

My research is in the areas of “behavioral finance” and “social interaction and communication.” All my work is empirical and aims at (a) testing behavioral finance theories and (b) improving our understanding of how financial market participants interact, communicate and learn from each other. I will first describe my work in behavioral finance and then move on to my work on social interaction and communication.

1. Behavioral Finance

The “traditional” framework in finance assumes that investors and managers form rational beliefs. Given such beliefs, investors and managers make decisions based on Expected-Utility preferences. The traditional framework largely neglects market frictions.

Behavioral finance asks whether the behavior of investors and managers can be better described using models that are “psychologically more realistic,” for example by allowing for less-than-fully rational beliefs and non-Expected-Utility preferences. Behavioral finance also asks whether important market frictions enable mispricing induced by “biased” beliefs and “non-traditional” preferences to persist.

I have three papers on less-than-fully rational beliefs, one paper on non-Expected-Utility preferences, and two papers on market frictions.

1.1. Less-Than-Fully Rational Beliefs and Non-Traditional Preferences

✓ **“Country-specific sentiment and security prices” (Journal of Financial Economics 2011)** provides evidence that a country’s popularity can cause prices of companies involved in that country to become meaningfully detached from their fundamental values. The paper looks at country closed-end equity funds (CCEFs). CCEFs are investment companies (“funds”) that are publicly traded in the United States and invest in non-US equity securities. The “New Germany Fund,” for instance, is publicly traded on the New York Stock Exchange and invests in German stocks. Because CCEFs are publicly traded, we can directly observe their market values. Because CCEFs regularly report their security holdings, we also know which and how many stocks these funds hold. This, in turn, allows us to compute the market values of the funds’ underlying assets against which the market values of the funds themselves may be compared. If a country’s popularity affects investor demand and security prices in a manner not fully justified by fundamentals, these two values may diverge.

To examine this possibility, I augment the CCEF data with country popularity scores compiled by Gallup Inc. These country popularity scores are based on surveys in which Americans are asked how they perceive other countries (ranging from “very unfavorable” to “very favorable”). My sample exhibits significant cross-sectional variation (e.g., between Russia and the UK) as well as meaningful time-series variation (e.g.,

regarding France and Germany around the emergence of “Freedom Fries” when the two countries did not support the US-led Iraq War in 2003). My paper finds that the market value of a CCEF (which, given that CCEFs are traded in the United States, is determined by US investors) tends to be substantially below the market value of the corresponding CCEF’s underlying assets if a CCEF invests in a country that is less popular among Americans. The economic significance of these country-sentiment-related “discounts” is substantial. For instance, around the beginning of the Iraq War in 2003, a period marked by sharp drops in France’s and Germany’s popularity among Americans, the average discount of French and German CCEFs temporarily increased from 14.54% to 27.77% before subsequently reverting. I make similar observations regarding American depository receipts, which are similar in structure to closed-end funds. I also find that there is less US money flowing into countries via mutual funds, M&As, and foreign direct investments when these countries are less popular among Americans.

- ✓ **“Price-based return co-movement” (Journal of Financial Economics 2009)**, written jointly with Clifton Green, specifies another biased belief that seems to importantly determine security prices. The stock market is one of the few areas where the price level per se is irrelevant, because a company worth \$100 may decide to issue 10 shares trading at \$10 each, or alternatively, issue 100 shares trading at \$1 each. That is, a company can adjust its stock price to any level it wants simply by adjusting its number of shares outstanding, rendering the construct of stock price levels relatively meaningless.

In reality, investors nevertheless appear to draw strong inferences from stock price levels, often perceiving companies with high stock price levels to be of higher quality and companies with low stock price levels to have more upside potential (practitioner excerpts available upon request). We also know that, in practice, investors group stocks into categories and move money in and out of these categories depending on which category is in fashion at a given point in time (e.g., Barberis and Shleifer 2003). If investors care about stock price levels and move money in and out of price-level categories, similarly priced stocks may move together in excess of what is warranted by fundamentals.

To test this idea, we look at stock splits; stock splits are events in which companies increase the number of shares outstanding, thereby (cosmetically) lowering the price of each share. Our paper finds that after stock splits the affected stocks co-move substantially more with lower-priced stocks and less with higher-priced stocks. Our pattern emerges only after the effective date of a split and not the announcement date. The shift in price-based return co-movement is robust to research design choices that are pertinent to matching firms, the windows over which co-movement is measured and the manner in which it is calculated.

- ✓ Sentiment may affect not only investors, but also corporate managers. A substantial portion of the largest firms in the United States and many other countries are still managed by their founders (e.g., Amazon,

Facebook, FedEx, L Brands, Netflix, Universal Health Services). Prior literature notes that these large founder-managed firms behave quite differently from non-founder-managed firms, for instance incurring much higher capital expenditures than their non-founder counterparts (e.g., Fahlenbrach 2009). One possible behavioral explanation for this difference is that founder CEOs are more overconfident than non-founder CEOs.

“Are founder CEOs more overconfident than professional CEOs? Evidence from S&P 1500 companies” (Strategic Management Journal 2017), written jointly with Hailiang Chen and Joon-Mahn Lee, gauges whether overconfidence can help explain the difference in behavior between founder-managed firms and non-founder-managed firms by directly testing whether founder CEOs are, indeed, more overconfident than their non-founder counterparts. We apply textual analysis to CEOs’ personal tweets and discussions of their firms’ earnings. We also analyze management earnings forecasts and CEOs’ personal trading decisions to infer their level of confidence. Across all four settings, our results indicate that founder CEOs are, indeed, more overconfident than non-founder CEOs.¹

- ✓ Investors may not only exhibit biased beliefs but also possess non-traditional preferences. The best-known model of non-traditional preferences is prospect theory. One unique prediction that comes out of prospect theory is that investors exhibit a strong preference for positive skewness such that even an asset’s idiosyncratic skewness becomes priced. Highly positively skewed securities therefore become overpriced and subsequently earn negative excess returns.

In **“IPOs as lotteries: Skewness preference and first-day returns” (Management Science 2012)**, written jointly with Clifton Green, we test this prediction. Initial Public Offerings (IPOs) are highly positively skewed. If skewness is priced, IPOs should subsequently earn negative excess returns; excess returns should be more negative when IPOs are more highly skewed. To the degree that, during the book-building process, underwriters do not fully factor investors’ skewness preferences into an offer price,² more highly positively skewed IPOs should also experience higher returns from offer price to first-day closing price.

We construct an industry-level measure of expected skewness, which computes whether any stock in the industry of the IPO recently had a breakthrough performance. IPOs in an industry with a recent breakthrough performance are considered to have higher upside than IPOs in an industry with no recent

¹ To us, this is an intuitive but not obvious finding. While many entrepreneurs in start-up companies have been found to be overconfident (e.g., Busenitz and Barney 1997, Forbes 2005), studies on corporate life cycles also argue that the characteristics required of entrepreneurs are significantly different from the characteristics required of successful CEOs of Fortune 500 companies. Founder CEOs who fail to adapt to become more like professional CEOs may therefore find themselves being replaced (Boeker and Karichalil 2002, Hambrick and Crozier 1986).

² In the book-building method, underwriters consult with institutional investors to gauge demand and set the offer price. Kumar (2009) provides evidence that institutional investors do not exhibit skewness preferences.

breakthrough performances (and, indeed, our measure of expected skewness strongly positively predicts actual IPO skewness). We find that our measure of expected skewness positively predicts returns from offer price to first-day closing price and negatively predicts long-run performance. For instance, IPOs in the bottom tercile of expected skewness experience first-day returns and subsequent one-year excess returns of 11.44% and 1.55%, respectively; the corresponding numbers for top-tercile expected skewness IPOs are 25.78% and -10.02%, respectively.

The evidence we report in our paper suggests that prospect theory is useful in explaining the behavior of investors and financial markets.

1.2. Market Frictions

Mispricing induced by less-than-fully rational beliefs and non-traditional preferences can persist only if there are market frictions. Perhaps the most widely discussed potential market friction is short-sale constraints. In “Offsetting Disagreement and Security Prices” (Working Paper 2017) and “Arbitrage Involvement and Security Prices” (Management Science 2017), we provide fresh perspectives on the relevance of short-sale constraints.

✓ In **“Offsetting Disagreement and Security Prices” (Working Paper 2017)**, written jointly with Dong Lou and Chengxi Yin, we make the novel argument that portfolios of securities are generally valued at less than the combined values of the securities in those portfolios. This occurs because companies liked by some investors are often not liked by other investors (“investor beliefs frequently cross”). This makes it almost impossible to construct a portfolio that perfectly pleases large groups of investors and contains (only) every investor’s most favorite companies. The maximum level of excitement that a portfolio of companies receives from investors, therefore, is almost always less than the sum of the level of excitement that the individual companies in the portfolio receive from their most fervent supporters. To illustrate this point by example, consider a portfolio containing Apple and Microsoft. One investor group may be enthusiastic about Apple, but not like Microsoft; another investor group may be enthusiastic about Microsoft, but not like Apple. So while both Apple and Microsoft have their own group of fervent supporters, since beliefs cross, no investor group will be excited about Apple and Microsoft as a package. The level of excitement that “Apple-soft” receives is therefore less than the sum of the level of excitement that Apple and Microsoft receive from their most fervent supporters.

We propose that in the presence of binding short-sale constraints, wherein prices are set by the most optimistic investors, this difference in the level of excitement can become priced and “the whole” can trade below “the sum of its parts.” The degree to which the whole trades below the sum of its parts increases as more investor beliefs cross. If investor beliefs do not cross and, for instance, in a portfolio of two securities one investor group is excited about both the first and second stocks, that same investor group will also be

very excited about the portfolio and the portfolio will not trade at a discount, no matter how strongly another investor group dislikes both the first and second stocks as what matters in the presence of binding short-sale constraints is not the average investor belief, but the belief of the most fervent supporter.

To test our idea, we turn to closed-end equity funds (CEFs). CEFs are investment companies holding portfolios of equity securities. CEFs themselves are publicly traded. If there is a high degree of belief crossing among a CEF's underlying assets, we should expect that CEF to trade at a discount relative to the sum of the values of its underlying equity securities.

We look at all stock-pairs in a given CEF portfolio and we gauge—via analyst earnings forecasts—whether investors who are excited about the first stock also are excited about the second stock (“beliefs do not cross”), or whether investors who are excited about the first stock do not like the second stock (“beliefs cross”). We then compute a portfolio-weighted average level of belief-crossing.

Consistent with our prediction, we find that a one-standard-deviation increase in our measure of belief-crossing comes with a 0.49% increase in the CEF discount. Our effect strengthens with estimates of short-sale constraints in the CEF's underlying assets. We make similar observations across three other settings where the whole and the sum of its parts may be separately evaluated: exchange-traded funds, M&As, and conglomerates.³

In the end, our paper adds to the literature by providing relatively clean evidence of the relevance of short-sale constraints. Our paper also contributes to the literature by making the novel point that, at least in financial markets, “the whole is frequently less than the sum of its parts.”

- ✓ The focus of the previously described paper (along with that of most of the literature) is on how the inability to short can cause prices of stocks that are short-sale constrained to be “too high.” **“Arbitrage Involvement and Security Prices” (Management Science 2018)**, written jointly with Baixiao Liu and Wei Xu, makes the new and perhaps initially surprising argument that the in-ability to short can also cause prices to be “too low.” This is because, in the presence of a well-functioning short-selling market, arbitrageurs (such as hedge funds) can buy seemingly underpriced stocks and simultaneously short industry peers to protect their stock purchases from industry-level and market fluctuations. This essentially lets hedge funds purchase a security without having to worry about the portion of the security's return volatility that is

³ While there already is a large body of literature pointing to the relevance of short-sale constraints, there are also papers offering alternative interpretations for the observations made by proponents of the relevance of short-sale constraints (e.g., Johnson 2004). Some recent studies question the relevance of short-sale constraints altogether (e.g., Kaplan, Moskowitz and Sensoy 2013). The key challenge in this literature is that stocks estimated to be more severely short-sale constrained likely differ along other dimensions; these other dimensions, in turn, may explain differences in prices. Since we are comparing prices of two securities that are composed of the same assets (the aggregate portfolio and the portfolio's underlying components), our analysis is less subject to this omitted-variable problem.

due to industry or market shocks, effectively increasing their “risk-bearing capacity” and allowing them to pursue “attractive buys” more aggressively (e.g., Shleifer and Summers 1990).

To test the relevance of this mechanism, we turn to Hong Kong. In the Hong Kong stock market, only stocks on a list of designated securities can be sold short. This short-sale list is revised on a (mostly) quarterly basis. Consider HSBC Holdings (HSBC) and Chong Hing Bank (CHB), both of which are publicly traded banks in Hong Kong. In May 1994, HSBC was added to the short-sale list and became the first publicly traded bank whose shares can be shorted, (HSBC \equiv “hedge stock”). CHB had a substantially lower market-to-book ratio (MB) than its industry peers and, at least by that metric, appeared undervalued (CHB \equiv “seemingly underpriced stock”). We argue that the addition of HSBC to the short-sale list produced substantial improvements in the ability to hedge risk specific to the banking industry and we test whether the addition of HSBC to the short-sale list encouraged hedge funds to (more) aggressively buy shares of CHB, leading to a permanent rise in CHB’s stock price.

We use a difference-in-differences analysis and find that hedge funds substantially increase their long holdings in seemingly underpriced stocks around the emergence of a hedge stock from 1.71% of shares outstanding to 2.22%. We observe no such pattern for high-MB stocks or among investors that are long only and for whom the change in the ability to hedge industry risk via short positions has little consequence. Our second-stage analyses investigate how this more aggressive buying relates to prices and market efficiency. In line with our hedge fund holdings results, we find that seemingly underpriced stocks experience strong abnormal trading activity and positive abnormal returns of 0.64% in the week following the addition event. No such pattern is observed among stocks with high MB ratios. The positive performance accrues only after the addition becomes effective and not when the addition is announced. The positive performance does not revert. Results from discontinuity analyses and placebo tests strongly suggest that our results are generated by an increase in risk-bearing capacity due to the relaxation of short-sale constraints rather than by industry news.

We see our key contribution to the literature on market frictions as follows: Prior literature argues that in the absence of a deep and liquid short-selling market, arbitrageurs cannot trade against overpricing. Our paper suggests that in the absence of a deep and liquid short-selling market, arbitrageurs also cannot trade aggressively on underpricing as they cannot hedge their long positions. As such, the commonly held view that imposing short-sale constraints helps elevate stock price levels is incomplete. By making it difficult for investors to hedge and aggressively trade on underpricing, imposing short-sale constraints can actually cause stock prices to go down or remain at depressed levels. To the best of our knowledge, we are the first to formally make this point and we hope it becomes part of the broader discussion on the implications of short-sale constraints.

2. Social Interaction and Communication

I am very much interested in how financial market participants interact with and learn from each other. My papers in this area examine (a) how social interactions play out in the corporate board room; (b) how investors—amid the rise of social media—increasingly inform each other; (c) how much employees in knowledge-based industries owe their success to insights generated from interactions with their colleagues; (d) how rapidly information and opinions travel via word of mouth; and (e) how poor communication in corporate annual reports affects asset prices.

2.1. Social Interaction in the Corporate Board Room

✓ My first paper in this area focuses on social interactions in the corporate world, in particular between CEOs and corporate directors. Currently, stock exchanges require the majority of directors to be “independent” of a CEO, i.e., not to have any business or family ties with the CEO. The thinking behind this rule is that independent directors are more effective monitors of the CEO and better representatives of shareholders’ interests.

In **“It pays to have friends” (Journal of Financial Economics 2009)**, written jointly with Seoyoung Kim, we conjecture that while business or familial ties clearly interfere with a director’s effectiveness as a monitor, a significant portion of corporate boards currently categorized as “independent” are not truly independent. The average board in the United States has twelve members. Given the small-group setting and the many social interactions that naturally arise, it is hard to imagine that some directors do not become friendly with the CEO over time, even absent any business or family ties. We conjecture that this holds particularly true if a CEO and a director share background similarities (e.g., both the CEO and the director are from Vermont and both went to Cornell).

To test this idea, we hand-collect biographical data for all CEOs and directors of Fortune 100 companies from 1996 through 2005. We find that background similarities (based on region of birth, alma mater, military service, college major, and primary industry of employment) are quite prevalent on corporate boards.

We find that among boards with greater background similarity, CEOs, on average, receive much higher pay; CEOs are also substantially less likely to be fired after performing poorly. These results are consistent with the idea that social ties negatively impact how directors monitor and discipline CEOs. Given the prevalence of background similarity on corporate boards, our findings therefore suggest that a considerable percentage of boards currently classified as “independent” are substantively not independent.

2.2. Social Interaction and Communication among Investors and Employees

- ✓ While the previously described paper focuses on challenges that arise from social interactions, social interactions can also produce huge benefits. In particular, social interactions can be a highly valuable source of information. This holds particularly true in the context of the rapid growth of social media and our increasing reliance on each other for news, opinions, and advice (e.g., switching to Wikipedia from Britannica). This reliance on the “wisdom of the crowd” has also begun to play a greater role in financial markets as investors increasingly turn to each other for stock recommendations on various investment-related social media outlets.

In **“Wisdom of crowds: The value of stock opinions transmitted through social media” (Review of Financial Studies 2014)**, written jointly with Hailiang Chen, Prabuddha De, and Yu Hu, we ask whether the content generated on such social media outlets is valuable to investors, or whether such content is best described as noise, or, worse, primarily represents attempts of investors’ to manipulate each other through fake news.

We turn to Seeking Alpha, one of the largest investment-related social-media websites in the United States. The channels through which investors can voice their opinions and exchange investment ideas on this website are twofold: Users can write up their thoughts on a stock of their own choosing and submit articles to Seeking Alpha for publication. In response to these articles, any interested user can write a commentary, sharing his or her own view, which may agree or disagree with the author’s view on the company in question. Users can choose to follow other users and receive e-mails whenever those users publish a new article. Users are paid by the number of page views their articles garner. Users thus have a monetary incentive (and perhaps also non-monetary incentives) to develop a positive reputation and a large following within the Seeking Alpha community.

We scrape the entire Seeking Alpha website, resulting in close to 100,000 articles written by around 6,500 distinct users from 2005 through 2012. These articles cover more than 7,000 firms. These articles are commented by more than 180,000 distinct users.

We find that stocks viewed favorably in Seeking Alpha articles and commentaries subsequently earn positive abnormal returns, whereas stocks viewed unfavorably subsequently earn negative abnormal returns. These abnormal performances do not revert. Stock opinions on Seeking Alpha also predict subsequent earnings surprises. These findings suggest that views expressed on Seeking Alpha are valuable to investors, providing a rationale for why investors increasingly turn to each other for stock recommendations on social media.⁴

⁴ Contrary to what we find asserted in some papers that cite our work, our paper **does not** provide evidence that Seeking Alpha articles and commentaries cause stock prices to move. If anything, insofar as we start computing future returns

✓ In “**Information sharing and spillovers: Evidence from financial analysts**” (**Management Science 2018**), written jointly with Jose Liberti and Jason Sturgess, we turn the emphasis from investors to employees and we ask to what degree employees, in knowledge-based industries, owe their success to social interactions with their colleagues.

This is a difficult question to tackle empirically: We cannot fairly contrast the performance of two individuals, one who is surrounded by great colleagues and another who is essentially on her own, and attribute any observed difference in performance to social interactions. This is because the assignment of individuals to firms itself cannot be thought of as occurring independently of an individual’s innate ability (the “selection effect”). We also cannot fairly assess a change in an individual’s performance as she moves from one company to another, because job changes also depend on an individual’s inherent ability.

Instead, our identification comes from sell-side analysts who, while working at the same brokerage, cover multiple M&As—in particular, the acquirer prior to an M&A and the new firm that comes out of the M&A. To illustrate this point by example, consider Analyst A from Goldman Sachs covering two M&As. For the first M&A, Analyst A covers the acquirer only; she has no colleague covering the target (= “Type 1”). In the second M&A, Analyst A covers the acquirer; she also has an in-house colleague covering the target prior to the M&A with whom she can discuss the newly formed firm (= “Type 2”). We compare Goldman Sachs Analyst A’s performance when she is of Type 1 with her performance when she is of Type 2.

Selection explains the matching of Analyst A with Goldman Sachs. Since we examine differences in performance between two M&As for the same analyst working at the same brokerage, we essentially hold selection constant. Furthermore, we include M&A fixed effects to mitigate the concern that analyst type and performance are correlated with M&A characteristics.

Under this empirical design, we find that an analyst’s forecast accuracy is roughly 35% higher when she is of Type 2 than when she is of Type 1. This difference in forecast accuracy is stronger when acquirer- and target-analysts reside in the same locale, when they are part of a smaller team, and when the target-analyst is of higher quality. This difference disappears when the target analyst leaves the broker immediately following the M&A.

To the best of our knowledge, our study is the first that is able to quantify the degree to which high-skilled employees in knowledge-based industries owe their success to interactions with their colleagues.⁵

from three days after an SA article’s publication, our results indicate that financial market participants react to SA articles only mildly (if financial market participants reacted strongly to SA content, there should be no predictability.)

⁵ Prior work, such as Mas and Moretti (2009), who study peer effects among cashiers at a supermarket, has only been able to provide evidence on productivity spillovers among “lower-skilled workers.”

✓ **“The speed of communication” (Work-in-progress 2018)**, written jointly with Shiyang Huang and Dong Lou, represents another attempt at “quantification.” Prior work finds evidence that information spreads via word of mouth as investment decisions made by local peers are highly positively correlated (Feng and Seasholes 2004, Hong, Kubik and Stein 2005, Brown, Ivkovic, Smith and Weisbenner 2008, Kaustia and Knuipfer 2012, Pool, Stoffman and Yonker 2015). We build on these studies and ask not whether, but how strongly and rapidly information spreads via word-of-mouth.

We think this is an important research question. To draw an analogy to epidemiology, if we think of the initial person with a disease as the infected (“patient zero”) and those around him/her as the susceptible, it is important to know whether or not the disease is contagious. It is even more important to know the exact contagion rate and what factors, and to what degree, determine that rate.

Our innovation, which allows us to address the abovementioned question, comes from the use of cross-industry stock-financed M&As as a source of plausibly exogenous variation in investors’ information-gathering activity. In particular, we exploit the fact that, at the completion of a cross-industry stock-financed M&A, investors in the target firm (from industry X) receive shares of the acquiring firm (from industry Y). We conjecture that, once they are endowed with shares of the acquiring firm, target investors more carefully study the acquirer industry, which leads to an increase in target investors’ trading activity in the acquirer industry (outside of the acquiring firm).

If target investors communicate their views to other investors in the same neighborhood via word of mouth, we may observe increased trading activity in the acquirer industry not only by target investors but also by their neighbors and their neighbors’ neighbors. Tracing out the “contagion” of trading in the acquirer industry then enables us to quantify the speed of communication among investors. Unlike correlational studies, our design allows us to pinpoint the initial information event along with patient zero and ask how the speed of communication varies with characteristics of the initial information event and characteristics of patient zero.

This research is ongoing, but our initial analysis reveals that in the year after the completion of a cross-industry stock-financed M&A, target investors more than double their trading frequency in the acquirer industry compared with other investors (again, we exclude the acquirer firm from our calculation). In our data, target investors are surrounded by ten neighbor investors who live within a three-mile radius. We find that, on average, one out of these ten neighbors disproportionately increases her trading in the acquirer industry; the average spike in trading frequency of “infected” neighbors is around 100% and, as such, very similar to the average spike in trading frequency of target investors. The spike in trading frequency abates after around eighteen months. In placebo tests, we find no spike in trading frequency for target investors and their neighbors around cross-industry cash-financed M&As. The degree to which information spreads

from target investors to their neighbors is higher if the initial information event and target investors display certain characteristics.⁶

2.3. The Effects of “Poor” Communication

- ✓ The last paper I will include in my research statement, **“It pays to write well” (Journal of Financial Economics 2017)**, written jointly with Hugh Kim, looks at the effect of “poor” communication on asset prices. Corporate annual reports represent one of the primary mechanisms through which information is disseminated to investors. Corporate annual reports come in the form of accounting numbers framed by a substantial amount of text. While most prior work on corporate disclosure has been on the informativeness of the accounting numbers, our study focuses on the informativeness of the text (“readability”).

In particular, we look at the readability of CEF annual reports. Like all publicly traded corporations, CEFs file annual reports with the Securities and Exchange Commission (SEC) and their shareholders. These reports generally contain a report to the firm’s shareholders; an investment review; an investment outlook; various pieces of information regarding the firm’s officers, directors, and voting policies; a financial statement; and the firm’s security holdings.

Again, because CEFs themselves are traded on stock exchanges, we can compare the market values of the funds against the market values of the funds’ underlying equity securities. That is, we can assess whether CEFs with poorly written annual reports are punished by investors and trade at a discount relative to the total value of their underlying assets.

Our measure of readability draws on the Plain English Handbook of the SEC (1998), which was developed to help firms make their disclosure documents easier to read. In the Handbook, the SEC discusses eight language-related factors that make a document less readable, of which we are able to capture five through the use of a copy-editing software designed to help people write more clearly: (1) passive voice, (2) weak/hidden verbs, (3) superfluous words, (4) legal and financial jargon, and (5) unnecessary details. Our measure of readability is based on the average number of the five abovementioned writing faults per sentence.

To assess the validity of our readability measure, we randomly assign undergraduate business students annual reports with either “high readability” scores or “low readability” scores. We find that students largely agree with the output generated by our readability measure, as they perceive reports with high readability scores to be significantly easier to read than those with low readability scores. We repeat our experiment with

⁶ For instance, if the initial information event represents “very good news” and/or if patient zero is older/ female/ in a high income bracket.

“traditional” measures of readability such as the Fog Index and the Flesch-Kincaid Index and we find that students agree significantly less, suggesting that our measure, which draws directly from the SEC’s Handbook as well as actual writing faults that writing classes teach us to avoid, is a cleaner measure of readability than those utilized in prior literature.

Consistent with the readability of disclosure documents affecting investors and security prices, our panel regression suggests that a 10-percentage-point increase in the number of writing faults per sentence on average increases the CEF discount by 2.7 percentage points. We arrive at the same conclusion when we conduct a difference-in-differences analysis of reports issued around Congress’s enactment of the Plain Writing Act of October 2010 (PWA). The PWA was designed to make documents produced by and filed with government agencies easier to read. We conjecture that while the PWA did not materially impact the readability of annual reports that already had high readability scores, it had an incremental positive impact on the readability of reports whose readability scores had been low. Consistent with this conjecture, we find that, after the PWA took effect, funds with previously low readability experienced sudden, disproportionate, and lasting improvements in their readability scores. Curiously, for these funds, we find a decline in the occurrence of every writing fault listed by the SEC. Our results further show that the disproportionate rise in readability is accompanied by an abnormal and lasting drop in CEF discounts.

In seminal work, Tetlock (2007) examines how the fraction of negative words (“negative sentiment”) in a widely read news column about the stock market relates to future stock market valuations. If we think of sentiment expressed in a text as the “first moment,” we can think of the difficulty with which a text can be processed as the “second moment.” To the best of our knowledge, our study is the first to examine how the second moment impacts stock market valuations and our evidence suggests that the effects are substantial.

References (other than my own work)

- Barberis N, Shleifer A, 2003. Style investing. *Journal of Financial Economics* 68(2): 161-199.
- Boeker W, Karichalil R. 2002. Entrepreneurial transitions: Factors influencing founder departure. *Academy of Management Journal* 45(4): 818-826.
- Brown JR, Ivković Z, Smith PA, Weisbenner S. 2008. Neighbors matter: Causal community effects and stock market participation. *Journal of Finance* 63(3): 1509-1531.
- Busenitz LW, Barney JB. 1997. Differences between entrepreneurs and managers in large organizations: Biases and heuristics in strategic decision-making. *Journal of Business Venturing* 12(1): 9-30.
- Fahlenbrach R. 2009. Founder-CEOs, investment decisions, and stock market performance. *Journal of Financial and Quantitative Analysis* 44(2): 439-466.
- Feng L., Seasholes MS. 2004. Correlated trading and location. *Journal of Finance* 59(5): 2117-2144.
- Forbes DP. 2005. Are some entrepreneurs more overconfident than others? *Journal of Business Venturing* 20(5): 623-640.
- Hambrick DC, Crozier LM. 1986. Stumblers and stars in the management of rapid growth. *Journal of Business Venturing* 1(1): 31-45.
- Hong H, Kubik JD, Stein JC 2005. Thy neighbor's portfolio: Word-of-mouth effects in the holdings and trades of money managers. *Journal of Finance* 60(6): 2801-2824.
- Johnson TC. 2004. Forecast dispersion and the cross section of expected returns. *Journal of Finance* 59(5): 1957-1978.
- Kaplan SN, Moskowitz TJ, Sensoy BA. 2013. The effects of stock lending on security prices: An experiment. *Journal of Finance* 68(5): 1891-1936.
- Kaustia M, Knüpfer S. 2012. Peer performance and stock market entry. *Journal of Financial Economics* 104(2): 321-338.
- Kumar A. 2009. Who gambles in the stock market? *Journal of Finance* 64(4): 1889-1933.
- Mas A, Moretti E. 2009. Peers at work. *American Economic Review* 99(1): 112-145.
- Pool VK, Stoffman N, Yonker SE. 2015. The people in your neighborhood: Social interactions and mutual fund portfolios. *Journal of Finance* 70(6): 2679-2732.
- Shleifer A, Summers LH, 1990. The noise trader approach to finance. *Journal of Economic Perspectives* 4(2): 19-33.
- Tetlock P.C. 2007. Giving content to investor sentiment: The role of media in the stock market. *Journal of Finance* 62(3): 1139-1168.