, long-term winners tend to be tomorrow's losers. Known as long-term mean reversion, this effect was documented by De Bondt and Thaler [1985] and Poterba and Summers [1988].[3]

By construction, a market-cap-weighted portfolio has higher weights in both short-term and long-term winners. Holding the former could potentially help the cap-weighted investor, while the latter create a drag on returns. In contrast, Rip Van Winkle forfeits all the benefits of short-term momentum, but he also progressively eliminates long-term price movements the longer he sleeps, removing this particular drag that plagues the cap-weighted investor.

Rip Van Winkle's performance is also consistent with a growing literature documenting that, for investors patient and resilient enough to stay the course, almost any non-price weighting scheme would generate excess returns compared with cap weighting. Arnott et al. [2005] noticed this phenomenon and suggested that we can add a dynamic value tilt—and retain the low turnover and vast capacity of cap weighting—by severing the link between a stock's price and the weight that it represents in our portfolio. Since the mid-2000s, a whole industry of smart beta[4] indices has appeared; the common denominator for these products is that company weights in the portfolio are determined in a fashion that is independent of price.

We are not the first to study stale cap-weighted portfolios' performance. Chen, Chen, and Bassett [2007] demonstrated that investors would earn better returns than would a cap-weighted index if they used a smoothed portfolio of stale cap-weighted indices. In this article, we examine a much longer history of data, use much longer cap-weighting lags, explore international results, study the term structure of the interaction between the past mispricing and the subsequent return drag on cap-weighted index performance, and compare the results of stale cap-weighting with clairvoyant future cap-weighting. That said, we owe a debt of primacy to Chen, Chen, and Bassett as pioneers in this space.

RIP VAN WINKLE'S STALE INDEX PORTFOLIO

Our simulations of Rip Van Winkle's results start with a cap-weighted index of the 1,000 largest market-cap stocks domiciled in the United States, looking back over various time lags. The portfolio holdings are rebalanced annually to the following year's cap weights.[5] We simulated the cap-weighted portfolios with lags of zero, one, two, three, five, seven, ten, fifteen, and twenty years. The zero-lag portfolio is, of course, the current cap-weighted portfolio, reconstituted and rebalanced annually at the beginning of each year. Let us call the cap-weighted portfolio from one year ago a 1-year stale portfolio and the portfolio that relies on 20-year-old cap weights a 20-year stale portfolio.

We used CRSP/Compustat and Worldscope/Datastream merged data sets, free of survivorship bias, to compute both ending market capitalizations and portfolio returns for each calendar year. We have data from year-end 1926 on the market capitalization and returns for U.S.-domiciled companies. Given that we need to wait 20 years for the longest-deferred portfolio, our return sample starts from the end of 1946 in the United States, so our first calendar year is 1947. To select and weight our 1,000-stock, zero-lag portfolio (an ordinary current cap-weighted portfolio) for use in calendar year 1947, we used the year-end 1946 market capitalizations. To select and weight our 1,000-stock, one-year stale portfolio for 1947, we used the year-end 1945 market capitalizations,

and to construct our 20-year stale portfolio for 1947, we used the year-end 1926 market capitalizations.

After measuring the return on each of these portfolios for calendar year 1947, we rolled forward and computed the 1948 portfolios in much the same way, and so forth. The difference in weights between the current and the one-year stale weights is dominated by the difference in price over the past year (corporate actions have a modest additional affect). As Rip Van Winkle holds his portfolio, the portfolio weights will accumulate some price drift. To eliminate this link between the price and the weights, we reconstitute and rebalance the stale-weight Rip Van Winkle portfolios each year. Naturally, the longer the lag, the more stale the portfolio weights, and the less correlated the weights are with current prices.

Of course, not all the companies that existed at year-end 1926 still exist in 1947, and not all companies that exist in 1947 existed at the end of 1926. What would Rip do? He's indolent, so he would ignore all the companies that didn't exist in 1926. Reciprocally, if a company happened to fail or was acquired in the intervening years, Rip would simply remove this stock and invest its resulting weight in the remaining companies, in proportion to their old capitalizations.

In investing, laziness sometimes pays. As Exhibit 1, panel A shows, not only do we find that the stale-weight portfolios actually perform better than the current

EXHIBIT 1

U.S. Stale Portfolio Return and Volatility, 1947–2013

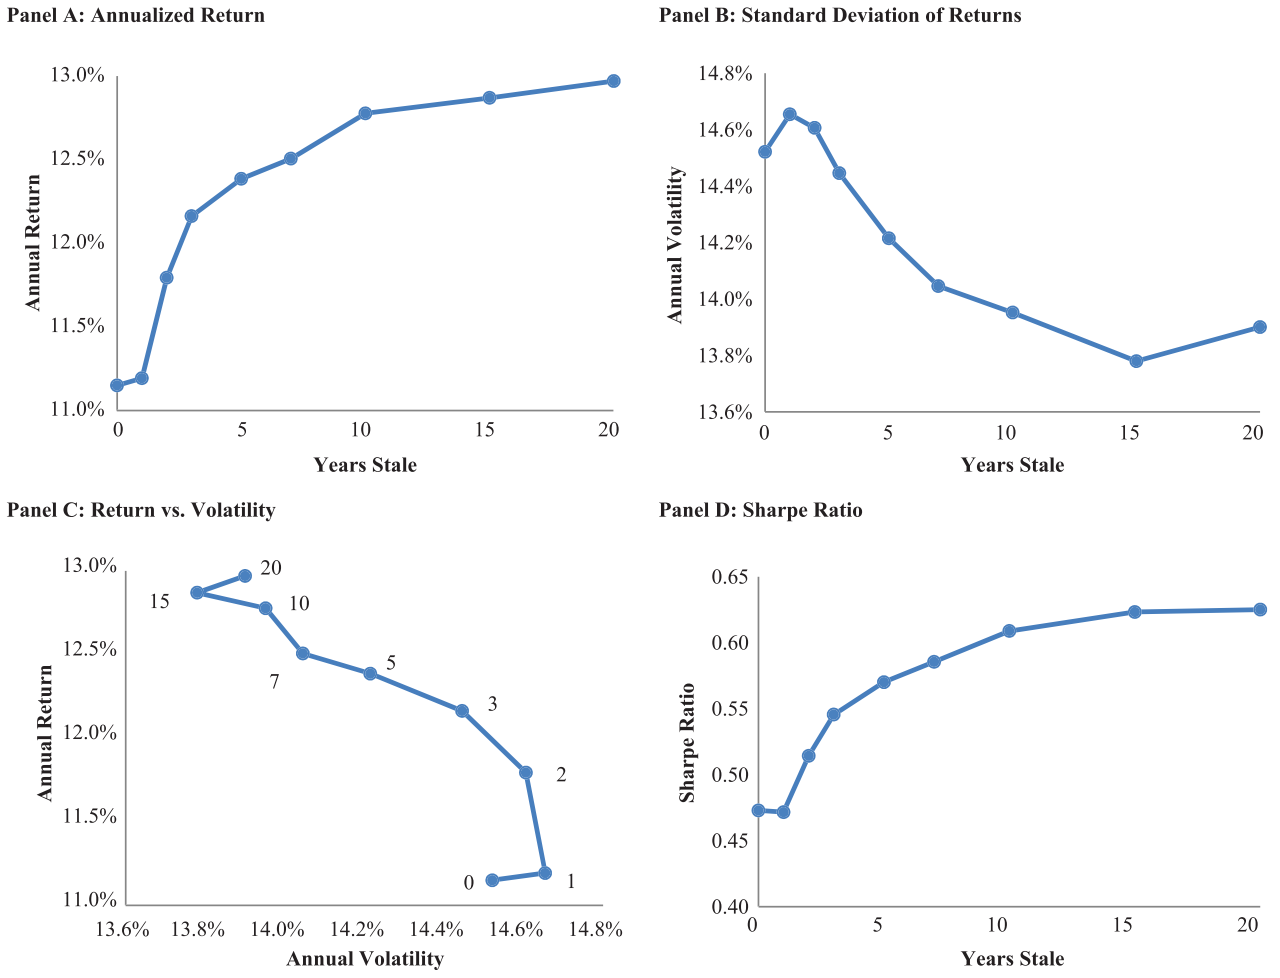

Source: Research Affiliates using data from CRSP, Compustat, Worldscope, and Datastream.

market portfolio, but also that the longer we sleep on the information, the better our performance!

Do the stale-weight portfolios achieve high returns by taking on unnecessary risks? After all, using outdated market information exposes the portfolio to potentially large overweights in companies that declined (and large underweights in companies that thrived). But as we can see in panel B, the risk does not rise; on the contrary, once we're past one year of stale data, the outperformance of the stale-weight portfolios is achieved with steadily lower volatility. Note that when we use five-year-old cap weights, we already garner over half of the performance advantage, and over half of the risk reduction, of using the full 20-year stale weighting. Panel C shows the path the portfolios take through return/volatility space as we form portfolios with progressively older data. Panel D shows the monotonic increase in Sharpe ratio after the first year, in which momentum holds us back.

RIP VAN WINKLE AND CLAIRVOYANT INDEX PERFORMANCE

Suppose, instead of Rip Van Winkle building us a portfolio based on 20-year-old cap weights, a genie offers to construct us a portfolio based on cap weights 20 years in the future. (The minor catch: he or she won't include stocks that don't yet exist.) My goodness! When Microsoft first went public, instead of weighting it in 1987 based on a year-end 1986 price of 12 cents (split-adjusted to today), we could give it 200 times that weight, based on a year-end 2006 price of $25. This kind of clairvoyance would have substantial and obvious value.

The notion of clairvoyance—examining how people might invest if they could foresee future prices—has proven to be a marvelous means for exploring the nature of the value effect (Arnott et al. [2009a]), the nature of the growth/value cycle (Arnott et al. [2009b]), and the discount rates that investors would have had to assume in order to justify past prices (Arnott et al. [2013]). It is no less powerful here. How does the value added by a stale cap-weighted portfolio compare with the value that might be added by a crystal ball enabling us to see future cap weights?

As many have observed and as the literature amply documents, we add far more value when we sever the link between a company's portfolio weight and its price than we cede by losing the information that today's price conveys about the fair value of a company's business—because that information is already in the price. However, introducing a link to the future price is another matter entirely. It seems self-evident that a portfolio constructed on the basis of future cap weights—a clairvoyant index fund—should handily outperform one based on current cap weights. Although the current cap weight contains valuable information about the market's consensus on a company's future prospects, that information is already in the price. By contrast, the future cap weight contains quite substantially more information about the future of a company, information that is not in today's price, because this cap weight is the company's future. Of course, clairvoyant cap-weighted portfolios win!

The left half of Exhibit 2 shows the results for clairvoyance. In contrast with stale Rip Van Winkle portfolios, where "more stale" means "more value added," the clairvoyant index funds do not add more value as we look further into the future. Weighting on one-year forward cap weights outperforms the current cap-weighted portfolio by about 6% per year. Looking ahead 10 or 20 years delivers about the same benefit, with a small increase in volatility and a significant increase in tracking error.

Exhibit 3 displays some of these results graphically. In panel A, we look at the comparative value-added relative to current cap weighting for the maximum overlapping period, maximum stale portfolio period, and maximum clairvoyant portfolio period.[6] The similarity across all three time spans is striking. Clairvoyance offer a 6% to 7% incremental return. There's a substantially lower return for using current cap weight and one-year stale cap weights. The power of the short-term momentum effect is evident in the fact that the one-year-old cap-weighted portfolio performs about the same as the current cap-weighted portfolio.

In panel B, we compare the risk of the clairvoyant and stale portfolios. Because the relative results are nearly identical across the different spans, we only used the span that works for both clairvoyant and stale portfolios: 1947 to 1994. Again, we see broad consistency. The clairvoyant portfolios are riskier than the current cap-weighted portfolio, while Rip Van Winkle's stale portfolios are generally less risky. Clairvoyant portfolios obviously are weighted toward the winners of the future; these tend to be the stocks of rapidly growing companies, which are disposed to be more volatile. The stale port-

EXHIBIT 2

U.S. Stale Portfolio and Clairvoyant Portfolio Return Attributes

	Clairvoyant Portfolios (Years Look-Ahead)					Cap Wgt	Stale Rip Van Winkle Portfolios (Years Lagged)							
	–20	–10	–5	–3	–1	0	1	2	3	5	7	10	15	20
1947–1994, Maximum Overlapping Span for 20-Year Stale and Clairvoyant Portfolios														
Annual Return	17.5%	17.9%	17.8%	17.8%	17.5%	**11.6%**	11.4%	11.9%	12.2%	12.6%	12.7%	12.9%	12.8%	13.1%
Std Deviation	15.0%	14.9%	14.5%	14.4%	14.4%	**14.1%**	14.0%	14.0%	13.9%	13.8%	13.6%	13.6%	13.4%	13.5%
Sharpe Ratio	0.85	0.87	0.90	0.90	0.88	**0.48**	0.46	0.50	0.52	0.56	0.57	0.59	0.59	0.61
Value Added	5.9%	6.3%	6.2%	6.2%	5.9%	**0.0%**	–0.2%	0.3%	0.6%	0.9%	1.1%	1.3%	1.2%	1.5%
Tracking Error	3.3%	3.0%	2.1%	1.8%	1.4%	**0.0%**	1.2%	1.8%	2.1%	2.5%	2.7%	2.9%	3.2%	3.5%
Info Ratio	1.80	2.11	3.01	3.46	4.37	**0.00**	–0.18	0.14	0.27	0.38	0.41	0.45	0.39	0.44
1927–1994, Maximum Span for 20-Year Clairvoyant Portfolios														
Annual Return	15.6%	16.0%	15.9%	16.1%	16.1%	**9.9%**								
Std Deviation	22.6%	22.2%	21.9%	21.6%	21.3%	**20.3%**								
Sharpe Ratio	0.53	0.55	0.56	0.58	0.58	**0.31**								
Value Added	5.7%	6.1%	6.1%	6.2%	6.3%	**0.0%**								
Tracking Error	3.7%	3.1%	2.5%	2.2%	1.7%	**0.0%**								
Info Ratio	1.56	1.98	2.46	2.84	3.63	**0.00**								
1947–2013, Maximum Span for 20-Year Stale Portfolios														
Annual Return						**11.2%**	11.2%	11.8%	12.2%	12.4%	12.5%	12.8%	12.9%	13.0%
Std Deviation						**15.3%**	15.5%	15.4%	15.3%	15.1%	14.9%	14.8%	14.6%	14.8%
Sharpe Ratio						**0.45**	0.45	0.49	0.52	0.54	0.55	0.57	0.59	0.59
Value Added						**0.0%**	0.0%	0.6%	1.0%	1.2%	1.4%	1.6%	1.7%	1.8%
Tracking Error						**0.0%**	2.4%	3.0%	3.2%	3.6%	4.0%	4.2%	4.3%	4.5%
Info Ratio						**0.00**	0.02	0.21	0.32	0.34	0.34	0.38	0.40	0.41

Source: Research Affiliates using data from CRSP, Compustat, Worldscope, and Datastream.

folios tend to be dominated by yesteryear's giants, which aren't growing as fast and therefore are less volatile.

Of all possible dates—past, present, and future—the current cap weights at $t = 0$ offer the very worst performance. Of course, this is exactly what a large (and increasing) number of investors use to create their portfolios.

Panel C shows the comparative Sharpe ratios (left axis) and information ratios (right axis). Here we observe the declining Sharpe and information ratios as the clairvoyant portfolio ventures further into the future. We observe the steadily increasing ratios for stale portfolios as Rip chooses to sleep longer. The performance benefit that Rip gets by sleeping is eventually a surprisingly large fraction of the benefits derived from perfect clairvoyance. Comparing the 20-year stale and 20-year clairvoyant portfolios, Rip would deliver about one quarter of the added value delivered by the genie, with about one quarter of the genie's information ratio. Adjusted for Rip's lower risk, Rip delivers more than one-third as much of an improvement in the Sharpe ratio as the genie, offering 20 years of clairvoyance, can deliver.

Finally, in panel D, we show the comparative wealth of an investor who settles for very stale cap weights, relative to the more conventional cap-weighted index fund investor. For this purpose, we focus on the stale Rip Van Winkle portfolios, because we know that clairvoyance is equally valuable and impossible. This allows us to carry the analysis through the span from 1947 to 2013. Several points merit comment.

- Adding 1.8% per year for two-thirds of a century compounds rather nicely. The top line shows that the 20-year-lagged Rip Van Winkle portfolio delivers nearly three times the end-point wealth of cap weighting over this span.
- A five-year-lagged portfolio accomplishes roughly half as much, leading to twice the wealth of a cap-weighted investor over this same span.
- The results are episodic. This strategy can lead to protracted periods of disappointment, as indeed can almost any strategy.
- We can see more or less flat results for the first 20 years. In other words, had we tested this idea over the past 40 years instead of the past 67 years,

EXHIBIT 3

U.S. Stale Portfolio and Clairvoyant Portfolio Performance

Panel A: Annualized Returns

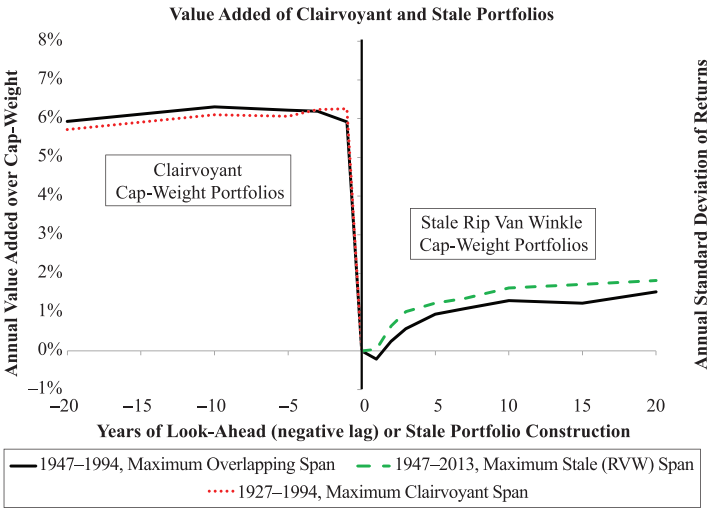

Panel B: Standard Deviation of Returns

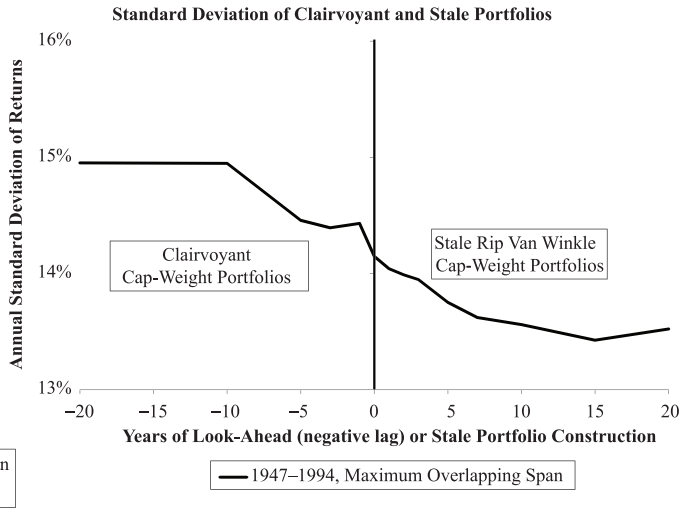

Panel C: Sharpe (Solid, LHS) and Information (Dashed, RHS) Ratios

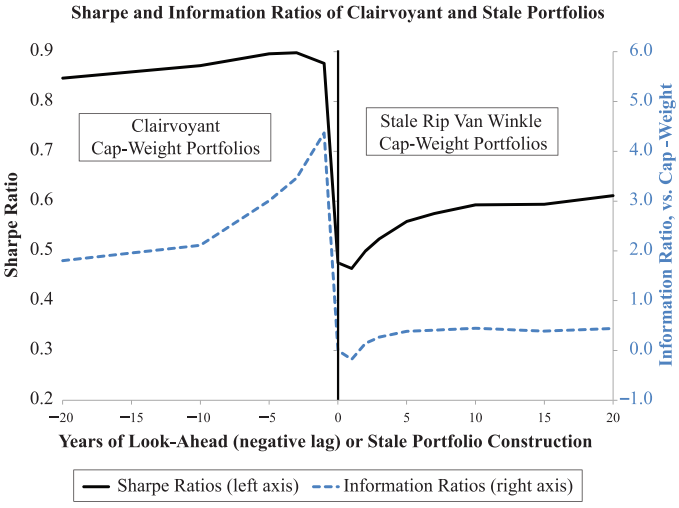

Panel D: Cumulative Relative Wealth (Value Added), 1950--2013

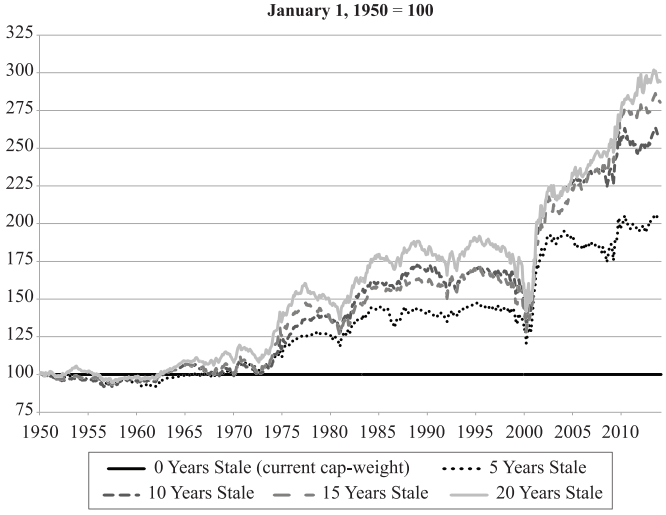

Source: Research Affiliates using data from CRSP, Compustat, Worldscope, and Datastream.

we'd have seen annualized incremental returns of nearly 300 basis points for the 20-year stale index portfolio. As always, start and end dates are very, very important.

- Even more notably, the Rip Van Winkle portfolio was oblivious to many of the companies that fueled the technology bubble in the late 1990s. It would seem that investing in the fads of yesteryear, or even past decades, is more profitable than investing in the fads of today, because today's fads are fully priced, while the fads of yesteryear rarely carry such a large valuation premium.
- During the global financial crisis, as the market was pricing in financial Armageddon, our slumbering investor awakes from his nap unaware of the calamity and allocates to financials as if it were any other year. Once again, Rip's obliviousness to the fears embedded in post-crash prices turns out to be an advantage, not a disadvantage.[7]

EXHIBIT 4

Developed and Emerging Market Rip Van Winkle and Clairvoyant Portfolio Performance

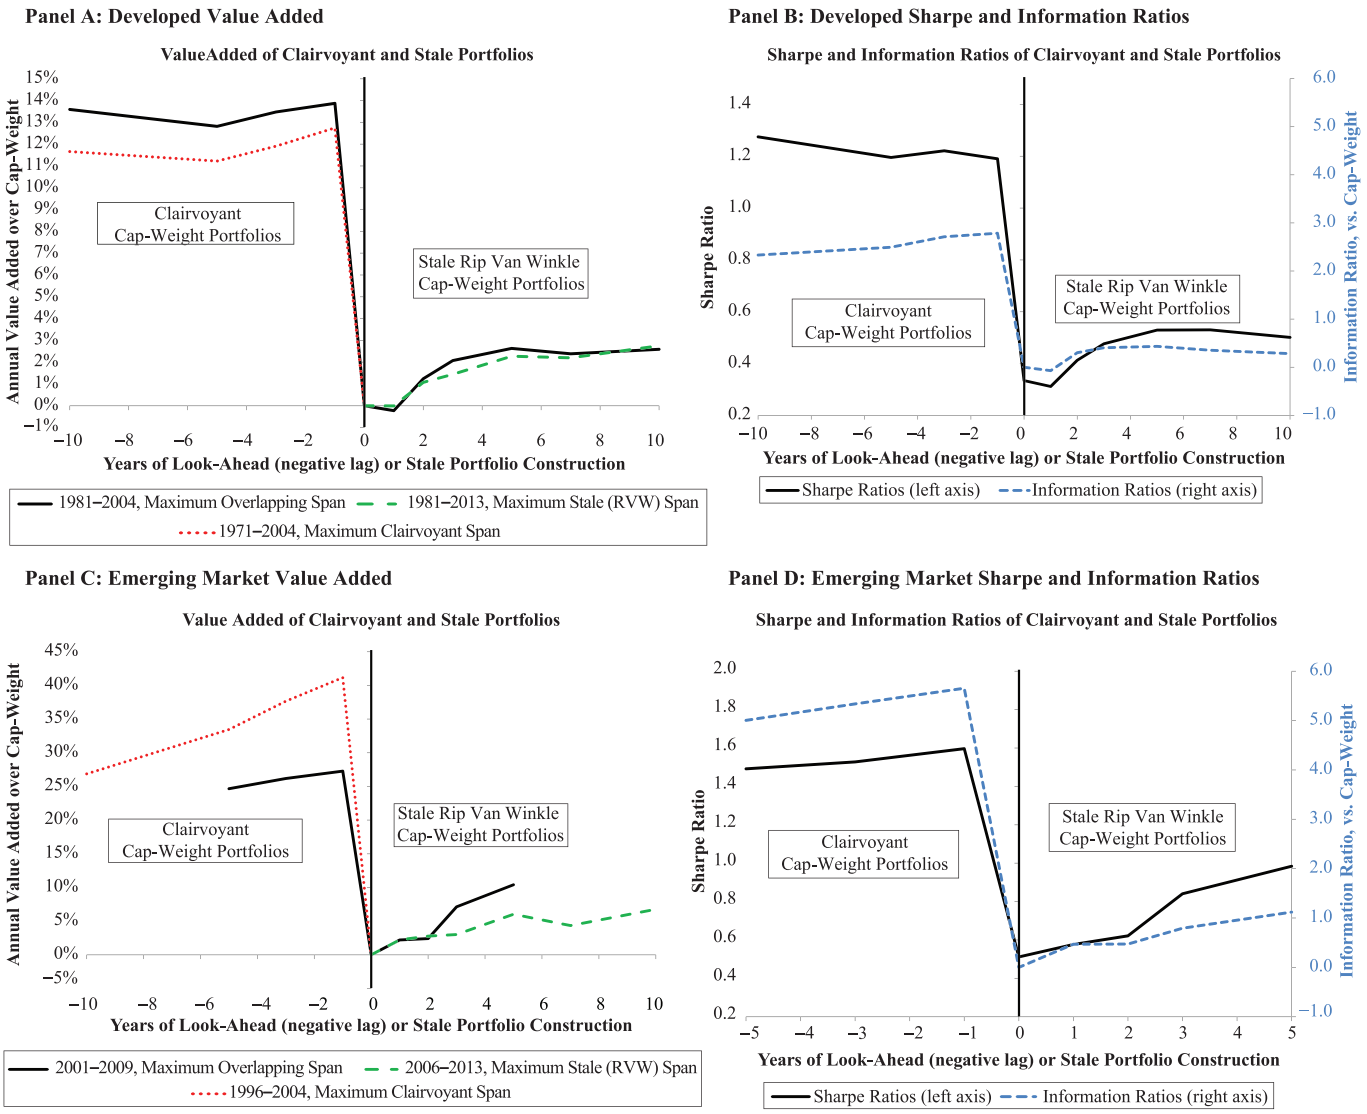

Source: Research Affiliates using data from CRSP, Compustat, Worldscope, and Datastream.

We repeat this exercise in the developed and emerging markets, where our market cap and returns data begin from year-end 1970 and year-end 1986, respectively. Because of these markets' shorter histories, we do not allow Rip Van Winkle to sleep any longer than 10 years; these simulations begin in 1981 and 1997, respectively. Exhibit 4 shows the results for developed (panels A and B) and emerging markets (panels C and D). Here we narrow our focus to the value-added, and to the Sharpe and information ratios.[8] To be sure, the time span in these markets is too short to have compelling statistical significance, but adding 240 to 700 basis points per year, over periods of 22 and 7 years, respectively, is certainly economically meaningful.

ISOLATING THE SOURCES OF RETURN

To understand the drivers of relative performance, we conducted an attribution analysis using a five-factor Fama–French model, as shown in Exhibit 5. In addition to the classic three-factor Fama–French attribution

EXHIBIT 5

Attribution of Returns, Fama–French Five-Factor Model, 1947–1994, Maximum Overlapping Span for 20-Year Clairvoyant and Stale Portfolios

	Clairvoyant Portfolios (Years Look-Ahead)					Cap	Stale Rip Van Winkle Portfolios (Years Lagged)							
	–20	–10	–5	–3	–1	Weight	1	2	3	5	7	10	15	20
Alpha	**5.29%**	**5.40%**	**4.72%**	**4.76%**	**4.40%**	*0.00%*	**0.31%**	**0.50%**	**0.54%**	**0.72%**	**0.79%**	**0.70%**	**0.77%**	**0.69%**
Mkt Beta	1.02	**1.03**	1.01	1.01	**1.01**	*1.00*	1.00	1.00	1.00	**0.98**	**0.98**	**0.97**	**0.96**	**0.97**
SMB	**0.08**	**0.05**	**0.03**	**0.04**	**0.03**	*0.00*	**–0.02**	–0.01	–0.01	**–0.02**	**–0.03**	**–0.04**	**–0.05**	–0.02
HML	**–0.12**	**–0.07**	–0.01	–0.01	0.00	*0.00*	**0.02**	**0.05**	**0.08**	**0.08**	**0.09**	**0.09**	**0.13**	**0.18**
MOM	–0.01	–0.01	**0.04**	**0.04**	**0.07**	*0.00*	**–0.07**	**–0.07**	**–0.07**	**–0.07**	**–0.07**	**–0.05**	**–0.06**	**–0.05**
LTMR	**0.18**	**0.16**	**0.13**	**0.10**	**0.03**	*0.00*	**0.04**	**0.08**	**0.12**	**0.18**	**0.20**	**0.25**	**0.19**	**0.19**

Bold denotes statistical significance, at the 95% level. For Mkt Beta it denotes statistical difference from 1.

Source: Research Affiliates using data from CRSP, Compustat, Worldscope, and Datastream.

(beta, size, and value), the five-factor model includes momentum, as suggested by Asness [1994] and Carhart [1997], and long-term mean reversion, as suggested by De Bondt and Thaler [1985] and Poterba and Summers [1988]. For this purpose, we focus on the span that allows direct comparison of up to 20 years of clairvoyant and stale index portfolios: 1947 to 1994.

Several results stand out in Exhibit 5.

- All the stale and clairvoyant portfolios have positive and statistically significant residual alphas after accounting for the effect of their Fama–French factor loadings. This is unsurprising for the clairvoyant portfolios but very interesting for the stale portfolios.[9]
- All of the clairvoyant portfolios have a small-cap tilt, which is statistically significant and economically meaningful. Of course, this is because we are overweighting companies that are smaller, relative to the contemporaneous cap-weighted portfolio, than they will be at the end of the measurement period. All of the stale portfolios have a large-cap tilt, net of long-term mean reversion, but not all are statistically significant and none is economically large.
- All of the stale portfolios have a value tilt that is statistically significant and, apart from the one-year stale portfolio, economically meaningful. None of the clairvoyant portfolios exhibit a value tilt; indeed, the 10-year and 20-year look-ahead portfolios have pronounced growth tilts.
- All of the stale portfolios have an anti-momentum bias, which is always statistically significant and economically meaningful. The clairvoyant portfolios exhibit a positive momentum bias out to five years, whereupon it oddly changes sign and becomes insignificant.
- All the portfolios, clairvoyant and stale alike, have statistically significant loadings on the fifth factor: long-term mean reversion; those loadings are economically significant for all look-ahead and look-back periods greater than one year.

The big surprise in this five-factor attribution analysis is the residual alpha. Roughly half of the value added by the Rip Van Winkle strategy, whether he sleeps for two years or twenty, is not explained by the five factors in aggregate and notably is not explained by the value effect. Moreover, the alpha is always positive and often highly statistically significant. Because no one would claim that Rip Van Winkle has any skill whatsoever, this alpha suggests that the five-factor model is incomplete; it is missing some priced factor that has economically meaningful scale.[10] As in previous articles (Chow et al. [2011] and Arnott et al. [2013]), we invite the finance community to join the hunt!

INVESTMENT CHARACTERISTICS OF RIP VAN WINKLE'S PORTFOLIOS

Cap weighting has excellent investment characteristics, including immense capacity, high liquidity, and low turnover. By these metrics, Rip Van Winkle's strategy is—unsurprisingly—not much worse than the cap-weighted index. In fact, as Exhibit 6 shows, those who lack an intimate knowledge of market history would have a hard time matching the four portfolios to their original timeframe. Large companies tend to stay large for many years. So, Rip Van Winkle's 20-year stale-weight

EXHIBIT 6

List of Top 10 Stocks for Current, 5, 10, and 20-Year Stale Cap-Weighted Portfolios, January 2014

Current Cap-Weighted Portfolio			5-Year Stale Cap-Weighted Portfolio		
Rank	Name	Weight (%)	Rank	Name	Weight (%)
1	Apple	2.62	1	Exxon Mobil	4.72
2	Exxon Mobil	2.31	2	Wal Mart	2.56
3	Microsoft	1.62	3	Procter & Gamble	2.15
4	Google	1.62	4	Microsoft	2.01
5	General Electric	1.48	5	General Electric	1.98
6	Johnson & Johnson	1.35	6	AT&T	1.95
7	Wal Mart	1.33	7	Johnson & Johnson	1.93
8	Chevron	1.25	8	Chevron	1.75
9	Wells Fargo	1.25	9	Pfizer	1.39
10	Procter & Gamble	1.16	10	JPMorgan Chase	1.37
10-Year Stale Cap-Weighted Portfolio			**20-Year Stale Cap-Weighted Portfolio**		
Rank	Name	Weight (%)	Rank	Name	Weight (%)
1	General Electric	3.34	1	General Electric	4.12
2	Microsoft	3.17	2	Exxon Mobil	3.61
3	Exxon Mobil	2.91	3	Coca Cola	2.67
4	Pfizer	2.89	4	Wal Mart	2.64
5	Citigroup	2.69	5	Altria	2.24
6	Wal Mart	2.46	6	Merck	2.03
7	Intel	2.23	7	Procter & Gamble	1.79
8	AIG	1.85	8	IBM	1.50
9	Cisco	1.79	9	Du Pont	1.50
10	IBM	1.71	10	Pepsico	1.50

Source: Research Affiliates using data from CRSP, Compustat, Worldscope, and Datastream.

portfolio would still be investing in a broadly diversified index of hundreds of recognizable names, with weights that provide broad economic representation even today.

Exhibit 7 shows some of the investment characteristics of the stale-weight portfolios. Panels A through D demonstrate that stale-weighted portfolios preserve very high weighted-average market capitalization and adjusted daily volume for many years. In fact, in the United States, the weighted-average capitalization and trading volume (hence liquidity) tend to increase with the age of our deferred portfolio, as smaller companies leave the universe due to bankruptcy or business combinations. In many now-developed countries, whose stock markets have matured since the 1970s and 1980s, the older stale-weight portfolios would generally have a somewhat smaller weighted-average market capitalization, but they still show better liquidity than today's cap-weighted portfolio.

Obviously, Rip Van Winkle's portfolios experience higher turnover compared to the current cap-weighted index. Cap weights are quite volatile, but cap-weighted portfolios rebalance automatically: as the price of a security increases, the weight in the cap-weighted index also increases, so the investor or fund manager need not execute any rebalancing trades. The turnover of the cap-weighted index comes mostly from corporate actions and from the addition and deletion of constituents due to migrations across size ranges.

The Rip Van Winkle strategy doesn't self-adjust because it is designed to rebalance every year to a new stale portfolio. Rip, therefore, needs to do more trading than a traditional cap-weighted index manager. Surprisingly, not much more! In the United States, turnover increased from 4% to about 14% and in the developed markets from 6% to 17%, regardless of how stale Rip's portfolio may be. This elevated turnover is minimal, when compared with a typical active manager's transactions. Given the strategy's liquidity and hence large capacity, the higher trading costs associated with this rise in turnover are likely to be modest, especially relative to the incremental performance of Rip Van Winkle's strategies.[11]

TERM STRUCTURE OF RIP VAN WINKLE EXCESS RETURNS

Past returns directly influence company weights in the cap-weighted index—the most popular vehicle for passive investing. Yet investors know very little about how the well-documented patterns of long-horizon mean reversion and short-term momentum interact in affecting performance of the cap-weighted index. Rip Van Winkle's portfolios with stale cap weights provide an interesting opportunity to study the term structure of the interaction between past price changes and subsequent cap-weighted index returns.

Exhibit 8 charts the incremental value added as we increase the lag period of the stale cap-weighted portfolios. The one-year stale index underweights the recent winners and overweights the recent losers. Given that

EXHIBIT 7

U.S. and Developed Deferred Market Cap Turnover and Weighted-Average Market Cap

Panel A: U.S. Relative Avg Mkt Cap (1947–2013)

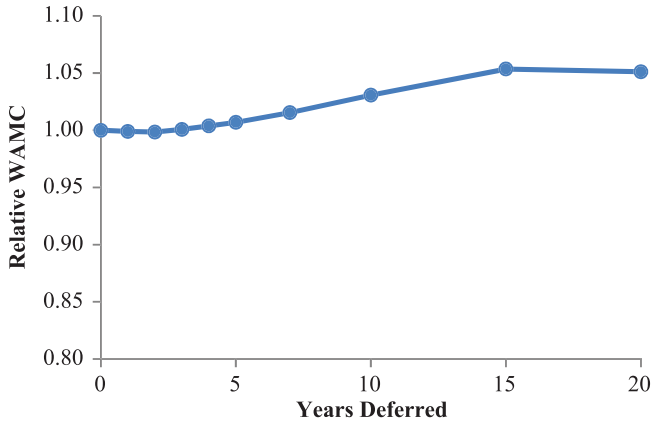

Panel B: Developed Relative Avg Mkt Cap (1981–2013)

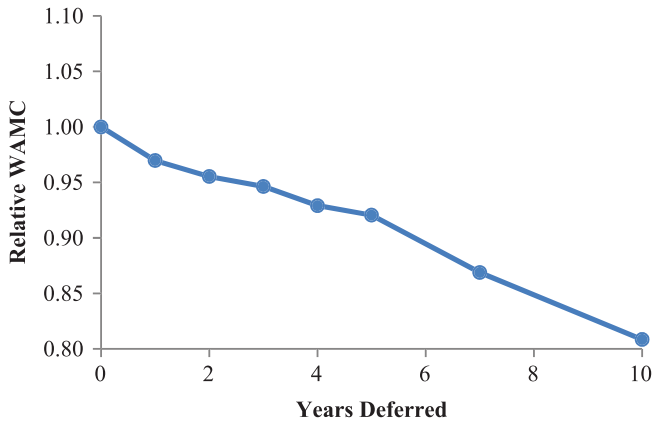

Panel C: U.S. Relative Avg Daily Volume (1947–2013)

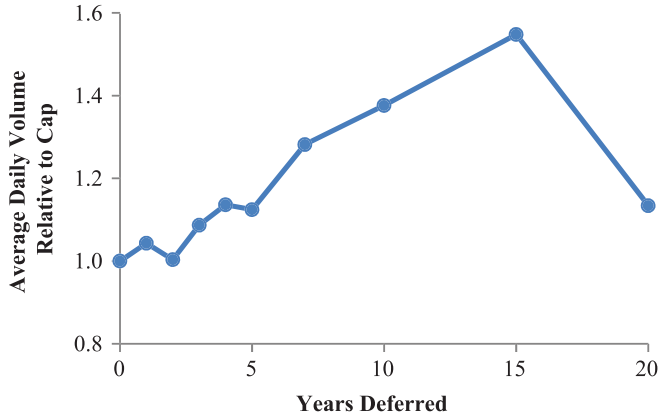

Panel D : Developed Relative Avg Daily Volume (1981–2013)

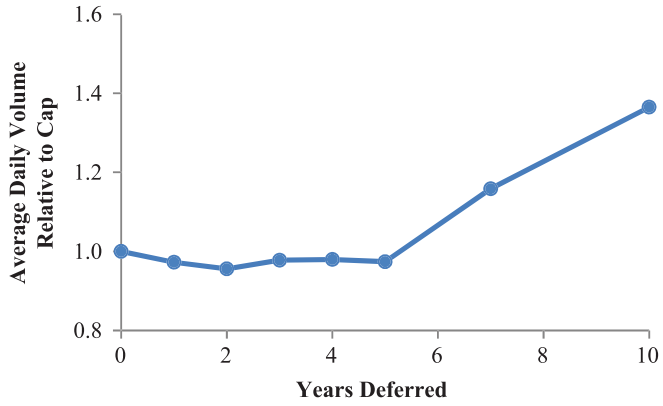

Panel E: U.S. Annual Turnover (1947–2013)

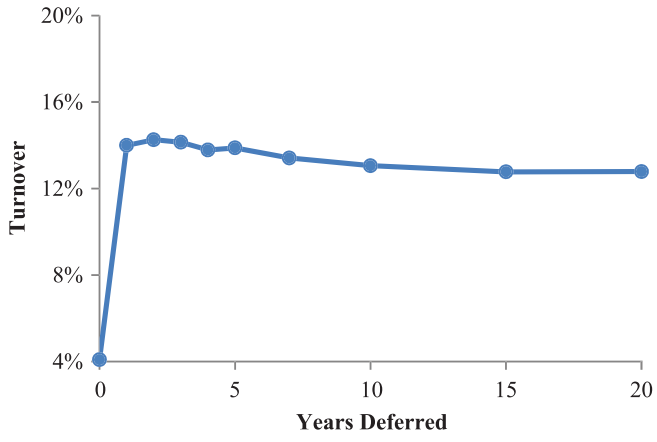

Panel F: Developed Annual Turnover (1981–2013)

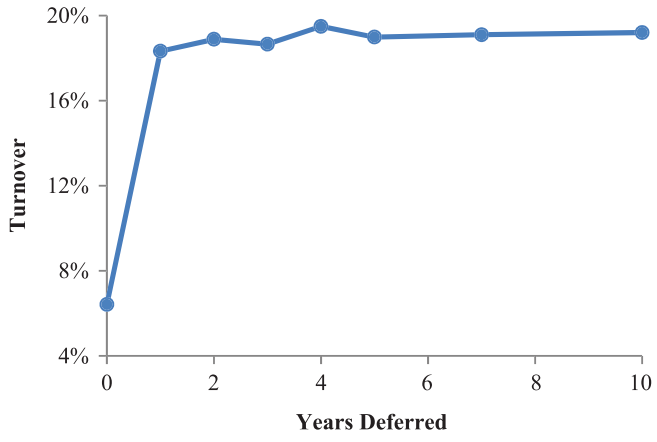

Source: Research Affiliates using data from CRSP, Compustat, Worldscope, and Datastream.

EXHIBIT 8

Term Structure of Value Added from Stale Cap Weighting

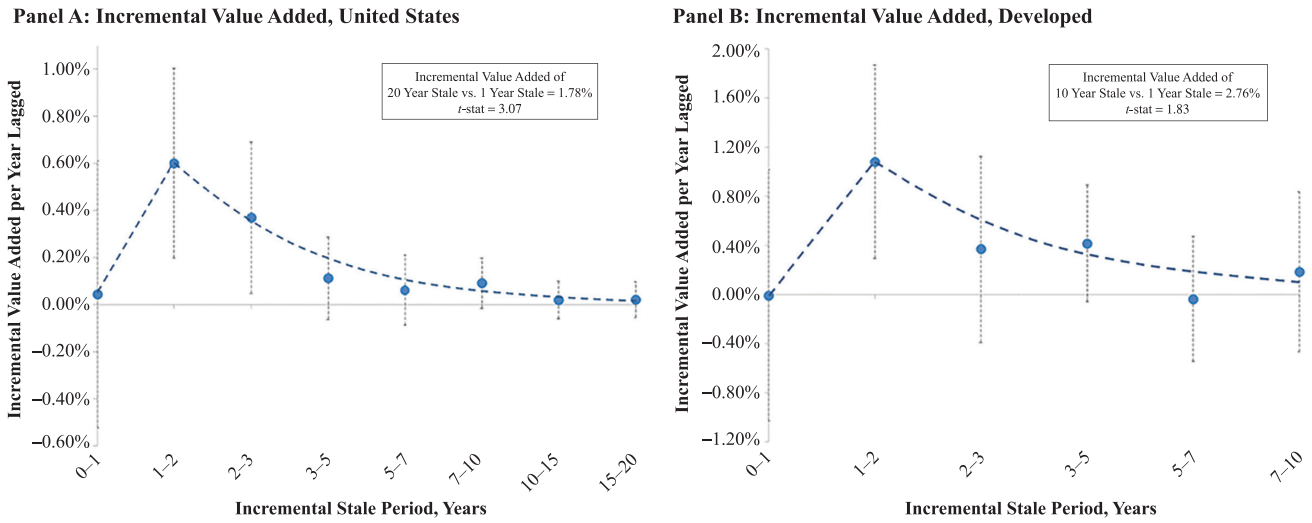

Source: Research Affiliates using data from CRSP, Compustat, Worldscope, and Datastream.

EXHIBIT 9

Cumulative Value-Added of One-Year Stale Index vs. Cap-Weighted Index

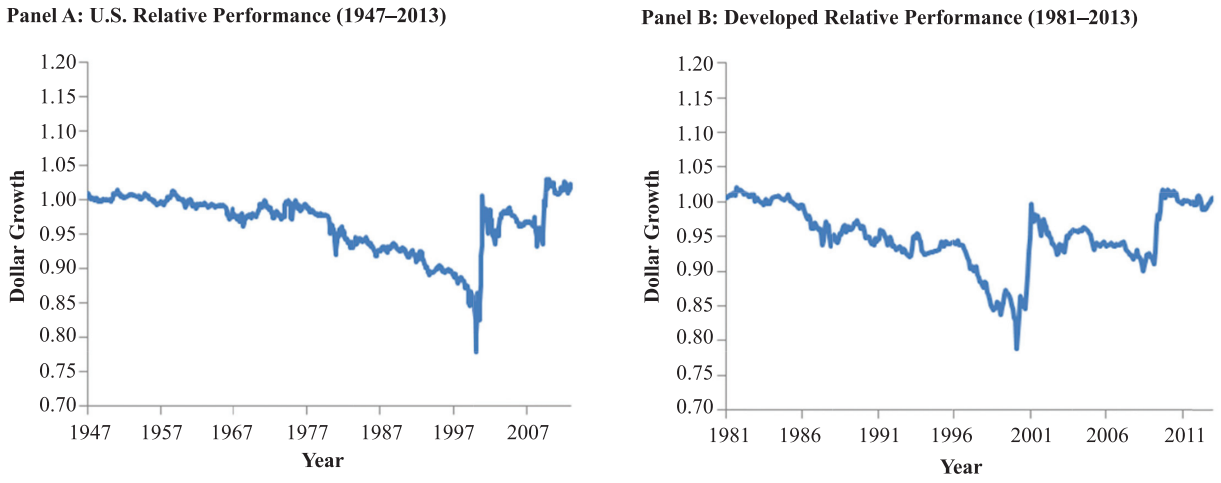

Source: Research Affiliates using data from CRSP, Compustat, Worldscope, and Datastream.

the momentum effect seems to exert its effect largely over spans of months (not years), we might expect the one-year stale cap-weighted index to hurt us. Instead, in both the United States and developed international markets, the one-year stale-weight index provides value added that is close to zero and lacks any statistical significance. The lack of underperformance is not so surprising, however, when we consider the well-documented tendency for serial correlation to be positive only over spans shorter than a year.

In Exhibit 9, we track the cumulative value added of the one-year stale index versus the regular cap-weighted index and observe that the stale portfolio actually underperforms most of the time. This result is consistent with our original expectation that negative momentum loading will hurt the one-year lagged stale

index. However, there are several episodes when the one-year stale index tends to win back all the previous years' accumulated underperformance. These episodes coincide with brief and severe momentum crashes, which have been amply documented in the literature.[12]

We know that the standard long–short momentum portfolio delivers positive returns. An important difference between the standard long–short portfolio and the active bets of the one-year stale-weight portfolio is the leverage. The long-short portfolio keeps constant leverage, while the active bets of the one-year stale-weight index relative to the benchmark vary over time. Active bets tend to increase in periods of heightened return dispersion and heightened volatility. Higher volatility has been demonstrated by several researchers, including Daniel and Moskowitz [2013], to negatively predict momentum premium. These rare momentum crashes coincide with periods of strong value factor performance and unambiguous market reversals.

Our hypothesis is that the one-year stale-weight portfolio seems to time these periods of momentum crashes—and superior value performance—just well enough not to be badly hurt by its anti-momentum character. As the result of this inherent momentum timing, the one-year stale index does not underperform, despite the strong negative momentum bet that it is taking.

When lagged for periods longer than one year, the stale-weight portfolios retain a diminishing negative momentum bet, but also start adding value by increasing bets on long-term mean reversion. The strongest incremental returns are observed in extending the lag from one year to two, or from two years to three, the very periods at which the literature documents the strongest mean reversion. That second year is a doozie; it has large reward and large statistical significance, with subsequent years of delay or staleness adding additional diminishing benefit with diminishing significance. But the benefits of Rip Van Winkle's lethargy continue to be economically powerful all the way out to 10 years.

Rip Van Winkle's investment style also gives us some insight into the standard cap-weighted portfolio. Because of its benchmark status, all other equity strategies are thought of as having factor tilts, with respect to the cap-weighted portfolio. By investing with cap weights from other points in time, however, we can perceive the tilts and tendencies of the current cap-weighted portfolio with respect to the broader opportunity set, and break away from the perspective where the current cap-weighted portfolio has—by definition—no style tilt. In other words, this work suggests that perhaps the cap-weighted market is making active bets—and taking on factor loadings—relative to its non-cap-weighted opportunity set,[13] favoring growth and momentum while shunning long-horizon mean reversion.

For those who are willing to view the world from a non-cap-weight-centric perspective, we can offer additional insights. We observed that all of Rip's stale-weight portfolios had a negative momentum tilt, and yet were able to overcome that disadvantage with their natural tendency to make a stronger anti-momentum bet (with respect to current cap weighting) at times when momentum was about to crash. In other words, the stale-weight portfolios all have a negative momentum tilt, but good momentum timing. From the opposite perspective, however, one could say that it is the current cap-weighted portfolio that has a positive momentum tilt, but poor momentum timing. Because the cap-weighted portfolio rewards recent (six- to twelve-month) winners by assigning them higher weights than before, it benefits from those stocks' tendency to outperform. However, that predilection for recent winners varies over time, getting stronger when cross-sectional return dispersion is high and weaker when stock returns are more clustered during less turbulent markets. When the market reverses direction (and momentum strategies crash), the cap-weighted portfolio's stronger-than-usual momentum bet causes it to give back any momentum-based gains it had amassed.

When considering the effect of long-term mean reversion, the results grow even bleaker for the cap-weighted index. We observe that Rip Van Winkle's stale portfolios benefit mightily from overweighting (relative to current cap weighting) the stocks that have come down in price since he fell asleep, and vice versa. From the opposite perspective, one could just as reasonably say that Rip was no market genius; it is the cap-weighted portfolio that underperforms due to overreliance on stocks that have risen most in price over the past one to twenty years. Overweighting those long-term winners (with respect to valuation-indifferent weighting strategies such as stale weights) means that the cap-weighted portfolio is making a bet against long-term mean reversion. This proves to be a very expensive bet, costing cap-weighted investors on the order of 180 basis points per year.

ISN'T THIS JUST THE VALUE EFFECT, ALL OVER AGAIN?

Of course it is, in part. But…

Why is it that an investment strategy based solely on stale information would outperform one weighted by current market capitalization? Stale cap weight contains less information than current cap weight. The answer should be linked to price—after all, we know that the stale-weight index differs from the current cap-weighted index primarily due to recent relative price changes. The longer we lag the cap-weighted index, the less biased we are toward today's expensive stocks. Of course, with lagged cap-weighted portfolios we are very biased toward yesteryear's expensive stocks; today, these may or may not carry lofty prices, market caps, or multiples.

To demonstrate the stale portfolios' relative cheapness, Exhibit 10 plots their price-to-book and price-to-dividend ratios, expressed as a multiple of the corresponding fundamental ratios of the cap-weighted index. The higher the ratio, the more expensive the portfolio. Not surprisingly, the market cap-weighted portfolio is the most expensive by both of these valuation measures. As the lag becomes longer, we own portfolios trading at successively lower multiples, with valuation discounts averaging 20% to 30% on the most

EXHIBIT 10

Price to Fundamental Ratios Relative to the Cap-Weighted Portfolio

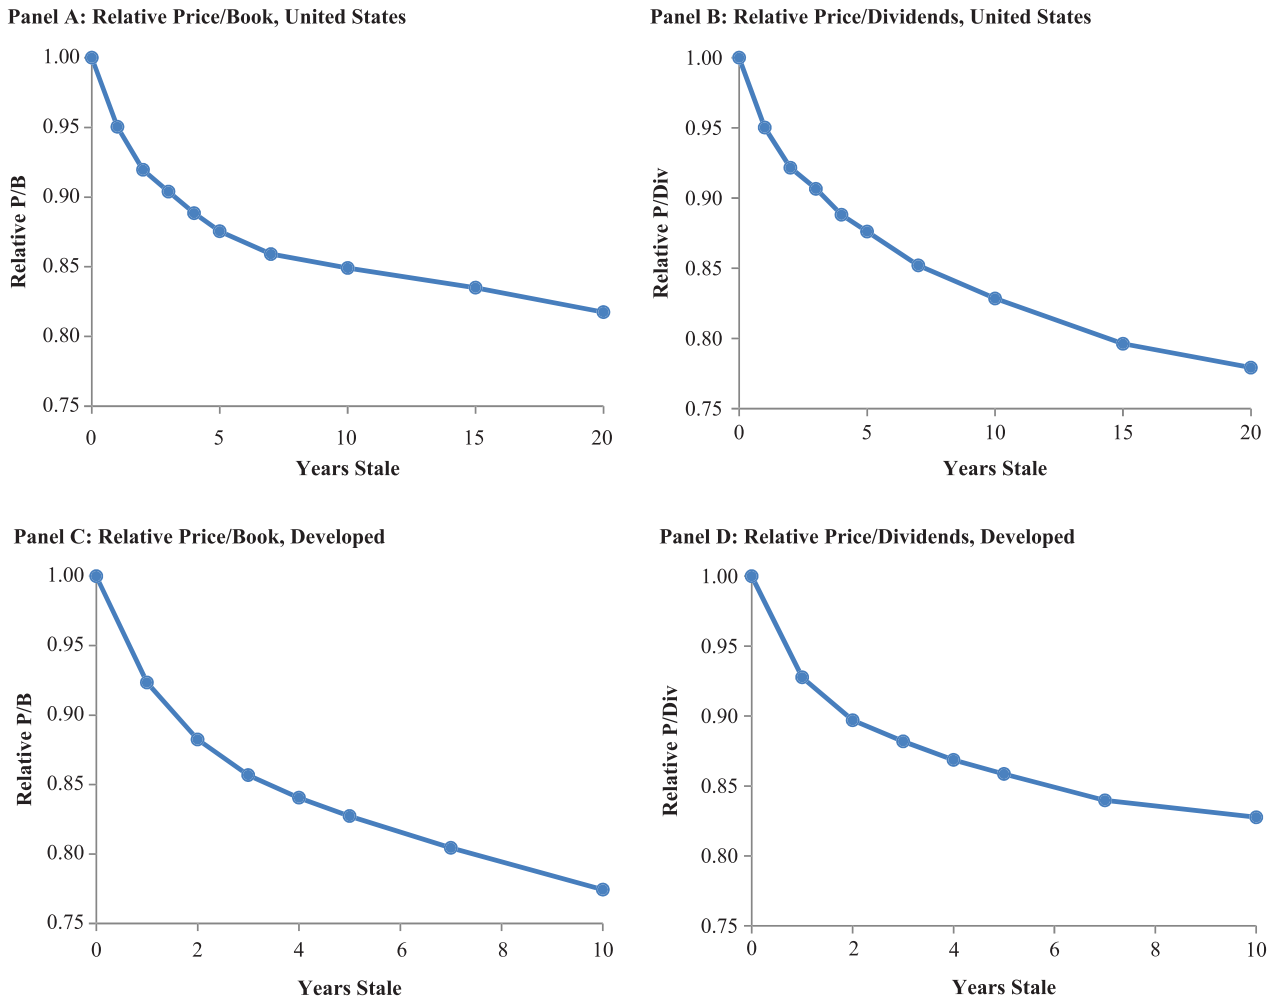

Source: Research Affiliates using data from CRSP, Compustat, Worldscope, and Datastream.

stale cap-weighted portfolios. This is consistent with the monotonically increasing HML factor we observed in Exhibit 5, which explains the increase in value added at longer lags.

We think it's very powerful to consider examining the capital markets not only from a cap-weighted market-centric perspective, but also from a macroeconomy-centric perspective. From the former vantage point, the cap-weighted markets are utterly neutral, by definition, and any departures from the cap-weighted market portfolio weights represent active bets. From a macro-economy perspective, the cap-weighted market is making obvious and sometimes large active bets, presuming (using a 2014 example) that Apple will be the largest source of risk-adjusted profits in the world, delivered to its shareholders in the decades ahead. Perhaps this will be true. But it is not yet true. From a macroeconomic perspective, the market is making an active bet on Apple today, relative to the much smaller current macroeconomic footprint that it occupies in the U.S. and global economy.

Viewed relative to its opportunity set, our use of stale price information to create portfolios is not inherently a growth or a value strategy. It is specifically using current (or very recent) prices that leads to a heavy reliance on stocks at the precise moment that they are most expensive. Viewed relative to the cap-weighted markets, we make an inherent value bet by choosing almost any method to sever the link between the current price and the weight in our portfolio. At the risk of voicing heresy, we believe that mispricing, not a hidden risk, is the value effect's key driver.[14]

RIP VAN WINKLE THROUGH THE LENS OF BEHAVIORAL FINANCE

The market has often been described as manic-depressive (although not so much in academic or indexing circles). Markets can get very optimistic—some would say manic—about a stock or sector. In March 2000, the market was ready to pay 182 times the then-current earnings for Cisco. The stock was priced to reflect a tacit market expectation that Cisco would deliver more risk-adjusted profit to its shareholders than any other company in the world in the decades to come. With the blessing of hindsight, it now seems the market was at best overreaching and perhaps even foolish. An even more extreme example was Nokia, priced at the time as if more than 70% of Finland's economic future rested on this cellphone company.

Conversely, the market can become deeply pessimistic—perhaps even depressive—about a stock or a sector. For instance, in March 2009 the market was willing to pay only 21% of the book value of Barclays Bank, pricing it for reasonably high odds of oblivion, perhaps even a "call option" on its very survival, even though many other British, American, and European banks were closer to the brink at that time. Again, in hindsight, the market might seem to have been unduly glum.

The market's bipolar character, first eloquently explored by Benjamin Graham [1949], means that any stock that eventually underperforms is cap-weighted above its eventual "fair value weight," and any stock that eventually outperforms is cap-weighted below its prospective weight. This means that the cap-weighted index suffers losses from overweighting overpriced stocks and underweighting underpriced stocks, relative to this unknowable fair-value-weighted portfolio. It's easy to dismiss this observation as an irrelevant tautology, because we cannot know the fair value and so cannot know the fair value weight for any investment.

This tautology matters.[15] It means that a cap-weighted index should exhibit a return drag relative to any strategy that can neutralize this over-reliance on future underperformers and under-reliance on future outperformers. Obviously, weighting a portfolio by the future cap weight would serve this purpose and win handily. Less obviously, stale portfolios also empirically serve this purpose rather nicely, as do other valuation-indifferent methods examined in many other articles in these pages in recent years and, less frequently, in past decades.

Behavioral finance accepts that stocks can be mispriced. Of course, to the adherents of efficient markets, this discussion about manias, depression, euphoria, pessimism, mispricing, and so forth will seem indistinguishable from gibberish. Even so, Rip Van Winkle's portfolios with stale cap weights provide us with an interesting opportunity to study the term structure of interaction between past mispricing and the subsequent return drag on cap-weighted index performance.

Let us re-examine Exhibit 8, showing the incremental value added as we extend the cap-weighted portfolios' lag period. The one-year stale-weight index provides very little value added. As we discussed, this would seem to be a manifestation of the well-docu-

mented momentum effect and the time-varying interaction between the momentum effect and the active bets of the one-year stale-weight index. An advocate of behavioral finance might suggest that, when the market feels euphoria (pessimism) about a certain stock, the feeling can build for some time, creating a momentum effect. As mean reversion kicks in, beyond the months-long span during which markets evince momentum, we begin to enjoy compensation for our willingness to favor yesteryear's favorites at today's prices.

Exhibit 9 seems to suggest that whether the stock becomes overpriced and overweighted or underpriced and underweighted in the cap-weighted index, short-term momentum favors the cap-weighted index in a normal year, but not reliably. There are outlier years with extreme volatility, when momentum crashes; these episodes significantly hurt the cap-weighted index performance relative to stale portfolios and other valuation-indifferent strategies. If we average the regular and outlier years in the short run, we neutralize the interaction between the past return and subsequent cap-weighted index performance. However, a year or two after the initial mispricing, when rational thought presumably gains sway, the stock prices begin to correct themselves; this is when stale momentum empirically hurts us the most.

This term structure analysis suggests that the return drag on the cap-weighted portfolio comes most strongly from the stocks that became over- or underpriced two or three years prior, but it does not stop there. The return drag from mispricing can be felt in the cap-weighted index 5, 10, or even 20 years after the mispricing occurred. That is why Rip Van Winkle's approach in sleeping for 20 years would work marvelously in investing—an overreliance on expensive stocks, which was harmful two decades ago, is now quite benign.

Advocates of behavioral finance would suggest that there is always a mania or a depression—a manifestation of crowd behavior—evident in some corners of the market. In clinical terms, our results suggest that each of these market manias and depressions, affecting any stock or sector, usually lasts for about a year, with subsequent remission. After an episode of mania or depression, it takes an average of one or two years for the market to begin to rationally seek fair value. But these remissions are incomplete, and the hangover from the mania or depression can be felt, albeit with ever-diminishing impact, for at least 20 years.[16]

CONCLUSION

Many believe that the cap-weighted index is, if not mean–variance optimal, a benchmark that's very hard to beat. Many believe that, in order to outperform, an investor needs superior analysis or timely access to better information. Even then, the belief is that the competitive nature of the investment industry and the stock market would quickly arbitrage away even these opportunities to earn superior returns.

We show that, contrary to these beliefs, if Rip Van Winkle settled in for a 20-year nap and, upon awakening, invested in the surviving companies in his 20-year-old cap-weighted index in proportion to their stale weights from 20 years ago, he would have enjoyed average historical outperformance of 180 basis points or more, depending on the geography (much more in emerging markets). This result plays out domestically, internationally, and in emerging markets, over spans up to two-thirds of a century.

Naturally, the success of Rip Van Winkle's strategy comes from its active bets, if we classically define any departures from the current cap-weighted market portfolio as active management. Stocks exhibit momentum in the short run and mean reversion in the long run. The cap-weighted index puts more of investors' money into securities that have increased in price the most (growth and momentum stocks) relative to other stocks, and when compared with any valuation-indifferent portfolio construction method. Rip Van Winkle's portfolio reduces our reliance on these stocks. Conversely, Rip Van Winkle's portfolio would hold more than cap-weighted indices in stocks whose prices have fallen to less extravagant levels, which gives Rip Van Winkle's portfolios an inherent value bias relative to today's cap-weighted portfolio.

Rip Van Winkle's strategy—or any other strategy which does not directly link the weight in our portfolio to the price—would favor cheaper securities. A Fama–French attribution suggests that a value bias is the largest source of Rip Van Winkle's superior performance. But Rip isn't trying to take on a value bias. He's using a full-fledged cap-weighted portfolio, just a few years behind the times. This invites a question that we've asked in other forms and forums, many times before. Does Rip Van Winkle have a value bias when he really doesn't care about growth or value, or does cap weighting have a growth bias, relative to its opportunity set? The ques-

tion matters. The answer depends on whether we are rooted in a cap-weight-centric worldview.

It is not our intent to propose 20-year-old stale cap weight as a sensible strategy. There are well-documented better ideas. It has been our intent to explore why this preposterously crude idea works—globally—so very well.

ENDNOTES

The authors would like to thank Philip Lawton, Katy Sherrerd, Jaynee Dudley, and Lillian Wu for their helpful suggestions and editorial assistance.

[1]Irving was a science fiction writer ahead of his time!

[2]Comparing this result of active and passive investments in the book *Common Sense on Mutual Funds* by John C. Bogle [2010] showed that among the 201 active mutual funds that they investigated, only 28 outperformed the benchmark by more than 200 basis points per annum and 52 outperformed the benchmark by 100 to 200 basis points. And these were the survivors! Rip Van Winkle would be in the top cohort of these 52 successful active funds and in the second decile of the surviving funds.

[3]Long-term mean reversion, documented by the studies of De Bondt and Thaler [1985] and Poterba and Summers [1988], means that winner stocks, which had the best performance in the past three to five years, tend to underperform in the future; while loser stocks, which had the worst performance in the past three to five years, tend to subsequently outperform. Jegadeesh and Titman [1993] studied shorter periods and demonstrated that most recent winner stocks, those which had the best performance in the most recent six to twelve months, tend to outperform loser stocks, which had the worst performance in the last six to twelve months. This is known as short-term momentum.

[4]The term "smart beta" was coined by the Towers Watson consultancy in the United Kingdom. When coming up with the term, they did not intend it to be derogatory toward the traditional beta. Instead they assumed that investors need to be smart in order to use these ideas correctly and so to fully benefit from the new betas hitting the market. Specifically, they argued that these stocks should be used as part of a core holding, to avoid the risk of trading into them after recent success and bailing out after recent disappointments. Their reciprocal term for cap weight was not "dumb beta," but "bulk beta," reflecting the fact that one could invest in these cap-weighted indexes in more or less unlimited scale. When we use the term "smart beta," we refer to rules-based indices that do not use market capitalization to weight securities.

[5]We are not using the S&P 500 Index or the Russell 1000 Index for our zero-lag cap-weighted index because those indices do not go back far enough for our purposes. The simulated returns of the index we constructed for this research are essentially indistinguishable over long periods from those of the S&P 500 or the Russell 1000. The rate of return differences since the launch of each are within 18 and 15 basis points, and the correlations are 0.995 and 0.999, respectively.

[6]We can't know what the 20-year clairvoyant portfolio for 1994 was until the end of 2014, and we can't know what the stale portfolio for 1946 was without data from 1925. As our data extend from year-end 1926 until year-end 2013, this constrains the overlapping time span to 1947 to 1994. That's still almost a half-century. If we want to consider only clairvoyant index funds, we can go back to 1927, but must stop in 1994. And if we want to consider only stale Rip Van Winkle index funds, we can go right up to our most recent year but cannot go back earlier than 1947. Nonetheless, two-thirds of a century seems a useful span for this sort of test.

[7]A cynic might correctly note that, if the financial crisis had destroyed the financial services industry, Rip's clients would suffer accordingly. We readily acknowledge this fact, while observing that counterfactuals and alternative realities are not a useful form of research, even though they are very popular in current discussions about the policy choices in the financial crisis.

[8]The corresponding attribution and performance graphs are available upon request.

[9]In Exhibit 5, we provide evidence that the residual alpha cannot be attributed to standard long-term mean reversion (taking effect over a span of years). The long-term mean reversion factor portfolio goes long (short) stocks with the worst (best) performance in the period from ten years prior to one year prior. We also found evidence (omitted due to space limitations but available upon request) that the residual alpha cannot be attributed to short-term mean reversion (taking effect in less than a month) documented by Jegadeesh [1990]. The short-term mean reversion factor portfolio goes long (short) stocks with the worst (best) performance in the prior month.

[10]It is our hypothesis that this missing factor may be some form of timing the existing factors that we discussed in the previous section. It may also be some idiosyncratic factor related to long-term mean reversion, which does not exhibit share price co-movement the way the growth/value effect does. However, then it would be difficult to isolate in conventional attribution or factor analysis methods. Let the hunt continue!

[11]We also considered a composite of 5-, 10-, 15-, and 20-year stale portfolios. Unsurprisingly, this has slightly nicer characteristics, with a bit more statistical significance than the single-span portfolios offer. More notably, it cuts the incremental turnover almost in half. Instead of contra-trading

against the latest price changes price changes from some years ago, the latter at least is averaged away to be less important.

[12]Momentum crashes, documented by Daniel and Moskowitz [2013] and Barroso and Santa-Clara [2013], are highly predictable phenomena that occur during market recoveries that are preceded by significant market crashes. When the formation period of a momentum portfolio coincides with a market crash, the momentum portfolio will be long (or overweight) the low-beta stocks that had fared the best in the downturn and short (or underweight) the high-beta stocks with the most negative prior-year return. When the market recovers, the prior year's losing, high-beta stocks tend to lead the recovery and the momentum investor experiences a momentum crash. The one-year stale portfolio, whose active weights are by definition inversely related to prior-year excess return, is an anti-momentum portfolio that loses to the market in most months and outperforms significantly during a momentum crash. Furthermore, the one-year stale portfolio takes on stronger active bets (and therefore stronger anti-momentum bias) following years of high dispersion in stock returns, such as market crash years. The result is a stronger anti-momentum bias in times just before a momentum crash and outsized outperformance during momentum crashes that, in our sample, almost perfectly negates the underperformance of the remaining months.

[13]Most factor work anchors on cap weight as neutral by definition, so the market cannot have any active tilts in any conventional sense. But, relative to the broad macroeconomy, it would today be hard to judge Apple as the world's largest company, or Cisco (ever so briefly) in 2000. It carries the world's largest market capitalization as a reflection of future growth expectations, not current economic footprint. We refer to these tilts as "relative to the opportunity set," because, whether that opportunity set is weighted by fundamentals, equally weighted, or weighted by any other means that ignores price, the cap-weighted market is indeed making a growth and momentum bet.

[14]See Kalesnik [2013], Chaves et al. [2013], and Arnott et al. (forthcoming).

[15]The series of articles on clairvoyant value (Arnott et al. [2009a and 2009b] and Arnott et al. [2013]) explores this idea in some detail from the perspective of a clairvoyant investor. What would an investor have been willing to pay, had she seen all future income distributions and the end-point share price? These three papers, respectively, document that 1) the market does a brilliant job of paying a premium for companies that subsequently deliver superior growth, with roughly a 50% correlation between the relative valuation paid for a stock and its relative NPV for anyone who could have seen the future; 2) the market then overpays for that superior growth, so that the correlation between premium valuation and subsequent IRR is roughly –50%; 3) the relative performance of growth and value stocks is highly and inversely correlated to the average premium that the market is willing to pay for growth relative to value; and 4) the much-vaunted size effect is an outlier effect, with the average small-cap stock performing a little worse than the average large-cap stock, but the few best small-cap stocks saving the day for the small-cap portfolio.

[16]Or more! We did not test longer spans, but the marginal efficacy clearly fades on these longer spans.

REFERENCES

Arnott, R., J. Hsu, J. Liu, and H. Markowitz. "Can Noise Create the Size and Value Effects?" *Management Science,* forthcoming.

Arnott, R., J. Hsu, V. Kalesnik, and P. Tindall. "The Surprising Alpha from Malkiel's Monkey and Upside-Down Strategies." *The Journal of Portfolio Management*, 39 (2013), pp. 91-105.

Arnott, R., J. Hsu, and P. Moore. "Fundamental Indexation." *Financial Analysts Journal*, 61 (2005), pp. 83-99.

Arnott, R., F. Li, and K. Sherrerd. "Clairvoyant Value and the Growth/Value Cycle." *The Journal of Portfolio Management*, 35 (2009a), pp. 142-157.

——. "Clairvoyant Value and the Value Effect." *The Journal of Portfolio Management*, 35 (2009b), pp. 12-26.

Arnott, R., F. Li, and G. Warren. "Clairvoyant Discount Rates." *The Journal of Portfolio Management*, 40 (2013), pp. 109-123.

Asness, C. "Variables that Explain Stock Returns." Ph.D. dissertation, University of Chicago, 1994.

Barroso, P., and P. Santa-Clara. "Momentum Has Its Moments." Working paper, 2013.

Bogle, J. *Common Sense on Mutual Funds*. Hoboken, New Jersey: John Wiley & Sons, 2010.

Carhart, M. "On Persistence of Mutual Fund Performance." *Journal of Finance,* 52 (1997), pp. 57-82.

Chaves, D., J. Hsu, V. Kalesnik, and Y. Shim. "What Drives the Value Premium? Risk versus Mispricing: Evidence from International Markets." *Journal of Investment Management*, 11 (2013), pp. 1-18.

Chen, C., R. Chen, and G. Bassett. "Fundamental Indexation via Smoothed Cap Weights." *Journal of Banking and Finance*, 31 (2007), pp. 3486-3502.

Chow, T., J. Hsu, V. Kalesnik, and B. Little. "A Survey of Alternative Equity Index Strategies." *Financial Analysts Journal*, 67 (2011), pp. 37-57.

Daniel, K., and T. Moskowitz. "Momentum Crashes." Unpublished manuscript, 2013.

De Bondt, W., and R. Thaler. "Does the Stock Market Overreact?" *Journal of Finance*, 40 (1985), pp. 793-805.

Graham, B. *The Intelligent Investor.* New York: Harper & Brothers, 1949.

Investment Company Institute. "2013 Investment Company Factbook." www.icifactbook.org/pdf/2013_factbook.pdf.

Jegadeesh, N. "Evidence of Predictable Behavior of Security Returns." *Journal of Finance*, 45 (1990), pp. 881-898.

Jegadeesh, N. and S. Titman. "Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency." *Journal of Finance*, 48 (1993), pp. 65-91.

Kalesnik, V. "Two Nobel Laureates…Two Tales of Value." Research Affiliates (2013). www.researchaffiliates.com/Production%20content%20library/F-2013-11-Two-Nobel-Laureates.pdf.

Poterba, J. and L. Summers. "Mean Reversion in Stock Prices: Evidence and Implications." *Journal of Financial Economics*, 22 (1988), pp. 27-59.

Sharpe, W. "Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk." *Journal of Finance*, 19 (1964), pp. 425-442.