Empirical Research Partners
Rochester Cahan
212 803-7973
Yu Bai
212 803-7919

August 3, 2017

Investment Ideas from the Ivory Tower August 2017

Big Data: Robot Newsreaders and Self-Driving Sectors

Charting the Wilderness Beyond the Last Row in Excel

- Since we launched our "Investment Ideas from the Ivory Tower" research series, almost four years ago now, we've used it as our scratch pad to test cutting-edge ideas gleaned from the latest academic literature. Think of it as our Triple-A affiliate in the Minors. Not all the ideas pan out, but our hope is that a few of the best eventually start slugging in the Big Leagues. Recent rookies that got the call-up include our work on hedge fund ownership and the steep equity yield curve, both of which have become integral parts of our starting lineup.
- In this edition, we turn our attention to Big Data, a topic so buzzworthy that even Smart Beta is getting jealous. In trawling through the academic research on the subject one thing stands out very clearly: most of the signals predict future stock returns at horizons of less than a month, and many only last for a day or even a tick. Of all the papers we read, less than a fifth were able to demonstrate alpha that persisted out to *investment* holding periods of a year or longer. However, for the handful of papers that did find longer-term efficacy the common theme was Natural Language Processing (NLP), the use of computer algorithms to interpret text-based content.

Hey, Robots Have Feelings Too

- We came across an interesting Fintech firm that's using NLP technology to comb through millions of news stories in real-time, looking for references to listed companies and other noteworthy economic topics like Fed commentary, non-farm payrolls, and chatter about corporate profits. When a story about one of the topics comes across the wires the robots assess the sentiment of the story in a few milliseconds, using machine learning algorithms to predict how likely it is that a human would interpret the story positively or negatively.
- Since saving a millisecond here or there doesn't matter for long-term investors we took the data and studied how it might be deployed in a fundamentally-driven, buy-and-hold portfolio instead. The first thing we noticed is that aggregate net sentiment, computed across the tens of thousands of stories processed each month, is a very close mirror image of our valuation spread metric. Spreads are widest when sentiment is dire.
- Computing net sentiment indexes for individual economic topics allows us to see directly and in real-time what's driving episodes of panic or euphoria. For example, when the valuation spread spiked to +1½ standard deviations in February of last year net sentiment indexes for the Fed, China, and commodities were all plunging at the same time. That's consistent with our view at the time that investors feared some combination of a Fed policy error that would stall the cycle prematurely, an implosion in China, and/or an emerging markets debt crisis triggered by collapsing commodity prices.

Self-Driving Sectors

- A second interesting use of NLP we came across is to use the text of 10-K filings to identify a company's natural peer group, by finding companies that have linguistically-similar business descriptions. Often the companies identified this way are different from those lumped into the same GICS industry. The academics found that a momentum signal derived from these text-based peers outperformed the standard own-firm momentum factor. We did our own work and found that on average that's true, but the advantage isn't consistent enough over time to make it a compelling addition to the toolbox.

The Best of the Rest

- Appendix 1 beginning on page 12 lists some of the other recent academic papers on Big Data that piqued our interest. It's a good summer reading list for those lazy days at the beach. Sun and productivity aren't mutually exclusive.

Sungsoo Yang (212) 803-7925 Nicole Price (212) 803-7935 Yi Liu (212) 803-7942 Yuntao Ji (212) 803-7920 Janai Haynes (212) 803-8005

Conclusions in Brief

- Most Big Data signals have short predictive horizons…

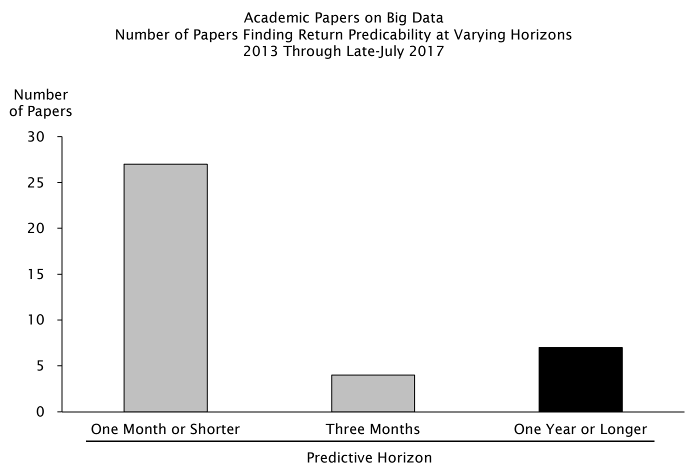

Source: Empirical Research Partners Analysis.

- …But text-mining shows some promise for long-dated alpha:

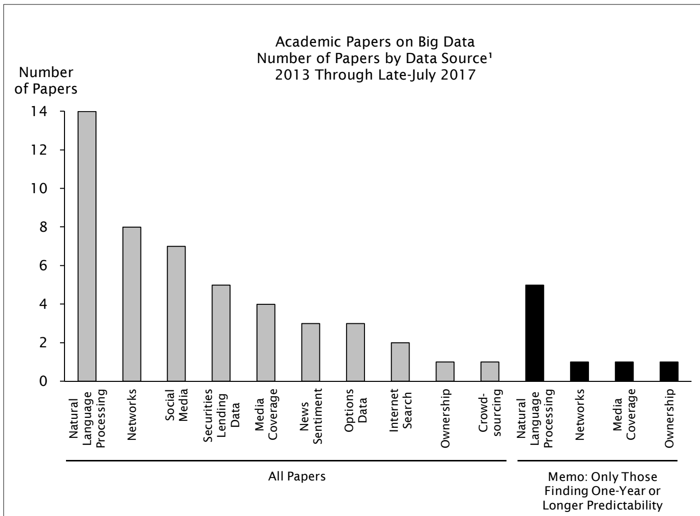

Source: Empirical Research Partners Analysis.
¹ Sums to more than total because papers can be classified in multiple categories.

- We studied real-time sentiment data scraped from news stories…

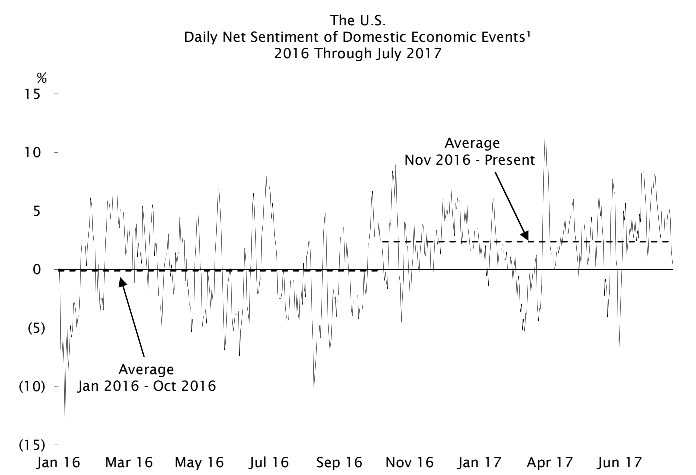

Source: Alexandria, Empirical Research Partners Analysis.
¹ Data smoothed on a trailing five-day basis.

- …And found that it mirrors the valuation spread:

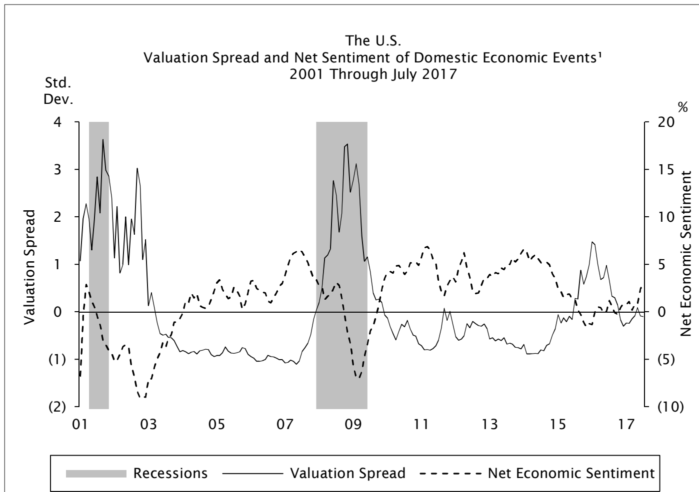

Source: Alexandria, National Bureau of Economic Research, Empirical Research Partners Analysis.
¹ Valuation spread measured as the cheapest quintile compared to the market average, based on the largest 1,500 stocks. Net sentiment is smoothed on a trailing three-month basis.

- China sentiment vis-à-vis the U.S. explained much of the panic early last year:

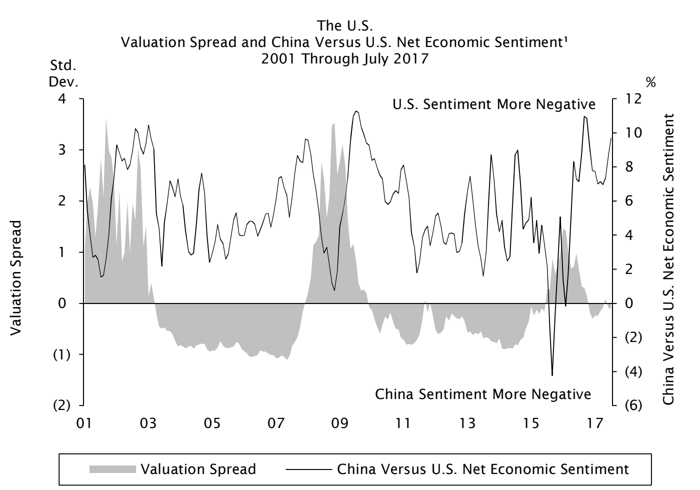

Source: Alexandria, Empirical Research Partners Analysis.
¹ Sentiment smoothed on a trailing three-month basis.

- Currently job-related sentiment is close to all-time highs:

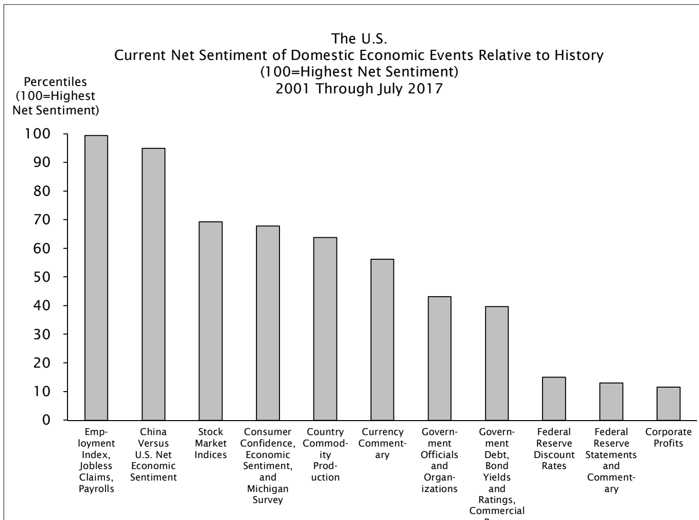

Source: Alexandria, Empirical Research Partners Analysis.

Big Data: Robot Newsreaders and Self-Driving Sectors

Charting the Wilderness Beyond the Last Row in Excel

Big Data has become the buzzword of the year, to the extent that Smart Beta is feeling quite jealous. Not to worry though, the first Big Data-driven Smart Beta ETF is surely rolling out to the launch pad as we speak. Like most things in our industry, separating hype from reality is half the challenge. Trading on Twitter sentiment sounds really cool, until you realize your portfolio needs to be completely refreshed every single day. On the other hand, burying your head in the sand and ignoring the march of technology usually ends with a mouthful of sand and not much else. In our work on the topic we've been looking for the happy median: which elements of the Big Data revolution actually add value in a real-world, fundamentally-driven investment process, a setting where long-term outcomes matter more than calling the next tick?[1]

The biggest practical hurdle to deploying Big Data is the short shelf life of most of the signals that have been discovered so far. We did a comprehensive review of the academic work that's been done on the topic and less than a fifth of the papers we read were able to demonstrate alpha that persisted out to investment holding periods of a year or longer (see Exhibit 1). The vast majority of signals predicted returns at intervals of less than a month, meaning one would need to constantly churn the portfolio to harvest them. Among the papers that did show some promise for long-term investors, the use of Natural Language Processing (NLP) was a common theme (see Exhibit 2). In a nutshell NLP uses computers to try to interpret text-based data automatically, often by training machine learning algorithms to recognized patterns in the words a human reader would associate with abstract quantities like sentiment, importance, and novelty.

Exhibit 1: Academic Papers on Big Data
Number of Papers Finding Return Predictability at Varying Horizons
2013 Through Late-July 2017

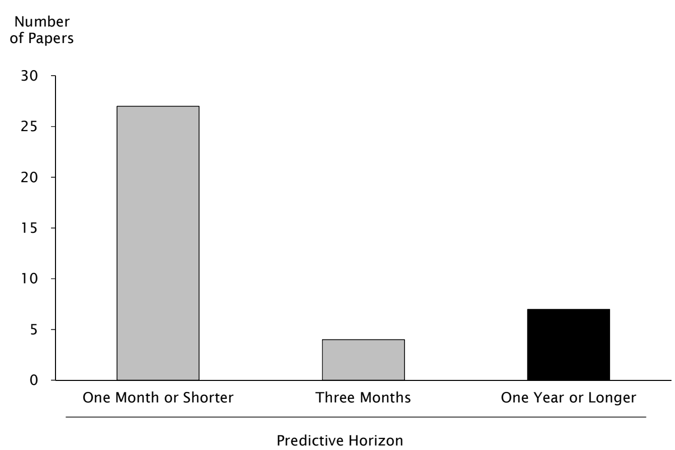

Source: Empirical Research Partners Analysis.

Exhibit 2: Academic Papers on Big Data
Number of Papers by Data Source[1]
2013 Through Late-July 2017

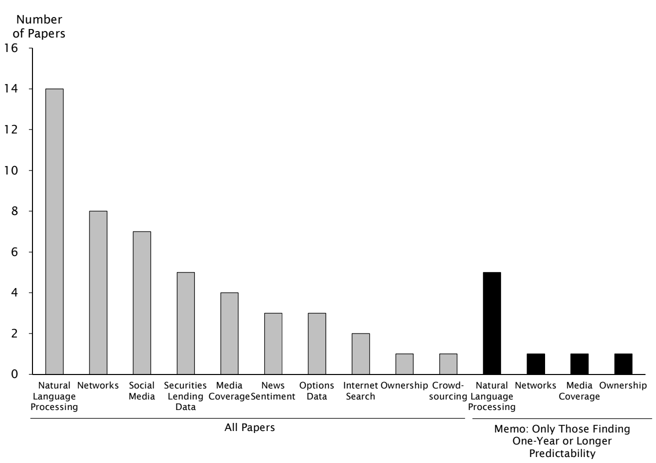

Source: Empirical Research Partners Analysis.

[1] Sums to more than total because papers can be classified in multiple categories.

Since NLP appears to be the most fruitful line of inquiry for long-term investors we did some work to see if we could find anything that might be additive to our investing toolbox. Luckily one of the two things we tried worked, so naturally we'll start with the win.

Hey, Robots Have Feelings Too

We recently came across an interesting Fintech firm called Alexandria that uses NLP technology to comb through millions of news stories in real-time, looking for references to listed companies and other important economic topics like non-farm payrolls, Fed commentary, and chatter about corporate profits (see Exhibit 3). Once the robots are confident they've found a story about one of the pre-defined topics they assess the sentiment of the story, meaning how likely is it that a human would interpret the news positively or negatively. All of this happens in something like 30 milliseconds.

[1] *Stock Selection: Research and Results* March 2017. "Big Data, Little Alpha?"

Exhibit 4 shows the monthly net sentiment of all stories about U.S. domestic economic events, computed over an average of around 25,000 stories published each month. It's worth noting that net economic sentiment, as extracted from the flow of news, isn't necessarily the same thing as an economic surprise index (see Exhibit 5). Sentiment is more like the *level* of expectations whereas surprises are the deviation from those expectations; for example it's possible and indeed common to have positive surprises when prevailing sentiment is negative, i.e., an outcome that's less-bad than already low expectations.

Exhibit 3: The U.S.
Top Ten Domestic Economic Events by Frequency
Twelve Months Ending July 2017

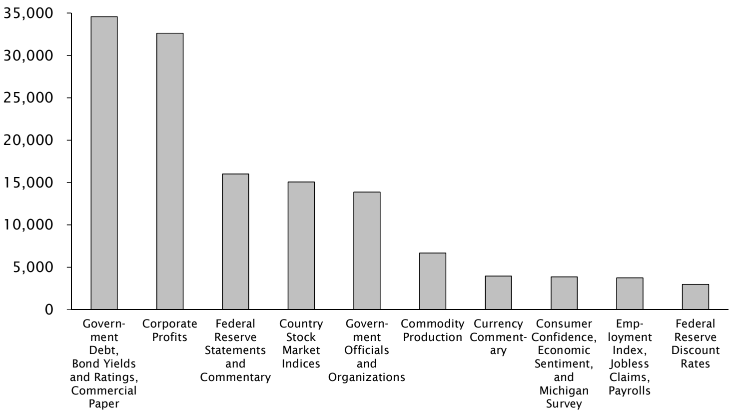

Source: Alexandria, Empirical Research Partners Analysis.

Exhibit 4: The U.S.
Net Sentiment of Domestic Economic Events
2001 Through July 2017

Source: Alexandria, National Bureau of Economic Research, Empirical Research Partners Analysis.

In our framework we've long used the valuation spread to capture the amount of fear baked into the market, with wide spreads signaling that investors are demanding a tremendous risk premium to hold whatever is looking scary at the time (see Exhibit 6). When spreads are wide we're being paid to take the other side of the market's fear, and usually that pays off because investors tend to overplay their panic when things are looking grim. It turns out net sentiment, as extracted from the voluminous flow of news stories, is almost the perfect mirror image of the valuation spread: when sentiment is dire spreads are wide and when it's sanguine spreads are narrow (see Exhibit 7).

In fact, there's historically been a perfectly monotonic relationship on average between net sentiment and the level of the valuation spread in the U.S., and the current spread is almost exactly where we'd expect it to be for the middle-quintile reading we have today (see Exhibit 8).

Exhibit 5: The U.S.
Economic Surprise Index and Net Sentiment of Domestic Economic Events[1]
2001 Through July 2017

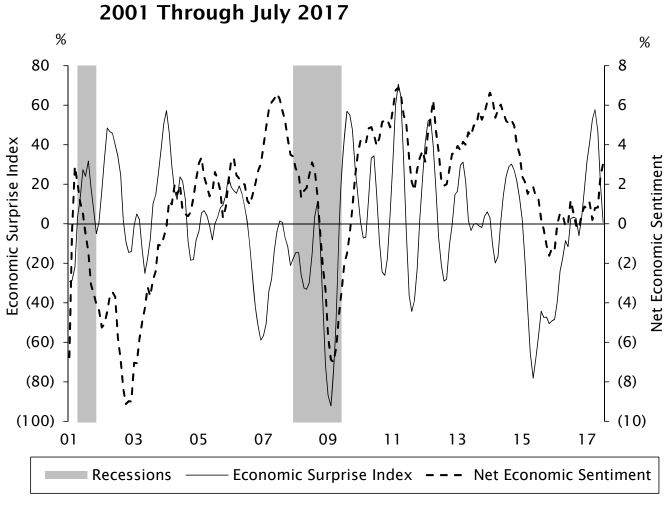

Source: Alexandria, Bloomberg L.P., National Bureau of Economic Research, Empirical Research Partners Analysis.

[1] Data smoothed on a trailing three-month basis.

Exhibit 6: U.S. Valuation Spreads
Expected Return of the Top Quintile Compared to the Average
1952 Through July 2017

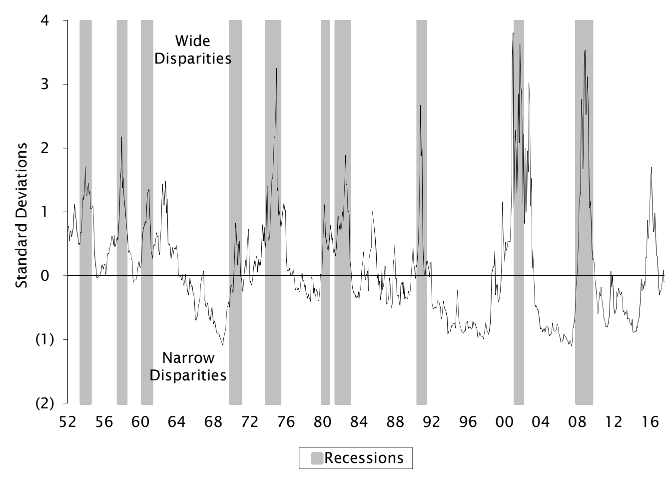

Source: National Bureau of Economic Research, Empirical Research Partners Analysis.

Exhibit 7: The U.S.
Valuation Spread and Net Sentiment of Domestic Economic Events[1]
2001 Through July 2017

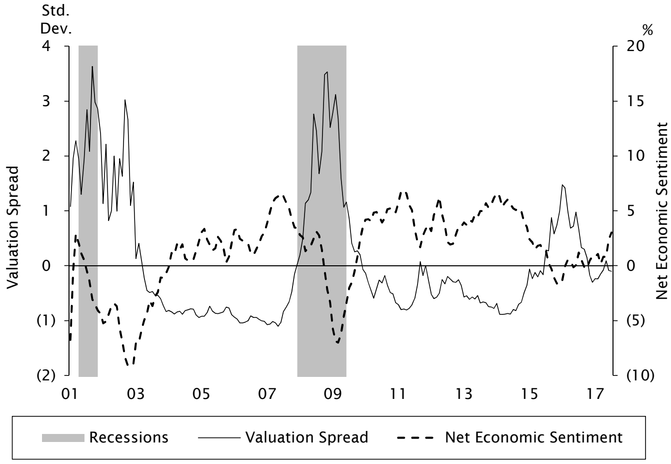

Source: Alexandria, National Bureau of Economic Research, Empirical Research Partners Analysis.

[1] Valuation spread measured as the cheapest quintile compared to the market average, based on the largest 1,500 stocks. Net sentiment is smoothed on a trailing three-month basis.

Exhibit 8: The U.S.
Average Valuation Spread by Quintile of Net Economic Sentiment
2001 Through July 2017

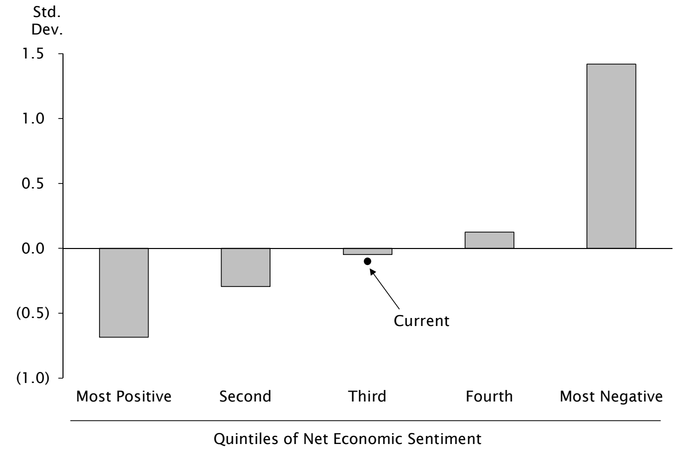

Source: Alexandria, Empirical Research Partners Analysis.

It's also interesting to overlay net sentiment with our Regime Indictor, which has proved itself adapt at pinpointing the prevailing mood of the market in live performance since 2007 (see Exhibit 9). As we'd expect, positive net sentiment tends to coincide with neutral-to-growth-driven regimes, and again the current level of sentiment is about what we'd expect for the growth-tilted regime we're in today (see Exhibit 10). Historically net sentiment tends to accelerate from a value-driven through a neutral regime and then starts to decline once the regime progresses to a growth-tilt and investors begin to fret about the impending end of the cycle (see Exhibit 11).

Exhibit 9: The U.S.
The Regime Indicator and Net Sentiment of Domestic Economic Events
2001 Through July 2017

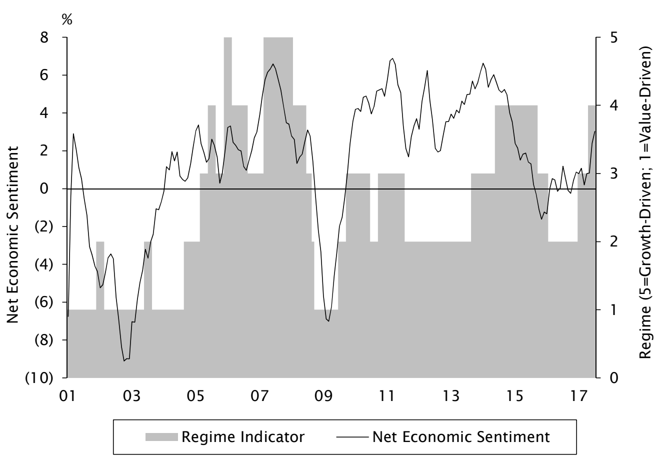

Source: Alexandria, Empirical Research Partners Analysis.

[1] Sentiment data smoothed on a trailing three-month basis.

Exhibit 10: The U.S.
Average Net Economic Sentiment by Regime
2001 Through July 2017

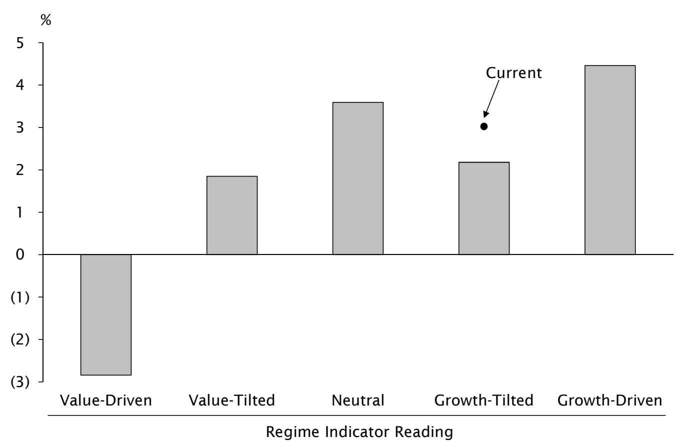

Source: Alexandria, Empirical Research Partners Analysis.

Be Greedy When Others are Fearful

The most obvious question is whether the level of net economic sentiment tells us anything about what might happen to equity market returns in the future. It turns out it does, and it's no surprise that Warren Buffet was right: we should buy when others are fearful (see Exhibit 12). In the chart we looked at the nominal returns to large-capitalization stocks over the following year, contingent on the starting level of net sentiment. As shown in the black bars, year-ahead returns have been twice as large on average when starting from negative net sentiment versus positive. Breaking things down more granularly, the grey bars show the future performance of the market based on quintiles of net sentiment. Today we're in the third quintile of the historical net sentiment distribution.

Exhibit 11: The U.S.
Average Monthly Change in Net Economic Sentiment
by Regime
2001 Through July 2017

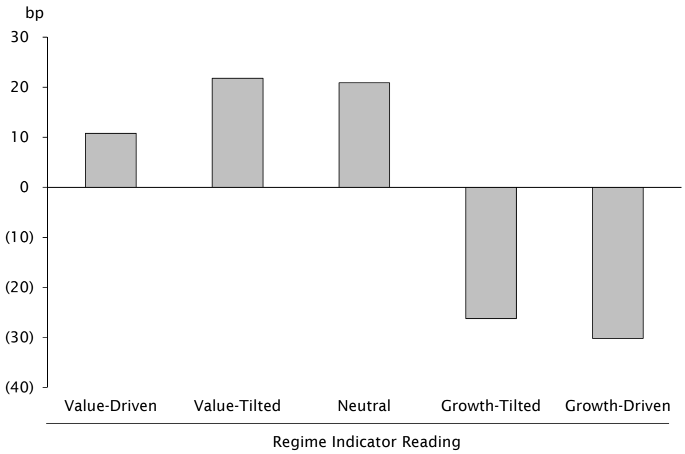

Source: Alexandria, Empirical Research Partners Analysis.

Exhibit 12: Large-Capitalization Stocks
Nominal Returns in Following 12-Months
by Direction and Quintile of Net Economic Sentiment
Monthly Data Compounded to Annual Periods
2001 Through July 2017

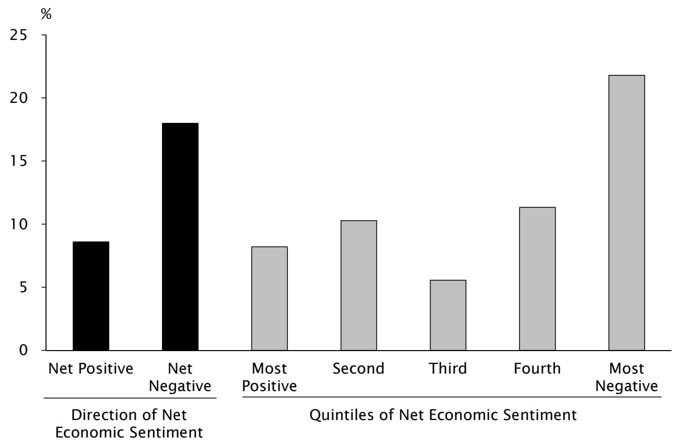

Source: Alexandria, Empirical Research Partners Analysis.

We repeated the analysis for value stocks, high momentum stocks, and Big Growers (see Exhibits 13 through 15). For value investing the starting point is critical and the best future returns come when the trade is a gut-wrenching contrarian bet at odds with the dismal prevailing mood. The value premium is earned with white knuckles. On the flip-side, momentum stocks and Big Growers, an elite cadre of stocks with the very best all-around growth credentials in the market, need a sentiment tailwind to fill their sails. Here positive sentiment is a sign investors are comfortable enough to extrapolate past trends in stock prices and cash flow growth into the future, without the gnawing fear that something will go wrong.

It's telling that daily net sentiment shifted positive after the election last year, and after an initial pop in the value-orientated cyclicals it's been the growth stocks that have mostly led since then (see Exhibit 16). We think real-time net sentiment will be a useful additional tool to help us better model regimes going forward, particularly in a world where factors are now tradeable and their fates can swing wildly on a whim.

Exhibit 13: Large-Capitalization Stocks
Relative Returns to the Best Quintile of Valuation
by Direction and Quintile of Net Economic Sentiment
Measured Over One-Year Holding Periods
2001 Through July 2017

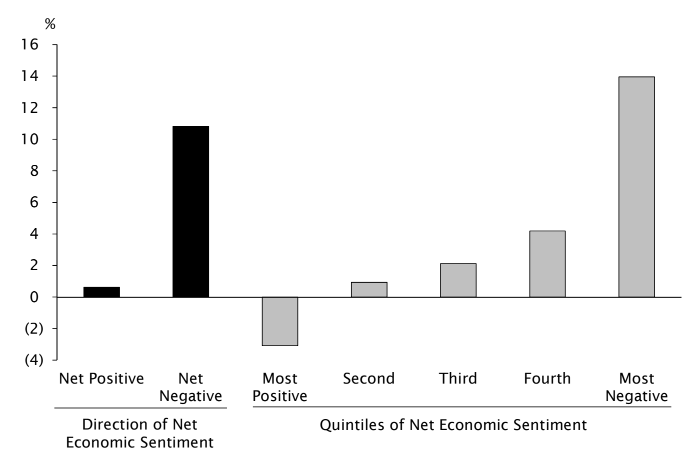

Source: Alexandria, Empirical Research Partners Analysis.

Exhibit 14: Large-Capitalization Stocks
Relative Returns to the Best Quintile of
Nine-Month Price Trends
by Direction and Quintile of Net Economic Sentiment
Measured Over One-Year Holding Periods
2001 Through July 2017

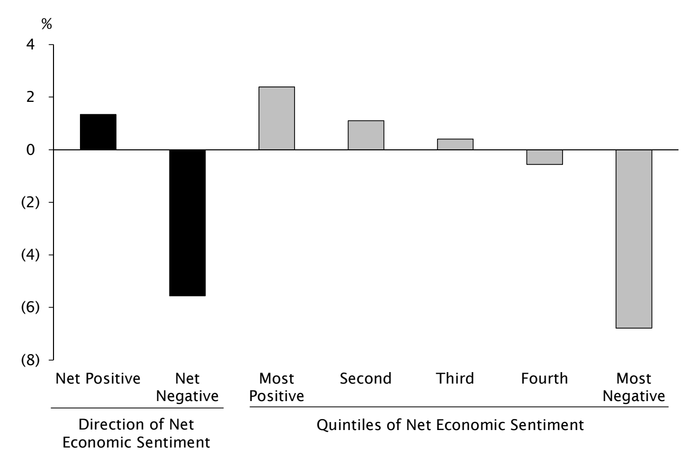

Source: Alexandria, Empirical Research Partners Analysis.

Exhibit 15: Big Growers
Relative Returns in Following 12-Months by Direction and Quintile of Net Economic Sentiment Monthly Data Compounded to Annual Periods 2001 Through July 2017

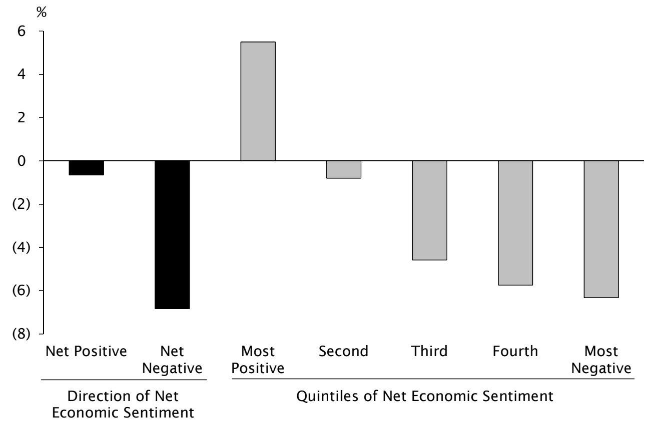

Source: Alexandria, Empirical Research Partners Analysis.

Exhibit 16: The U.S.
Daily Net Sentiment of Domestic Economic Events[1] 2016 Through July 2017

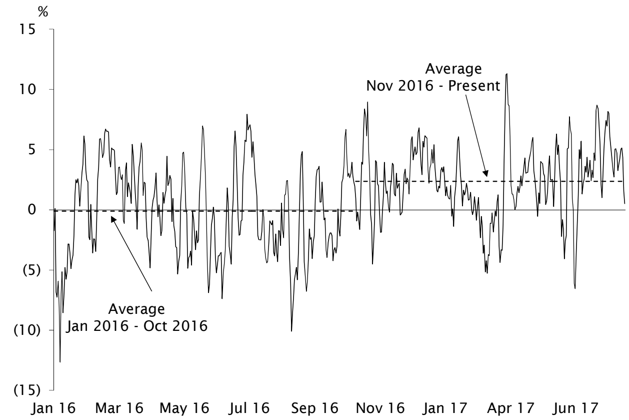

Source: Alexandria, Empirical Research Partners Analysis.
[1] Data smoothed on a trailing five-day basis.

The Why of Things

One nifty feature of Alexandria's data is that news stories are automatically tagged to specific "economic event" categories, as shown back in Exhibit 3. That means we can compute net sentiment indexes for each concept, giving us a real-time measure of the ebb and flow of the prevailing mood towards each. For example, in Exhibit 17 we plotted separate net sentiment indexes based on the news flow about the U.S. and China. There's a clearly a common component, the series are 68% correlated, but there are also times when they are quite different; that's easier to see if we plot the difference between the two series (see Exhibit 18). Currently the net sentiment extracted from stories about China is near the top-end of its historical range versus the U.S.

Exhibit 17: China and the U.S.
Net Sentiment of Economic Events[1] 2001 Through July 2017

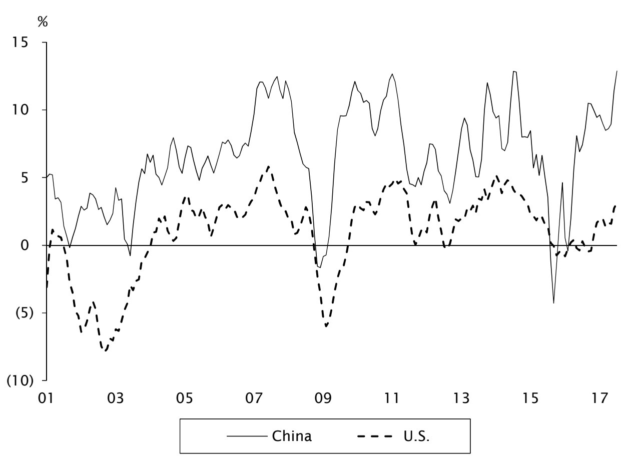

Source: Alexandria, Empirical Research Partners Analysis.
[1] Data smoothed on a trailing three-month basis.

Exhibit 18: China
Net Sentiment of Economic Events Relative to that of the U.S.[1] 2001 Through July 2017

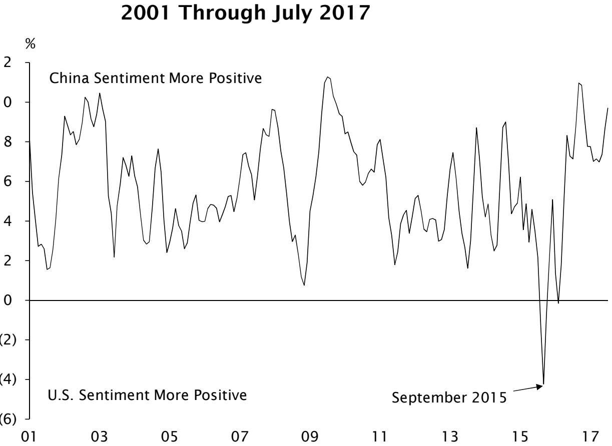

Source: Alexandria, Empirical Research Partners Analysis.
[1] Data smoothed on a trailing three-month basis.

The big dive in China-versus-U.S. sentiment in the third quarter of 2015 is also worth a closer look because it coincided with a sharp rise in the valuation spread that ultimately peaked at around +1½ standard deviations in mid-February last year, the highest-ever reading without a domestic U.S. recession (see Exhibit 19). At the time we argued that fear of an implosion in China was at the root of investors' rising trepidation, a conclusion that's supported by the net sentiment data from that period.[2]

[2] *Portfolio Strategy* August 2015. "China: The Dark Side of Bretton Woods II."

In addition to China we thought there were two other concerns contributing to the spike in spreads that peaked in mid-February last year: fear that the Fed would make a monumental policy error, stalling the cycle in its tracks, and alarm at the unfolding commodities bust and the risk it might prove contagious, perhaps via the large amount of Dollar debt tied to suddenly-worthless commodity assets in the emerging markets. Exhibit 20 zooms in on the turbulent period that began in 2014 and shows how net sentiment towards Fed policy and commodities evolved, in addition to the China sentiment series we've already seen. It was only when all three showed some signs of turning that spreads began to come back down to earth.

Exhibit 19: The U.S. Valuation Spread and China Versus U.S. Net Economic Sentiment[1] 2001 Through July 2017

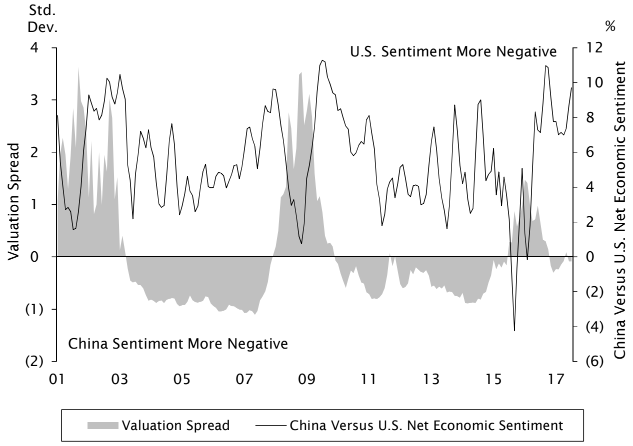

Source: Alexandria, Empirical Research Partners Analysis.

[1] Sentiment smoothed on a trailing three-month basis.

Exhibit 20: The U.S. Valuation Spread and Select Net Sentiment Indices[1] 2014 Through July 2017

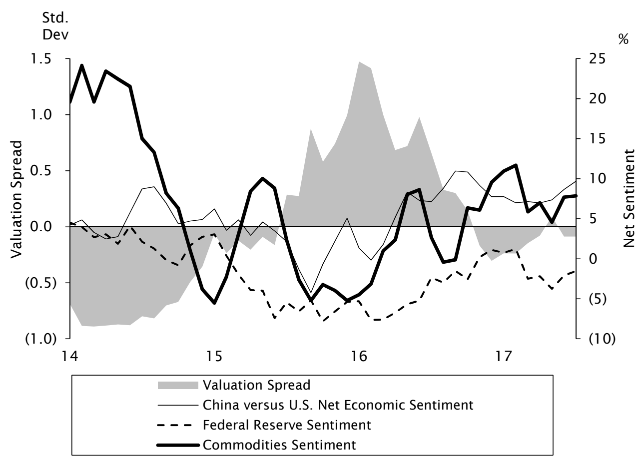

Source: Alexandria, Empirical Research Partners Analysis.

[1] Sentiment smoothed on a trailing three-month basis.

Over the long-run the net sentiment for the ten most-frequent economic events plus China-versus-U.S. net sentiment have explained almost all of the variation in the valuation spread (see Exhibit 21). Practically speaking that means when spreads are widening we can get a fairly good decomposition of what's driving them by looking at the individual sentiment indexes. Since the start of our data the sentiment surrounding corporate profits that has been most powerful in explaining the spread: a one standard deviation increase in optimism towards profits has been consistent with a contraction in the valuation spread of about 6/10ths of a standard deviation (see Exhibit 22). Exhibit 23 shows the history of the net sentiment surrounding corporate profits; it started to decline in mid-2014 as the oil price slid and the Dollar skyrocketed, and then continued to trend down until bottoming at its lowest-ever reading in January of this year.

Exhibit 21: The U.S. Actual and Estimated Valuation Spread[1] 2001 Through July 2017

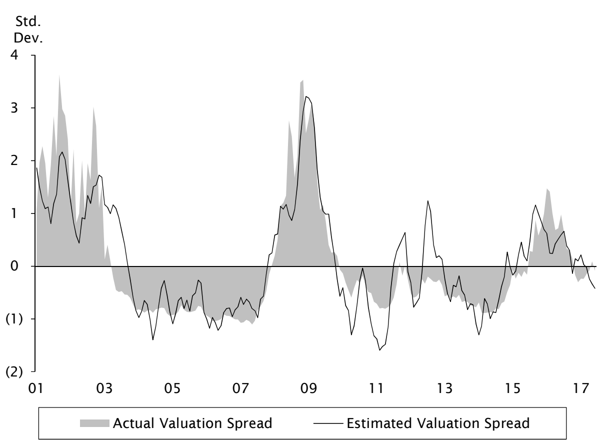

Source: Alexandria, Empirical Research Partners Analysis.

[1] Predicted spread based on regression model using net sentiment for the 10 most common economic events and China versus U.S. net sentiment as explanatory variables.

Exhibit 22: The U.S. Predicted Change in Valuation Spread for a One Standard Deviation Increase in Net Sentiment[1] 2001 Through July 2017

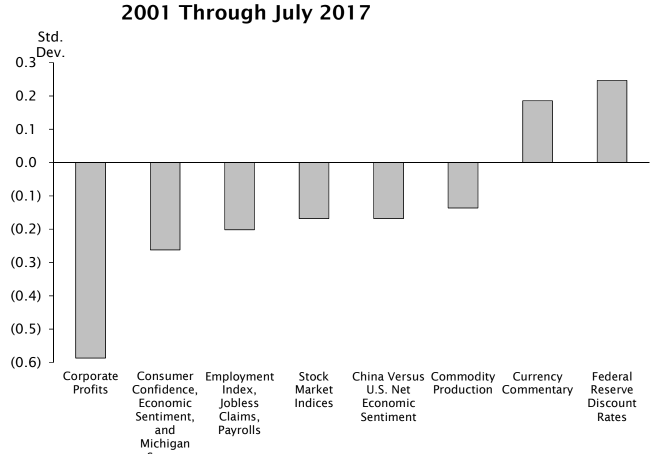

Source: Alexandria, Empirical Research Partners Analysis.

[1] All factors statistically-significant at the 5% level.

Exhibit 23: The U.S.
Valuation Spread and Net Corporate Profit Sentiment[1]
2001 Through July 2017

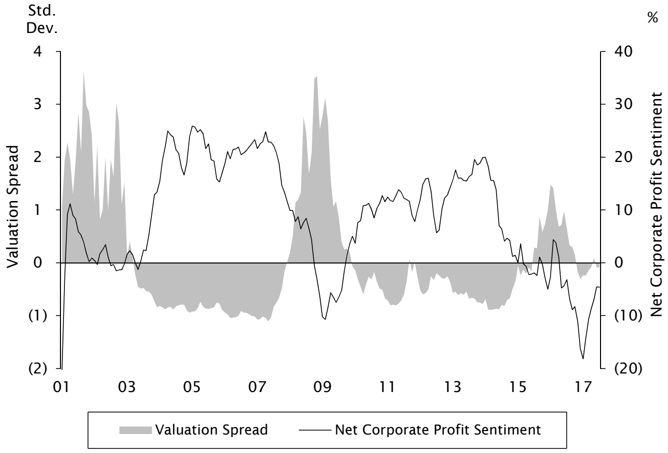

Source: Alexandria, Empirical Research Partners Analysis.
[1] Sentiment smoothed on a trailing three-month basis.

Exhibit 24: The U.S.
Predicted Change in Valuation Spread for a One Standard Deviation Increase in Net Sentiment[1]
2009 Through July 2017

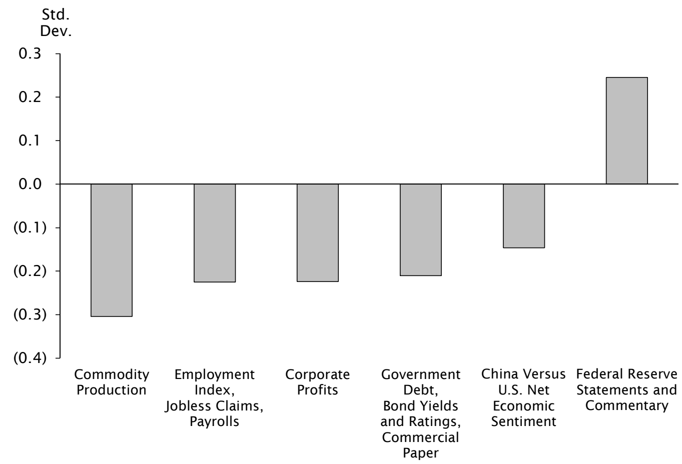

Source: Alexandria, Empirical Research Partners Analysis.
[1] All factors statistically-significant at the 5% level.

In the post-Crisis era the picture has been a little different, with sentiment towards commodities and job-related metrics mattering more than they have over the long-run (see Exhibit 24). The commodities sentiment index in particular has been on a topsy-turvy ride, following oil off the cliff in 2014 before staging a more sustained comeback in the past year (see Exhibit 25).

We took a look at how the current reading of each economic sentiment index compares to what we've seen historically (see Exhibit 26). In each case a reading of 100 would mean net sentiment is the highest it's ever been for that category. Job-related sentiment in particular is noteworthy because its current reading is close to the highest ever recorded (see Exhibit 27). Meanwhile the net sentiment associated with stories about consumer confidence has slipped recently, but are still around the 70th percentile of history (see Exhibit 28).

Exhibit 25: The U.S.
Valuation Spread and Net Commodities Sentiment[1]
2001 Through July 2017

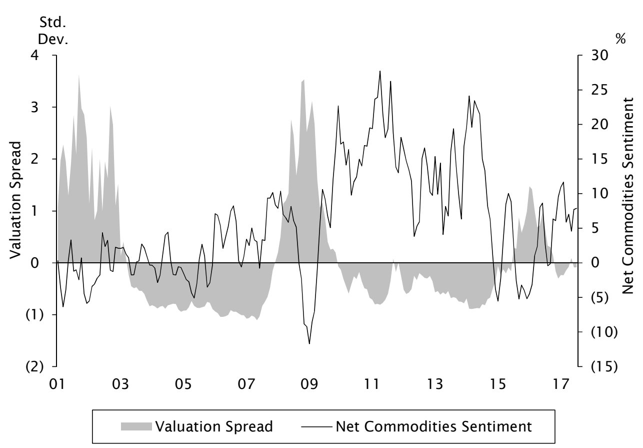

Source: Alexandria, Empirical Research Partners Analysis.
[1] Sentiment smoothed on a trailing three-month basis.

Exhibit 26: The U.S.
Current Net Sentiment of Domestic Economic Events Relative to History
(100=Highest Net Sentiment)
2001 Through July 2017

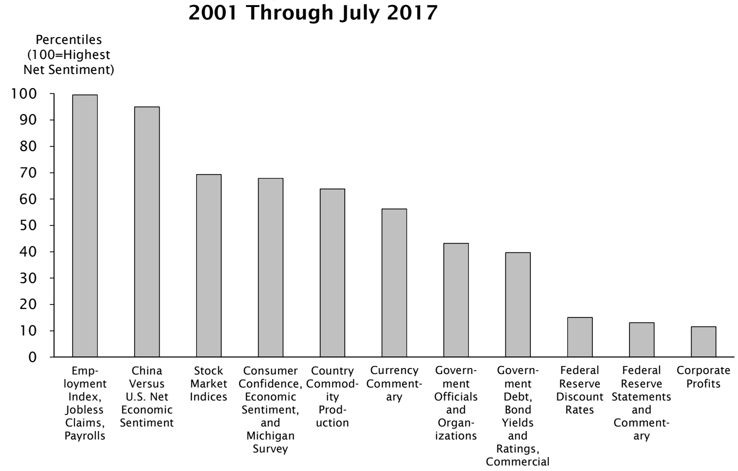

Source: Alexandria, Empirical Research Partners Analysis.

One of the popular narratives these days is that Big Data and machine learning in tandem will prove so powerful that *why* things happen won't matter much anymore. In a globalized world there are so many moving parts and obscure relationships that humans can't hope to understand them all, rather the robots will sniff them out and exploit them even if they can't translate their ideas into dummy-speak for their human overlords. Maybe, but we tend to think that the better use of Big Data will ultimately be to better understand the why of things. For example, when there's blood in the water and valuation spreads are exploding, what real-time fears are boiling to the surface? The machines can help us answer that by combing through millions of news articles in milliseconds, but it doesn't mean they can tell us what to do about it.

Exhibit 27: The U.S.
Net Sentiment Towards the Employment Index, Jobless Claims, and Payrolls[1]
2001 Through July 2017

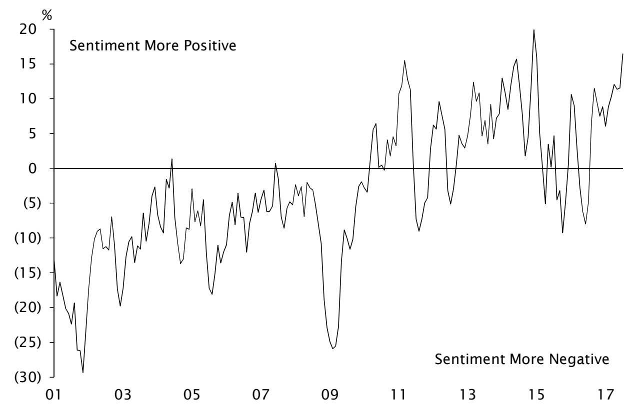

Source: Alexandria, Empirical Research Partners Analysis.
[1] Data smoothed on a trailing three-month basis.

Exhibit 28: The U.S.
Net Sentiment Towards Consumer Confidence, Economic Sentiment, and the Michigan Survey[1]
2001 Through July 2017

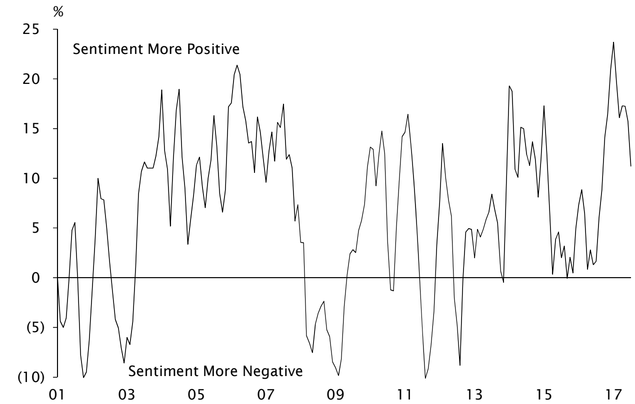

Source: Alexandria, Empirical Research Partners Analysis.
[1] Data smoothed on a trailing three-month basis.

Self-Driving Sectors

A second use of Natural Language Processing (NLP) that we came across recently used the text of 10-K filings to identify companies with "similar" business operations, where similarity in this case is based on how linguistically similar the business description section of one firm's 10-K is with another.[3] The idea is that grouping companies based on what they actually do, as reported to the SEC in their 10-Ks, might be better than relying on artificial industry classifications, like the popular GICS taxonomy.

For some companies the usual GICS industries (i.e., GICS Level 3) give a pretty good read on what they do. For example, consider United Airlines (see Exhibit 29). Obviously the company is an airline and therefore falls into the GICS Airlines industry (203020). It turns out the using the academics' methodology to define a text-based peer group gets to mostly the same place, in the sense that all the obvious airlines that appear in the GICS industry are also picked up by the text-based algorithm, see the grey bars in the chart. Where it gets a little more interesting is that the text-based algorithm also capturers air cargo carriers that compete with United for air freight; on a strictly GICS-basis the likes of FedEx fall into the Air Freight & Logistics (203010) industry. Still, one doesn't really need Big Data to come up with this fairly obvious list of competitors and near-competitors.

In contrast, a company like Apple is quite a different beast, operating in many different business lines (see Exhibit 30). Here the disparity between its GICS industry of Technology Hardware, Storage & Peripherals (452020) and its text-based peers is much larger. In fact, the only peer of note that appears in both Apple's GICS industry and is tagged as a text-based peer is Western Digital, see the grey bars. The white bars show that many text-based peers, like Microsoft for example, don't appear in Apple's GICs industry while the black bars show that there are also lots of companies in Apple's GICS industry that don't appear all that similar based on the company-disclosed descriptions of what they actually do.

The whole point of the exercise was to study whether text-based peers can be used to form a better momentum signal. The idea is simple: rather than looking at a firm's own past performance, perhaps it works better to base the momentum signal on the past performance of its peers instead, in the hopes that if the peers are rallying the firm itself will eventually join the party. Of course, that raises the question of which firms are legitimately peers. In past research, academics have found some evidence that industry-based momentum performs better than own-firm momentum, so the hypothesis behind using text-based peers is that it might do better still, if the NLP algorithms are doing a good job of finding firms that are more similar than what rigid industry taxonomies might suggest.

[3] Hoberg, G. and Gordon Phillips, 2017. "Text-Based Industry Momentum." Working Paper.

Exhibit 29: United Continental Holdings
GICS Industry- and Text-Based Peers
2016 Sales
As of Late-July 2017

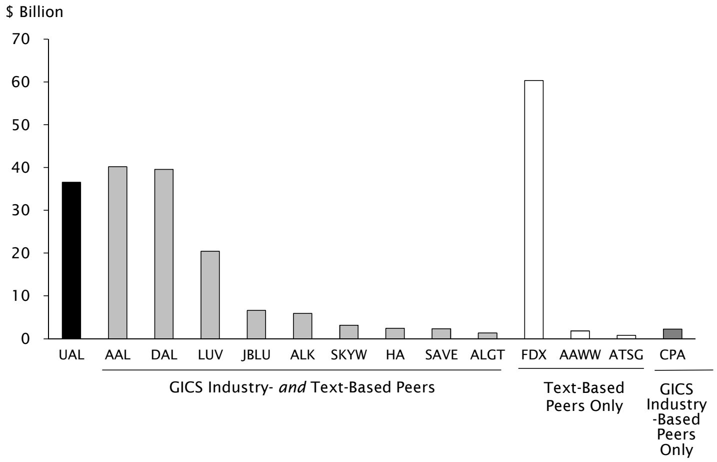

Source: Empirical Research Partners Analysis.

Exhibit 30: Apple
GICS Industry- and Text-Based Peers
2016 Sales
As of Late-July 2017

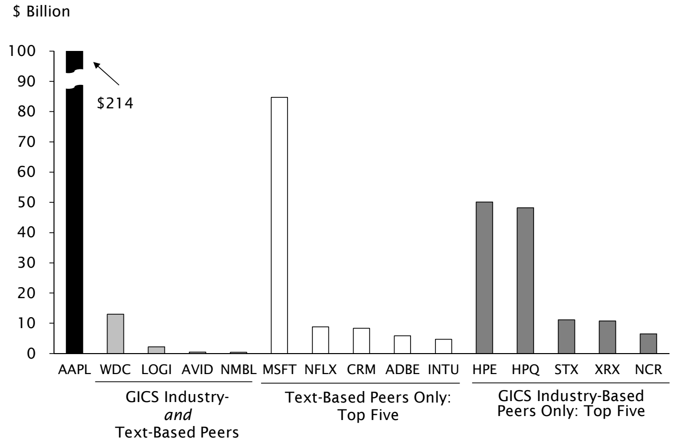

Source: Empirical Research Partners Analysis.

We tested three momentum signals side-by-side to see if there's anything to any of this (see Exhibit 31). The first signal was just standard nine-month price momentum, followed by the nine-month momentum of each firm's GICS industry peers, and finally the nine-month momentum of each firm's text-based peers. On average since 1997 the text-based signal did indeed generate the best returns in our large-cap universe. However, digging deeper reveals some big differences over time. In the late-1990s, at the peak of the New Economy era, pretty much any momentum signal worked as long as it put you into the high-flying Dot Coms (see Exhibit 32). But things flipped around in the decade of the 2000s, when all three signals were perverse, with poor momentum stocks, the black bars, faring better on average than good momentum stocks, the grey bars (see Exhibit 33). In the latest decade companies with good text-based momentum have outperformed marginally but those with poor text-based momentum have lagged by less than those with poor scores on the other momentum metrics (see Exhibit 34).

Exhibit 31: Large-Capitalization Stocks
Relative Returns to the Best and Worst Quintiles
of Nine-Month Price Momentum Factors
Measured Over One-Year Holding Periods
1997 Through Late-July 2017

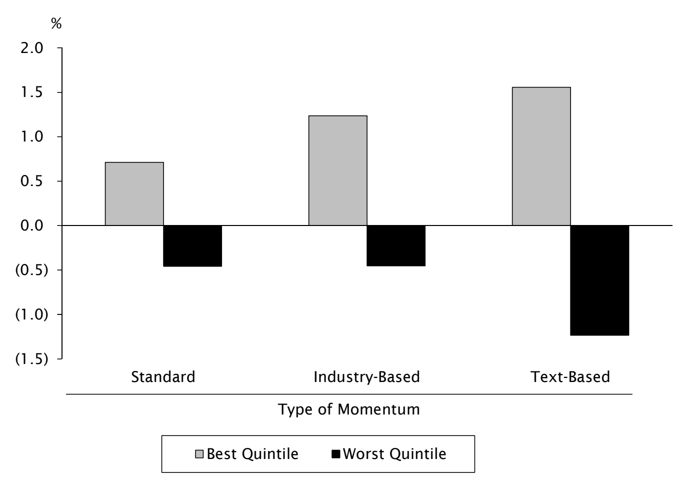

Source: Empirical Research Partners Analysis.

Exhibit 32: Large-Capitalization Stocks
Relative Returns to the Best and Worst Quintiles
of Nine-Month Price Momentum Factors
Measured Over One-Year Holding Periods
1997 Through 1999

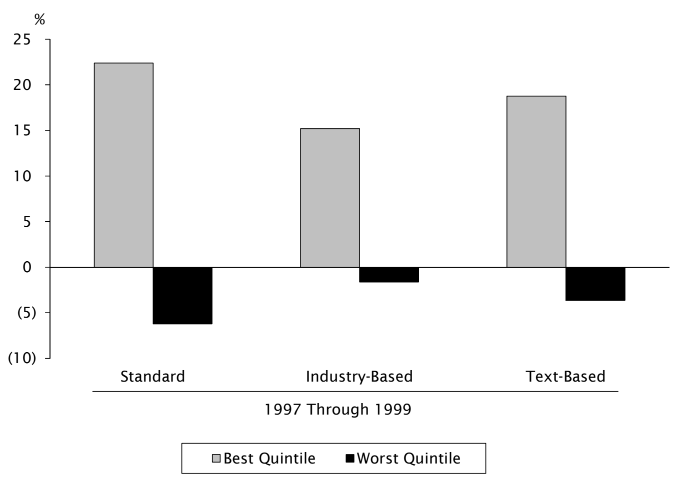

Source: Empirical Research Partners Analysis.

All in all the pick-up in performance by moving to a text-based peer group isn't that compelling, with most of it coming from being less-bad in the 2000s when momentum was terrible anyway. We think this will be par for the course going forward: much of the burgeoning Big Data research seems to involve lots of effort and expense for small gains at the margins of existing signals. What's totally missing in this particularly strand of research is whether one should spend time on momentum at all; we've argued that structural changes, like the rise in passive ownership, has made the signal less effective than in the past. While hacking the p-values for price momentum a little higher using Big Data is an interesting academic exercise, it misses the cost-benefit analysis of whether it's even worth trying to jury-rig a signal that's mostly fizzled. With Big Data the bar for the coolness factor is low, but the bar for real-world value-add is a lot higher.

Exhibit 33: Large-Capitalization Stocks
Relative Returns to the Best and Worst Quintiles
of Nine-Month Price Momentum Factors
Measured Over One-Year Holding Periods
2000 Through 2009

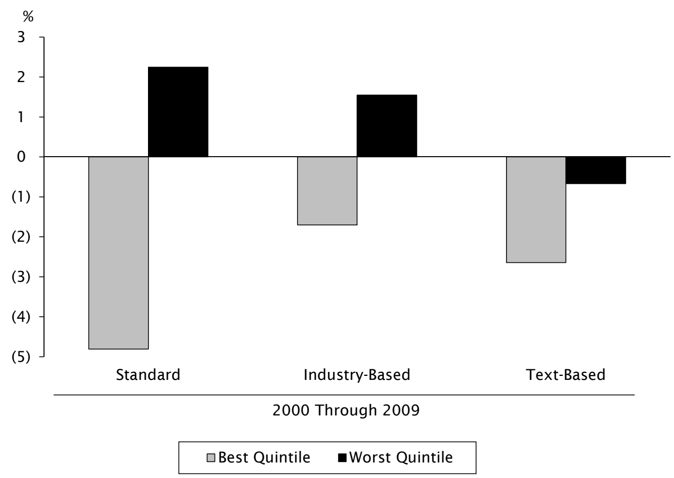

Source: Empirical Research Partners Analysis.

Exhibit 34: Large-Capitalization Stocks
Relative Returns to the Best and Worst Quintiles
of Nine-Month Price Momentum Factors
Measured Over One-Year Holding Periods
2010 Through Late-July 2017

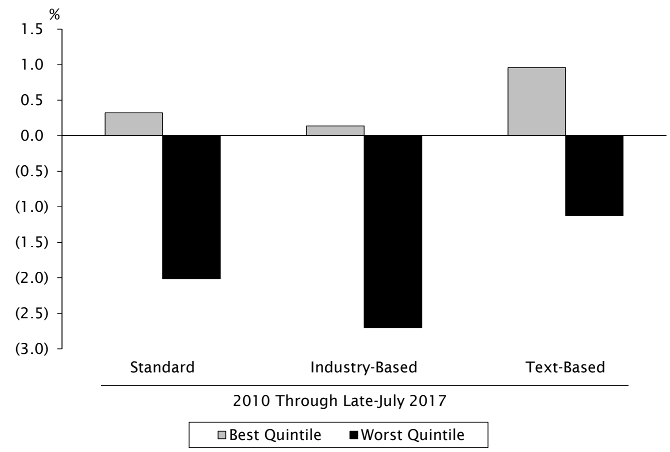

Source: Empirical Research Partners Analysis.

Appendix 1: Recent Academic Papers on Big Data

Agarwal, S., Jensen, J., and Ferdinando Monte, 2017. "The Geography of Consumption." SSRN Working Paper, available at https://ssrn.com/abstract=3002231.

Abstract: We use detailed information from U.S. consumers' credit card purchases to provide the first large scale description of the geography of consumption. We find that consumers' mobility is quite limited and document significant heterogeneity in the importance of gravity across sectors. We develop a simple model of consumer behavior, emphasizing the role of the durability/storability of products, to organize the main stylized facts. Heterogeneity in the storability of products across sectors generates a positive correlation between the strength of gravity and the frequency of transactions at the sector level; this correlation is a clear feature of the data. Using daily rain precipitation from thousands of weather stations in U.S., we show that shocks to travel costs change the spatial distribution of expenditure, and they do so differentially across sectors: hence, the level and heterogeneity of travel costs' shape the level and elasticity of any merchant's demand. This evidence suggests that incorporating the demand-side is essential to analyzing the distributional consequences of local and aggregate shocks across regions. These results also suggest the demand-side is critical to understanding the location of firms and employment in the large and understudied service sector.

Born, J., Myers, D., and William Clark, 2017. "Trump Tweets and the Efficient Market Hypothesis." SSRN Working Paper, available at https://ssrn.com/abstract=2973186.

Abstract: In a Semi-Strong Form (SSF) Efficient Market, asset prices should respond quickly and completely to the public release of new information. In the period from his election on 11/8/16 and his swearing in ceremony on 1/20/17, President-elect Trump posted numerous statements ('tweets') on his Twitter messaging service account that identified ten publicly traded firms. In the absence of new information, the Efficient Market Hypothesis (EMH) predicts that these announcements should have little or no price impact on the common stocks of these firms. Using standard event study methods, we find that positive (negative) content tweets elicited positive (negative) abnormal returns on the event date and virtually all of this effect is from the opening stock price to the close. Within five trading days, the CARs are no longer statistically significant. President-elect Trump's tweets were associated with increases in trading volume and Google Search activity. Taken as a whole, the price and trading volume response, combined with Google Search activity is consistent with hypothesis that it was small/noise traders who were acting on President-elect Trump's tweets and that their impacts were transitory.

Gan, Q., and Buhui Qiu, 2017. "Do Corporate Managers Manipulate Disclosure Through Changing 10-K File Size?" SSRN Working Paper, available at https://ssrn.com/abstract=3004403.

Abstract: File size is a simple measure of disclosure document readability. This study shows that 10-K file size change has negative and robust cross-sectional stock return predictability. A hedge portfolio based on 10-K file size change generates an abnormal return spread of more than 3% per annum. 10-K file size change also has negative predictability on future cash flow news and the return predictability of 10-K file size change reflects mainly its information content on future cash flow news. Consistent with disclosure manipulation, the return predictability of 10-K file size change is found to be stronger for firms with positive file size changes, high information asymmetry, or low recent investor attention, and it derives from the discretionary component of file size change that reflects mainly managerial disclosure discretion. Overall, the findings strongly suggest that corporate managers engage in disclosure manipulation through changing 10-K file size.

Jung, M., Wong, M., and Frank Zhang, 2017. "Buy-Side Analysts and Earnings Conference Calls." *Journal of Accounting Research*, Forthcoming.

Abstract: Companies' earnings conference calls are perceived to be venues for sell-side equity analysts to ask management questions. In this study, we examine another important conference call participant — the buy-side analyst — that has been underexplored in the literature due to data limitations. Using a large sample of transcripts, we identify 3,834 buy-side analysts from 701 institutional investment firms that participated (i.e., asked a question) on 13,332 conference calls to examine the determinants and implications of their participation. Buy-side analysts are more likely to participate when sell-side analyst coverage is low and dispersion in sell-side earnings forecasts is high, consistent with buy-side analysts participating when a company's information environment is poor. Institutional investors trade more of a company's stock in the quarters in which their buy-side analysts participate on the call. Finally, we find evidence that buy-side analyst participation is associated with company-level absolute changes in future stock price, trading volume, institutional ownership, and short interest.

Swanson, N., and Weiqi Xiong, 2017. "Big Data Analytics in Economics: What Have We Learned So Far, and Where Should We Go from Here?" SSRN Working Paper, available at https://ssrn.com/abstract=2998299.

Abstract: Research into predictive accuracy testing remains at the forefront of the forecasting field. One reason for this is that rankings of predictive accuracy across alternative models, which under misspecification are loss function dependent, are universally utilized to assess the usefulness of econometric models. A second reason, which corresponds to the objective of this paper, is that researchers are currently focusing considerable attention on so-called big data, and on new (and old) tools that are available for the analysis of this data. One of the objectives in this field is the assessment of whether big-data leads to improvement in forecast accuracy. In this survey paper, we discuss some of the latest (and most interesting) methods currently available for analyzing and utilizing big data when the objective is improved prediction. Our discussion includes a summary of various so-called dimension reduction, shrinkage, and machine learning methods, as well as a summary of recent tools that are useful for ranking prediction models associated with the implementation of these methods. We also provide a brief empirical illustration of big-data in action, in which we show that big data are indeed useful when predicting the term structure of interest rates.

Tarik Umar, 2017. "Complexity Aversion When Seeking Alpha." SSRN Working Paper, available at https://ssrn.com/abstract=3006405.

Abstract: I provide causal evidence that complexity and sentiment matter for attention to news and market reactions. First, using field data with randomization from Seeking Alpha, I find a standard-deviation increase in headline length (negativity) leads to 12%-fewer (2%-more) views. The effects are larger for less-sophisticated investors. Second, using company-earnings-release headlines, I find complexity has a market effect by instrumenting headline length with company-name length. A standard-deviation increase in length leads to 5%-fewer trades, 35-basis-points-tighter-intraday-price ranges, and 40-basis-points-return underreactions, correcting within two months. Complexity matters more for less-surprising news released on quieter days to less-sophisticated investors.