

Trading on Coincidences*

Alex Chinco[†]

October 18, 2015

Abstract

This paper proposes that traders use coincidences—events where two or more stocks with the same feature realize top-10 returns—to spot promising trading opportunities. If people trade on coincidences, then they will only notice an industry-specific shock, for example, once a pair of stocks from that industry realizes top-10 returns. As a result, the industry as a whole will experience positive abnormal returns the month after a coincidence. The data confirm this prediction. Industries with a coincidence in the previous quarter have abnormal returns of 1.48% in the following month, even after excluding the coincident firms. The effect is strongest in smaller industries and in industries that didn't display a coincidence two quarters ago. Crucially, industries with one top-10 stock and one 11:20 stock don't exhibit any abnormal returns. Results are robust to controlling for momentum, industry momentum, intra-industry cross-autocorrelation, analyst coverage, and institutional trading volume. Industry-wide trading volume rises after a coincidence, and there is evidence of trading on geography-based coincidences.

JEL CLASSIFICATION. D83, D84, G02, G12, G14

KEYWORDS. Coincidences, Cognitive Control, Attention, Behavioral Finance

*I am extremely indebted to Xavier Gabaix for his advice and continuous support. I have received many helpful comments and suggestions from Viral Acharya, Nicholas Barberis, Alexander Ljungqvist, Chris Mayer, Lasse Pedersen, and Jeff Wurgler as well as seminar participants at the NYU, Arizona State, Ohio State, Illinois, Wharton, Yale, and Toronto. All errors are my own.

Current Version: www.alexchinco.com/trading-on-coincidences.pdf

[†]University of Illinois at Urbana-Champaign. alexchinco@gmail.com. (916) 709-9934.

1 Introduction

Open up any financial news site and you'll find people talking about coincidences, events where two or more stocks with the same feature realize top-10 returns. For example, a 2011 Wall Street Journal article noted that the "health-care companies Biogen Idec, Aetna, Humana, and UnitedHealth Group rose more than 40%. . . 4 of the top 10 performers."¹ Why do people care? Shouldn't they just ignore coincidences in favor of more conventional performance measures? Aren't coincidences just meaningless chance events?

At first glance, this logic seems to be baked into the very definition of a coincidence:

Coincidence, *n.* A notable concurrence of events or circumstances having no apparent causal connection. —OED Online. Sep. 2014. Oxford Univ. Press.

Yet, not all coincidences are *mere* coincidences. While a coincidence might just be a notable but meaningless concurrence, it might also suggest a deeper causal connection that wasn't immediately apparent. For instance, after spotting that health-care-industry coincidence, the Wall Street Journal reporter dug a little deeper. He called up Citigroup's chief U.S. equity strategist and asked for an interpretation: "Investors sought out defensive names in areas like health care and staples as well as dividend yielders to protect themselves."

Traders have to constantly be on the lookout for these kinds of feature-specific patterns, but companies have far too many features to investigate every single one. So, how do they solve this cognitive-control problem? That is, how do traders figure out which features are worthy of further attention and which are safe to ignore? This paper proposes that people use top-10 coincidences as a heuristic for spotting promising trading opportunities.

Illustrative Example. Before discussing the empirical analysis in more detail, let's take a look at an illustrative example to see what can be learned from a top-10 coincidence. Imagine sitting at your desk at the start of Q2 2011 and noticing that National Semiconductor posted a 72% return last quarter, putting it among the 10 NYSE stocks with the highest returns. This is a huge quarterly return. But, should you bother checking the semiconductor industry as a whole for a shock to fundamentals?

Probably not. There are always going to be 10 stocks in the list of 10 stocks with the highest returns last quarter. If you're going to check the semiconductor industry in this instance, then you should also do a deep dive into the 8 or 9 other industries also listed in the top 10. What's more, National Semiconductor's performance might not have anything to do with its industry. The company was the subject of an unsolicited takeover offer from Texas Instruments early in 2011, and much of the company's media attention at this time

¹Zuckerman, G. (Jun. 2011) The Market at Midyear: What's Coming Next? *The Wall Street Journal*.

WSJ Articles with “National Semiconductor” in Abstract



Figure 1: National Semiconductor keywords. Reads: “If you pick a Wall Street Journal article about ‘National Semiconductor’, there’s a 5% chance that ‘Layoffs’ is a listed keyword.”

was the result of M&A rumors. Figure 1 shows the subjects used in Wall Street Journal articles about National Semiconductor. It’s peppered with all sorts of topics like “layoffs”, “acquisitions and mergers”, and “appointments and personnel changes”. If you’re going to dig into the semiconductor industry in this case, then you also need to analyze all of the other features for every other stock in the top 10, too! This is a lot of work.

However, if you see *both* National Semiconductor *and* Sequans Communications among the top-10 performers, then this calculus changes. Now, you might actually want to investigate the semiconductor industry. Coincidences like this one are much less likely to occur by pure chance. More importantly, because you observe both firms in the top 10, you can drastically narrow down the list of possible explanations for their shared behavior. For example, you can rule out any M&A-based stories since this wasn’t an issue for Sequans. Quantitatively, if assets display 1 out of every 50 features, then traders focusing on top-10 coincidences have to analyze 15-times fewer features than traders considering every feature displayed by one of the top-10 stocks. In short, coincidences are good heuristic tools for identifying trading opportunities worthy of further investigation.

There are, of course, many different kinds of coincidences. For instance, several of the 10 best-performing IPOs might all be underwritten by the same bank, or several of the top-10 mutual funds might all hold the same asset. Each of these coincidences suggests an unexpected connection that you might want to explore in more detail. It doesn’t matter whether it’s an industry-specific connection, a bank-specific connection, or an asset-specific connection. I study traders’ reactions to a particular species of coincidence, events where two or more stocks from the same industry all earn top-10 returns, because they receive a lot of attention and are easy to collect data on. Nevertheless, I show in the empirical analysis that the paper’s predictions generalize to other kinds of coincidences, such as when two stocks headquartered in the same country both realize top-10 returns.

NYSE's Top 10 Stocks of Q1 2011




Figure 2: NYSE stocks with largest price increases in Q1 2011. Reads: “Barnes & Noble Inc. had the 2nd highest price appreciation during Q1 2011 with their price rising by 80%.”

Hypothesis Development. I begin by analyzing the information content of top-10 coincidences. I show that, if the data are randomly generated, then coincidences contain no useful information and, in fact, occur more often than most people would guess. This is why peoples’ reactions to coincidences are “often cited as an illustration of the irrationality of human reasoning about chance,” in the words of [Griffiths and Tenenbaum \(2007\)](#). But, stock returns are non-random. If stocks realize feature-specific shocks, then the shocked feature is more likely to display a top-10 coincidence. What’s more, there are only a handful of coincidences to investigate each period. So, top-10 coincidences can help traders figure out which features to investigate in more detail when it’s too costly to check every single possibility. They can serve as heuristic cognitive-control devices.

I then use a simple extension of the [Hong, Stein, and Yu \(2007\)](#) regime-switching framework to develop the asset-pricing implications of trading on coincidences. The key insight is that, if people trade on coincidences, then they should only recognize a feature-specific shock once a coincidence has brought it to their attention. So, for example, because an industry that displays a coincidence is more likely to have realized an industry-specific shock, prices in that industry should jump following a top-10 coincidence on average. This effect should be stronger in smaller industries where top-10 coincidences are less likely to occur by pure chance and in industries with no previous coincidences.

Empirical Analysis. These predictions are borne out in the data. Stocks from industries with a top-10 coincidence in the previous quarter have abnormal returns of 1.48% in the subsequent month, even after excluding the coincident firms. Put differently, this 1.48%-per-month point estimate reflects how people revise their beliefs about all the semiconductor stocks except for National Semi and Sequans after seeing both National Semi and Sequans realize top-10 returns in the previous quarter. Just as predicted, the results are stronger following coincidences in smaller industries and following fresh coincidences.

What's more, because traders focus on top-10 coincidences rather than top-9, top-11, or top-12 coincidences, it's possible to distinguish between the effects of trading on coincidences and the effects of an unobserved shock to industry fundamentals. An industry containing the 5th and 9th highest returns in the previous quarter should have identical fundamentals to an industry containing the 3rd and 11th highest returns. The only difference between them is that the first industry happened to display a top-10 coincidence while the second industry didn't—that is, the second industry realized a near miss rather than a coincidence. Thus, any difference between the two industries' subsequent returns must be solely due to the existence (or absence) of a coincidence. Consistent with the trading-on-coincidences hypothesis, while industries display positive abnormal returns after a coincidence, they don't after a near miss.

Robustness Checks. These results are robust to a variety of alternative explanations. First, the findings are unlikely to be explained by firm-specific momentum since I remove the stocks involved in a coincidence when computing abnormal returns. The positive post-coincidence abnormal returns also remain when I include controls for the fraction of each industry held in a [Jegadeesh and Titman \(1993\)](#)-style momentum strategy. Second, the results differ from the industry-specific momentum studied in [Moskowitz and Grinblatt \(1999\)](#) since the post-coincidence abnormal returns exist even in industries that aren't held in an industry-momentum strategy. If an industry displays a top-10 coincidence, then it will have positive abnormal returns in the subsequent month, even if it didn't have particularly stellar returns in the previous quarter. Third, the findings aren't due to intra-industry cross-autocorrelation or analyst coverage. In fact, most of the post-coincidence effect comes from coincidences involving pairs of smaller stocks. Fourth, the results aren't specific to stocks with lots of retail trading, suggesting that both sophisticated and unsophisticated traders use top-10 coincidences as a cognitive-control device. Fifth, the effect operates via increased trading volume. Finally, the results aren't specific to industry groupings. For instance, the other NYSE-listed companies headquartered in the Netherlands tend to realize abnormal returns of 0.93% in the months after at least two Dutch companies earn top-10 returns.

2 Hypothesis Development

This section characterizes what returns should look like if people trade on coincidences and gives some intuition as to why. First, I use numerical simulations to show why traders might want to use coincidences to spot promising trading opportunities. Then, I develop some of the empirical implications of trading on coincidences in the context of the [Hong,](#)

Stein, and Yu (2007) regime-switching framework.

2.1 Numerical Simulations

Why might a trader pay attention to coincidences? Are there some situations where coincidences are more informative than others?

Top-10 Coincidences. To answer these questions, we first need to define a particular kind of coincidence in more detail. Suppose there are N stocks that can display any one of I different features. Let $x_{n,i}$ be an indicator variable such that

$$x_{n,i} = \begin{cases} 1 & \text{if } n\text{th subject has } i\text{th feature, and} \\ 0 & \text{else,} \end{cases} \quad (1)$$

where $\sum_{i=1}^I x_{n,i} = 1$. So, for example, if the features are the $I = 69$ GICS industries, then $x_{n,i} = 1$ if the n th stock belongs to the i th industry.

If we rank each stock according to its period $(t - 1)$ return with $(1:N)$ denoting the stock with the highest return, then we can define the number of stocks with the i th feature among the top 10 performers as

$$top10_i = \sum_{n=1}^{10} x_{(n:N),i}. \quad (2)$$

A top-10 coincidence in feature i occurs whenever there are two or more stocks with the i th feature among the 10 stocks with the highest returns in period $(t - 1)$,

$$coincidence_i = 1_{\{top10_i \geq 2\}}. \quad (3)$$

Data-Generating Process. Suppose there's been a shock to exactly one of the I features. Let's call this shocked feature, i^* . Each stock's returns are given by

$$r_n = \sum_{i=1}^I \theta_i \cdot x_{n,i} + \epsilon_n, \quad (4)$$

where $\epsilon_n \stackrel{\text{iid}}{\sim} N(0, \sigma_\epsilon^2)$ is a noise term and θ_i represents a feature-specific shock,

$$\theta_i = \begin{cases} \bar{\theta} & \text{if } i^* = i, \text{ and} \\ 0 & \text{else.} \end{cases} \quad (5)$$

I study a setting where the shocked feature, i^* , is always selected independently at random from the set of I possibilities—that is, $\Pr(i^* = i) = 1/I$.

Summary Statistics. To get an idea of how much information is conveyed by a top-10 coincidence, we need to know how many coincidences a trader should expect to see,

$$E[\sum_{i=1}^I coincidence_i] \quad (6)$$

Coincidences as Cues




Figure 3: *x-axis:* size of a feature-specific shock, $\bar{\theta}$, in Equation (4). *y-axis, left panel:* average number of top-10 coincidences a trader should expect to see. *y-axis, right panel:* probability that the i th feature has actually realized a shock conditional on observing a coincidence in feature i . Data are generated using 100,000 simulations. There are $N = 3,600$ stocks and $I = 60$ features. Each of the first 20 features affects 30 stocks, each of the next 20 features affects 60 stocks, and each of the final 20 features affects 90 stocks. Reads: “Coincidences become not only more common but also more informative as the size of the feature-specific shock increases.”

and how likely is it that the i th feature actually realized a shock conditional on observing a coincidence in feature i ,

$$\text{Pr}(i^* = i | \text{coincidence}_i). \quad (7)$$

The first estimate captures how many features you’re going to have to investigate if you trade on coincidences. This is really important for real-world traders who can’t afford to spend the time and energy needed to investigate every possible feature in the data. The second estimate captures how likely you are to find something when you investigate a top-10 coincidence.

Cognitive-Control Devices. Figure 3 uses 100,000 simulations to plot both of these statistics as the size of the feature-specific shock in Equation (4) ranges from $\bar{\theta} = 0$ (the data are random) to $\bar{\theta} = 1$ (feature-specific shocks are large). In the simulations, there are $N = 3,600$ stocks and $I = 60$ features. Each of the first 20 features affects 30 stocks, each of the next 20 features affects 60 stocks, and each of the final 20 features affects 90 stocks.

The left panel of Figure 3 shows that, when feature-specific shocks are larger, traders should expect to see more top-10 coincidences. Notice how even when there are no feature-specific shocks, $\bar{\theta} = 0$, traders should still expect to see a top-10 coincidence more than 75% of the time. Most people find this fact surprising. Their common-sense intuition (erroneously) tells them that coincidences should happen less often. This is the usual spurious-coincidence type of result that researchers typically focus on.

But, if we turn our attention to the right panel, which shows how likely it is that the i th feature actually realized a feature-specific shock conditional on observing a top-10

Coincidences at Cues, by Feature Size




Figure 4: *x*-axis: size of a feature-specific shock, $\bar{\theta}$, in Equation (4). *y*-axis, left panel: average number of top-10 coincidences a trader should expect to see when the feature is displayed by $\text{size}_i = 30$ stocks (solid, red), $\text{size}_i = 60$ stocks (dotted, blue), and $\text{size}_i = 90$ stocks (dashed, red). *y*-axis, right panel: probability that the *i*th feature has actually realized a shock conditional on observing a coincidence in feature *i* when the feature is displayed by $\text{size}_i = 30$ stocks (solid, red), $\text{size}_i = 60$ stocks (dotted, blue), and $\text{size}_i = 90$ stocks (dashed, red). Data are generated using 100,000 simulations. There are $N = 3,600$ stocks and $I = 60$ features. Each of the first 20 features affects 30 stocks, each of the next 20 features affects 60 stocks, and each of the final 20 features affects 90 stocks. Reads: “Coincidences in smaller industries are more informative.”

coincidence in feature *i*, we find that coincidences can actually be quite informative if the data are non-random. One out of every 5 coincidences is backed up by a feature-specific shock by the time the shock size has reached $\bar{\theta} = 0.50$. Thus, coincidences are useful cognitive-control devices. While coincidences aren’t sure-fire signals that a feature-specific shock has occurred, traders can use coincidences to intelligently narrow down the list of features that they need to investigate. This pruning function is really valuable in situations, like in the real world, where assets have many overlapping features and only a few of these features matter at any point in time.

More Informative Coincidences. These simulations also reveal useful cross-sectional patterns. There are clearly some situations where coincidences are more informative signals than others. For instance, Figure 4 replicates the analysis in Figure 3, but this time the statistics are calculated separately based on the number of stocks displaying the feature. In the left panel, we can see that a feature is more likely to realize a top-10 coincidence when it’s displayed by more stocks. After all, if more stocks display a feature, then it’s more likely that some of these stocks will end up among the top 10 performers by pure chance. But, continuing with this line of reasoning, we can see in the right panel that coincidences are more likely to be backed up by feature-specific shocks when the industry is smaller. Coincidences in smaller industries are more informative.

2.2 Asset-Pricing Framework

Let's now look at what these statistics mean for asset prices. If people trade on coincidences, then they should only recognize a feature-specific shock once a coincidence has brought it to their attention. So, on average prices should jump following a top-10 coincidence. I use a simple extension of the [Hong, Stein, and Yu \(2007\)](#) regime-switching framework to make this intuition more precise.

Market Structure. Consider a market with N risky assets. Time is discrete, and there is a discount rate of $\bar{r} > 0$. Just as before, let $x_{n,i} \in \{0, 1\}$ represent an indicator variable that is 1 if the n th risky asset has exposure to the i th feature. Each risky asset has exposure to one of I possible features, so $\sum_{i=1}^I x_{n,i} = 1$. At time $t = 0$, one of the I possible features, i^* , begins to realize feature-specific shocks, $\{\theta_{i^*,t}\}_{t \geq 0}$. These shocks follow an AR(1) process,

$$\theta_{i^*,t} = \rho \cdot \theta_{i^*,t-1} + \hat{\theta}_{i^*,t}, \quad (8)$$

with $\rho < 1$ and $\hat{\theta}_{i^*,t} \stackrel{\text{iid}}{\sim} N(0, \sigma_{\theta}^2)$. For all $i \neq i^*$, $\theta_{i,t} = 0$.

Risky assets pay out a dividend each period,

$$d_{n,t} = \sum_{i=1}^I \theta_{i,t} \cdot x_{n,i} + \hat{d}_{n,t}, \quad (9)$$

that is the sum of two components: First, all risky assets realize an idiosyncratic dividend shock, $\hat{d}_{n,t} \stackrel{\text{iid}}{\sim} N(0, \sigma_d^2)$. In addition, stocks with exposure to feature i^* also realize the feature-specific shock, $\theta_{i^*,t}$.

Information and Preferences. Traders are risk neutral. If they are aware of the shock to feature i^* , then immediately after the time t dividend is paid, they see the realization of $\hat{\theta}_{i^*,t+1}$, which they can use to estimate the next dividend payout, $d_{n,t+1}$. Assuming a constant discount rate of \bar{r} , this dividend forecast maps directly into an ex-dividend present value for each stock at time t . If they are not yet aware of the shock to feature i^* , then they naïvely assume that no shock has taken place.

Benchmark Cases. Before looking at trading on coincidences, let's start by analyzing a pair of benchmark cases. First, if traders are fully rational and already know exactly which feature has realized a shock, i^* , then the price of each asset—that is, its ex-dividend present value—will be given by

$$p_{n,t}^{\text{fullyRational}} = \begin{cases} \left(\frac{1}{[1+\bar{r}] - \rho} \right) \times \theta_{i^*,t+1} & \text{if } x_{n,i^*} = 1, \\ 0 & \text{else.} \end{cases} \quad (10)$$

Assuming traders are fully rational is one extreme case. The other extreme case to consider occurs when traders never learn about feature-specific shocks. In this sort of market,

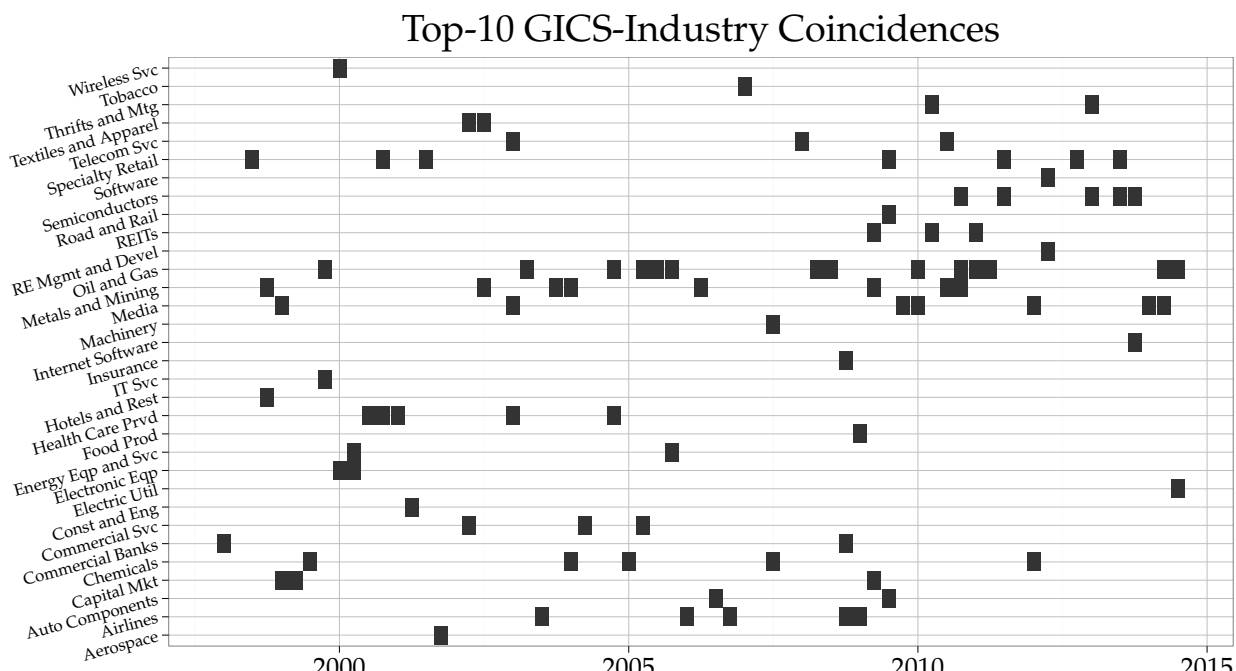


Figure 5: Top-10 GICS-*industry* coincidences using NYSE data from Q1 1998 to Q3 2014. Shaded boxes denote quarters in which at least 2 stocks in an industry had top-10 returns. Industries that didn't realize any top-10 coincidences are omitted to make the figure readable. Reads: "The semiconductor industry realized five top-10 coincidences."

traders price all of the assets the same,

$$p_{n,t}^{\text{neverLearn}} = 0, \tag{11}$$

and each asset's excess return is identical to its dividend stream,

$$r_{n,t}^{\text{neverLearn}} = (d_{n,t} + p_{n,t}^{\text{neverLearn}}) - [1 + \bar{r}] \cdot p_{n,t-1}^{\text{neverLearn}} = d_{n,t}. \tag{12}$$

Almost-Bayesian Updating. There are many different cases sitting in between these two extremes. [Hong, Stein, and Yu \(2007\)](#) explores one of these possibilities, which is based on almost-Bayesian updating. In that paper, people use an over-simplified univariate model for returns when trading. They initially (and wrongly) assume that they have the correct model. Nevertheless, people still track of their forecasting errors. If their original model performs sufficiently poorly over a period of time, then people discard it in favor of an alternative model that would have done better over the same period. The resulting model-switching dynamics can explain some of the puzzles we see in real-world returns, such as the difference in returns between value and glamour stocks.

While this in-between case has many attractive features, it also has some drawbacks. In particular, it doesn't scale well. The original analysis explores a binary setting where people only consider two possible models for returns, *A* and *B*. But, even in this binary

setting, the authors note that traders have to be “capable of performing quite sophisticated multivariate operations when evaluating which model is better, but. . .” still be “unable to make dividend forecasts based on more than a single variable at a time, which sounds somewhat schizophrenic.” This schizophrenia gets worse as the number of possible models grows. In fact, [Natarajan \(1995\)](#) shows that the Bayesian model-selection problem—that is, choosing which K models to use out of $Q \gg K$ possibilities—is computationally intractable. Choosing between simple models causes more headaches than using a simple model cures. This is the essence of traders’ cognitive-control problem.

Trading on Coincidences. This paper proposes a way around this scaling problem: trading on coincidences. Initially, people believe that there hasn’t been any shock and price all the assets accordingly. Each period, they observe assets’ realized returns, which we know from the analysis above are just the realized dividends. But, instead of keeping track of the relative merits of each of the I feature-specific shocks at each point in time, they do something much simpler. Prior to trading each period, people only investigate the collection of features that displayed a coincidence in the previous period. If the feature has realized a shock, then traders immediately recognize it and adjust their beliefs. Investigating a feature is an eye-opening experience as in [Tirole \(2008\)](#). Thus, instead of having to keep track of a running total of the relative merits of each feature-specific shock, traders just have to check for coincidences each period.

If people trade on coincidences, then prices are given by

$$p_{n,t}^{\text{coincidence}} = \begin{cases} \left(\frac{1}{[1+\bar{r}] - \rho} \right) \times \theta_{i^*,t+1} & \text{if } x_{n,i^*} = 1 \text{ and } \max\{(\text{coincidence}_{i^*,\tau})_{\tau=0,1,\dots,t}\} = 1, \\ 0 & \text{else.} \end{cases} \quad (13)$$

Equation (13) says that, if the shocked feature has already displayed a coincidence, then all the stocks with that feature will be correctly priced; however, if the shocked feature hasn’t displayed a feature yet, then all stocks will be priced as if no shock has occurred. So, there will be a price response whenever a shocked feature displays a coincidence.

Post-Coincidence Returns. In the empirical analysis, I study how stock returns change following a coincidence. We just saw that a feature is more likely to display a top-10 coincidence if it’s realized a positive feature-specific shock. So, if the i th feature realizes a top-10 coincidence, then stocks with that feature should see positive returns on average in the following period:

$$0 < E[r_{n,t}^{\text{coincidence}} | x_{n,i} = 1, \text{coincidence}_{i,t-1}] \quad (14a)$$

$$= \Pr(i^* = i | \text{coincidence}_{i,t-1}) \times \begin{cases} \left(\frac{1}{[1+\bar{r}] - \rho} \right) \cdot \theta_{i^*,t+1} & \text{if } \max\{(\text{coincidence}_{i^*,\tau})_{\tau < (t-1)}\} = 0, \\ 0 & \text{else.} \end{cases} \quad (14b)$$

Equation (14b) shows that a feature’s average returns following a coincidence can be written as the product of two terms. The first term represents the probability that a feature has realized a shock given a coincidence. The second term represents the change in the price conditional on realizing that a shock has occurred.

This two-term decomposition formalizes the empirical predictions outlined above. On average, post-coincidence abnormal returns should be positive. Since we know that spurious coincidences are less likely to occur when there are fewer stocks displaying the feature, post-coincidence returns should be larger when features are smaller. They should also only exist when the feature hasn’t displayed a coincidence in the past. If traders already know about a feature-specific shock from a coincidence in a previous period, then telling them about the shock again shouldn’t change how they price it.

3 Empirical Analysis

The data confirm the main predictions of trading on coincidences. Stock prices simultaneously adjust following a coincidence. That is, the prices of all stocks in the semiconductor industry rise after traders notice that National Semiconductor and Sequans Communications both realized top-10 returns in the previous quarter.

3.1 Data Description

I begin by describing the various data sources and variables I use in the analysis below.

Top-10 Coincidences. To identify coincidences, I use daily CRSP data from January 1st, 1998 to September 30th, 2014 to compute the raw quarterly returns for each NYSE-listed stock that’s actively traded on the last day of the quarter. For example, this means computing the raw returns during Q1 2011 for all stocks actively traded on March 31st, 2011. Next, I examine the 10 stocks with the highest returns each quarter and look for GICS-industry coincidences, that is, events where two or more stocks with the same GICS-industry code earn top-10 returns. Figure 5 displays all of the realized coincidences. The figure shows that, for instance, in addition to the coincidence in Q1 2011 mentioned earlier, the semiconductor industry had another coincidence in Q2 2010. Table 1 shows that there are 94 GICS-industry coincidences in this sample.

I restrict the sample in the main analysis to the period after January 1st, 1998 so that I can use the GICS-classification system. This is the standard industry-classification system used by practitioners, and it is important to the analysis that I group stocks into industries in the same way that traders do. As a robustness check, however, I examine the results

Summary Statistics

(a) Full Sample [N = 4,393]

	μ	σ	min	q_{25}	q_{50}	q_{75}	max
$abnormalReturns_{i,t+1}$	0.37	5.91	-63.87	-2.46	0.23	2.75	90.71
$\log_2(\#firms_{i,t})$	4.55	1.48	1.00	3.70	4.46	5.49	9.21
$inMomentumFrac_{i,t}$	0.32	0.20	0.00	0.18	0.31	0.44	1.00
$inIndustryMomentum_{i,t}$	0.31	0.46	0.00	0.00	0.00	1.00	1.00
$\#analysts_{i,t}$	7.91	3.73	0.00	5.50	7.52	10.00	25.67
$instVlmFrac_{i,t}$	0.38	0.15	0.00	0.28	0.37	0.49	1.00
$\log_2(vlm_{i,t+1})$	0.07	0.33	-1.44	-0.14	0.05	0.26	3.67

(b) Coincidence [N = 94]

	μ	σ	min	q_{25}	q_{50}	q_{75}	max
$abnormalReturns_{i,t+1}$	1.82	10.29	-12.06	-1.56	0.98	3.31	90.71
$\log_2(\#firms_{i,t})$	5.98	1.19	3.32	5.24	5.99	6.52	9.17
$inMomentumFrac_{i,t}$	0.51	0.18	0.10	0.38	0.50	0.66	0.92
$inIndustryMomentum_{i,t}$	0.85	0.36	0.00	1.00	1.00	1.00	1.00
$\#analysts_{i,t}$	8.02	2.97	0.62	6.24	8.03	9.77	14.98
$instVlmFrac_{i,t}$	0.36	0.14	0.12	0.25	0.34	0.48	0.62
$\log_2(vlm_{i,t+1})$	0.12	0.31	-0.70	-0.08	0.10	0.39	0.80

(c) No Coincidence [N = 4,299]

	μ	σ	min	q_{25}	q_{50}	q_{75}	max
$abnormalReturns_{i,t+1}$	0.34	5.78	-63.87	-2.48	0.21	2.74	65.41
$\log_2(\#firms_{i,t})$	4.52	1.47	1.00	3.70	4.39	5.43	9.21
$inMomentumFrac_{i,t}$	0.32	0.20	0.00	0.17	0.30	0.44	1.00
$inIndustryMomentum_{i,t}$	0.29	0.46	0.00	0.00	0.00	1.00	1.00
$\#analysts_{i,t}$	7.90	3.74	0.00	5.50	7.51	10.00	25.67
$instVlmFrac_{i,t}$	0.39	0.15	0.00	0.28	0.37	0.49	1.00
$\log_2(vlm_{i,t+1})$	0.06	0.33	-1.44	-0.14	0.05	0.25	3.67

Table 1: Each observation is an (industry, month) pair. Coincidence occurs whenever there are at least two stocks from a GICS industry with top-10 returns in the previous quarter. $abnormalReturns_{i,t+1}$: Industry-average abnormal returns relative to the [Fama and French \(1993\)](#) 3-factor model, excluding coincident firms. $\log_2(\#firms_{i,t})$: Log of the number of stocks in an industry on the last day of the previous quarter. $inMomentumFrac_{i,t}$: Fraction of firms in an industry that would be held in a momentum portfolio based on their returns in the previous quarter. $inIndustryMomentum_{i,t}$: Indicator variable that is one if an industry would be held in an industry-momentum portfolio based on its returns in the previous quarter. $\#analysts_{i,t}$: Average number of analysts covering each firm in an industry in the previous quarter. $instVlmFrac_{i,t}$: Fraction of an industry's trading volume coming from institutional traders in the previous quarter. $\log_2(vlm_{i,t+1})$: Log of the industry-average trading volume, where each stock's volume has been normalized to be mean zero and have unit standard deviation. Sample: Q1/98 to Q3/14.

when using an alternative industry-classification system based on SIC codes. Using SIC codes allows me to study a longer sample period dating back to January 1st, 1978. The results persist in this longer sample using the SIC industry-classification system, although they are less pronounced since traders don't directly use SIC codes to define a company's industry anymore.

Abnormal Returns. To measure each industry's performance following a coincidence, I start by computing each stock's abnormal returns relative to the [Fama and French \(1993\)](#) 3-factor model in the first month of each quarter. The factor loadings come from the daily returns in the previous quarter. So, for example, this means estimating each stock's abnormal returns in April using factor loadings computed using its daily data from January through March. The abnormal return for each GICS industry is then the average of these firm-level estimates. Importantly, when computing these averages in the main analysis, I omit any firms involved in a coincidence. For instance, following the Q1 2011 semiconductor-industry coincidence involving National Semiconductor and Sequans Communications, I compute the mean abnormal return for the semiconductor industry for April 2011 using all semiconductor stocks except for National Semiconductor and Sequans. The resulting panel is organized by (industry, month) with information on 68 dates ranging from January 1998 to October 2014 for each of the 69 GICS industries.

Horizon and Feature Selection. I look at quarterly coincidences because this ranking horizon is the most empirically relevant. [Figure 6](#) shows that Google search intensity for terms like "best performing" among all US investing topics peak in the first month of each quarter—that is, January, April, July, and October. So, looking at each stock's returns in the previous quarter is the natural choice.

I study industry groupings because it is one particularly salient feature. Financial news publications like *Investor's Business Daily* regularly report the top-10 industry movers. Industry-specific facts often show up on the Bloomberg-terminal news feed. The *Wall Street Journal* runs a quarterly segment called the "Winner's Circle" where they take a look at the top-10 performers in the previous quarter and often discuss industry-specific connections. Industry-specific coincidences are a good place to test the theory. In the robustness checks below, I show that the results are robust both to using alternative industry-classification schemes and to looking for country-specific coincidences.

Why Top-10 Coincidences? I focus on top-10 coincidences rather than top-9 or top-11 coincidences since these are the most relevant empirically. For example, while there is a steady stream of Google searches for the term 'top 10' in financial news as reported in [Figure 6](#), Google reports no measurable search results for the terms 'top 8', 'top 9', 'top 11', or 'top 12'. This distinction helps identify the effects of trading on coincidences from

Google Search Intensity for “Top 10”




Figure 6: Monthly search intensity reported by *Google Trends* from January 2007 to January 2014 for U.S. financial news. Raw data is a dimensionless index relative to total number of searches each month. Top panel: Realized (solid) and predicted (dashed) changes in the index value via an AR(6) model with month-of-quarter fixed effects. Bottom-left panel: Month-of-quarter fixed effects, $\gamma_{\text{textMofQ}(t)}$, from the predictive search-intensity auto-regression. Bottom-right panel: Average search intensity for the terms ‘Top 8’, ‘Top 9’, ‘Top 10’, ‘Top 11’, and ‘Top 12’ over the whole sample period. Reads: “People search for the term ‘Top 10’ but not for similar terms like ‘Top 9’ or ‘Top 11’. What’s more, the intensity with which they query Google for ‘Top 10’ rises by 0.53 standard deviations in January, April, July, and October relative to other months.”

the alternative hypothesis that traders are learning about econometrically-unobserved feature-specific shocks in some other way. I elaborate on this test in Subsection 3.3 below.

3.2 Post-Coincidence Abnormal Returns

Table 2 shows that an industry with a top-10 coincidence in the previous quarter has abnormal returns of 1.48% in the subsequent month. Moreover, just as predicted, the table also reports that this effect is strongest in smaller industries where coincidences are least likely to occur by pure chance and in industries that didn’t have a coincidence two quarters ago. Let’s now walk through each of these results in more detail.

Post-Coincidence Abnormal Returns. I start by regressing the abnormal returns of the i th industry in month $(t + 1)$ on an indicator variable that is one if at least two stocks from that industry realized top-10 returns in the previous quarter and zero otherwise,

$$abnormalReturns_{i,t+1} = \alpha + \beta \cdot coincidence_{i,t} + \epsilon_{i,t+1}. \tag{15}$$

There are 94 such coincidences when looking at GICS-industry groupings from January 1998 to October 2014. The first column of Table 2 reports a statistically-significant point estimate of $\beta = 1.48\%$ per month. This means that, for instance, the rest of the semiconductor industry should expect to earn a 1.48%-higher abnormal return in April 2011 because Sequans and National Semi both realized top-10 returns in Q1 2011.

The statistical significance of this point estimate confirms the main prediction of trading on coincidences. If people trade on coincidences, then they should only recognize an industry-specific shock once a coincidence has brought it to their attention. So, because an industry that displays a coincidences is more likely to have realized an industry-specific shock, prices in that industry should jump following a top-10 coincidence on average. This is exactly what I find in the data.

What’s more, this point estimate is also large and economically meaningful. It is on par with the 1%-per-month abnormal returns generated by canonical asset-pricing anomalies like momentum. For example, see [Jegadeesh and Titman \(1993\)](#) or [Asness, Moskowitz, and Pedersen \(2012\)](#). In addition, this point estimate reflects how the prices of the rest of the stocks in an industry react after a coincidence in the previous quarter. It reflects the price movements in April 2011 of all the other semiconductor stocks besides Sequans and National Semi. So, there is less concern about transactions costs related to trading stocks with rapid price movements.

Number of Firms. Let’s now explore some of the cross-sectional implications of trading

Post-Coincidence Abnormal Returns

Dependent Variable:	$abnormalReturns_{i,t+1}$			
<i>intercept</i>	0.34 (0.09)	0.60 (0.29)	0.34 (0.09)	
<i>coincidence_{i,t}</i>	1.48 (0.62)	10.30 (3.15)		[94]
$\log_2(\#firms_{i,t})$		-0.06 (0.06)		
$coincidence_{i,t} \times \log_2(\#firms_{i,t})$		-1.46 (0.52)		
<i>freshCoincidence_{i,t}</i>			1.95 (0.68)	[77]
<i>staleCoincidence_{i,t}</i>			-0.65 (1.44)	[17]
Observations	4,393			
Sample	Q1/98-Q3/14			

Table 2: Each observation is an (Industry, Month) pair. $abnormalReturns_{i,t+1}$: Industry-average abnormal returns relative to the [Fama and French \(1993\)](#) 3-factor model, excluding any coincident firms. $coincidence_{i,t}$: Indicator variable that is one if at least two stocks from the industry realized top-10 returns over the previous quarter and zero otherwise. $\log_2(\#firms_{i,t})$: Log of the number of stocks in an industry on the last day of the previous quarter. $freshCoincidence_{i,t}$: Indicator variable that is one if an industry displayed a coincidence in the previous quarter but not the quarter before that. Point estimates have units of percent per month. Numbers in parentheses are standard errors. Numbers in square brackets are event counts. Reads: “Because Sequans Communications and National Semiconductor realized top-10 returns in Q1 2011, the rest of the semiconductor industry excluding Sequans and National Semi should be expected to earn a 1.48%-per-month higher abnormal return in April 2011.”

on coincidences. First, if there are more firms in an industry, then the industry is more likely to display a top-10 coincidence by pure chance. So, the post-coincidence abnormal returns should be smaller. To test this hypothesis, I regress GICS-industry abnormal returns on the coincidence indicator variable from Equation (15) interacted with the log of the number of firms in that industry,

$$\begin{aligned} abnormalReturns_{i,t+1} = & \alpha + \beta \cdot coincidence_{i,t} \\ & + \gamma \cdot \log_2(\#firms_{i,t}) \\ & + \delta \cdot \{coincidence_{i,t} \times \log_2(\#firms_{i,t})\} + \epsilon_{i,t+1}. \end{aligned} \quad (16)$$

The second column of Table 2 shows that, if you doubled the number of firms in an industry, then the size of its post-coincidence abnormal returns would decrease by $\gamma + \delta = (-0.07\%) + (-1.46\%) = -1.52\%$ points. So, consistent with the theory, the data show that, when coincidences are more likely to be spurious, traders respond to them less aggressively.

Fresh vs. Stale Coincidences. Next, let's look at the impact of how new and unexpected a coincidence is. An industry should only display post-coincidence abnormal returns when traders face a cognitive-control problem—that is, when they weren't already analyzing the industry. To test this prediction, I decompose the 94 top-10 coincidences realized in the sample period into two sets: fresh and stale. I say that an industry is displaying a 'fresh' coincidence if there are at least two stocks from that industry among the top-10 returns in the previous quarter and the industry didn't display a coincidence in the quarter before that. I classify all remaining coincidences as 'stale'. For example, no semiconductor companies listed on NYSE earned top-10 returns in Q4 2010, so the event where Sequans and National Semiconductor realized top-10 returns in Q1 2011 is a fresh coincidence.

I then regress the abnormal returns of the i th industry in month $(t + 1)$ on indicator variables for "fresh" and "stale" coincidences,

$$abnormalReturns_{i,t+1} = \alpha + \beta \cdot freshCoincidence_{i,t} + \gamma \cdot staleCoincidence_{i,t} + \epsilon_{i,t+1}. \quad (17)$$

The third column in Table 2 shows that all of the post-coincidence abnormal returns occur after fresh coincidences. There are 17 instances where traders already knew to analyze an industry due to a previous coincidence, and in these cases there is a slightly negative post-coincidence effect, -0.65% per month. By contrast, in the 77 instances where traders faced a cognitive-control problem, the post-coincidence abnormal returns jump up to over 1.95% per month.

3.3 Top-10 Discontinuity

Finally, because traders tend to focus on top-10 coincidences rather than top-9, top-11, or top-12 coincidences, it's possible to distinguish between the effects of trading on coincidences and the effects of learning about an unobserved shock to industry fundamentals in some other way.

Omitted-Variable Problem. It would be nice if we could interpret the point estimates in Table 2 causally: after observing a coincidence, traders drive up the returns in an industry by 1.48% in the following month. But, we can't do this because of an omitted-variables problem. We saw in Subsection 2.1 that industries which display a top-10 coincidence are more like to have realized a positive shock to fundamentals. That's why it makes sense for traders to pay attention to them in the first place. But, this also means that the typical industry with a coincidence and the typical industry without a coincidence are fundamentally different from one another. So, the post-coincidence abnormal returns in Table 2 might just be due to differences in the fundamentals of the coincident and non-coincident industries rather than traders' reactions to the coincidences themselves.

Summary Statistics, Ctd.

(d) Near Miss	[N = 82]						
	μ	σ	min	q_{25}	q_{50}	q_{75}	max
$abnormalReturns_{i,t+1}$	0.85	5.04	-10.93	-1.90	0.57	2.66	18.07
$\log_2(\#firms_{i,t})$	5.88	1.05	3.70	5.18	6.07	6.25	9.21
$inMomentumFrac_{i,t}$	0.44	0.18	0.09	0.30	0.44	0.57	1.00
$inIndustryMomentum_{i,t}$	0.68	0.47	0.00	0.00	1.00	1.00	1.00
$\#analysts_{i,t}$	7.80	3.03	0.72	5.88	7.66	9.31	15.00
$instVlmFrac_{i,t}$	0.39	0.13	0.15	0.28	0.37	0.50	0.65
$\log_2(vlm_{i,t+1})$	0.04	0.30	-0.69	-0.16	0.01	0.21	1.04

Table 3: Each observation is an (industry, month) pair. A near miss occurs whenever there is one stock from a GICS industry that realizes top-10 returns in the previous quarter and one stock from that industry that realizes 11:20 returns in the previous quarter. $abnormalReturns_{i,t+1}$: Industry-average abnormal returns relative to the *Fama and French (1993)* 3-factor model, excluding coincident firms. $\log_2(\#firms_{i,t})$: Log of the number of stocks in an industry on the last day of the previous quarter. $inMomentumFrac_{i,t}$: Fraction of firms in an industry that would be held in a momentum portfolio based on their returns in the previous quarter. $inIndustryMomentum_{i,t}$: Indicator variable that is one if an industry would be held in an industry-momentum portfolio based on its returns in the previous quarter. $\#analysts_{i,t}$: Average number of analysts covering each firm in an industry in the previous quarter. $instVlmFrac_{i,t}$: Fraction of an industry's trading volume coming from institutional traders in the previous quarter. $\log_2(vlm_{i,t+1})$: Log of the industry-average trading volume, where each stock's volume has been normalized to be mean zero and have unit standard deviation. Sample: Q1/98 to Q3/14.

Near Miss vs. Coincidence. To address this concern, I exploit the fact that real-world traders disproportionately focus on top-10 coincidences rather than top-9, top-11, or top-12 coincidences. See the bottom-right panel of Figure 6. Intuitively, a positive shock to the semiconductor industry is equally likely to push National Semi and Sequans to have the 5th and 9th highest returns as it is to push National Semi and Sequans to have the 3rd and 11th highest returns. The only difference between these two outcomes is that the semiconductor industry would display a top-10 coincidence in the first case but not in the second. Thus, any difference between the features that display a coincidence and those that display a near miss must be due solely to traders' reactions to coincidences and not due to any underlying differences in industry fundamentals. If people trade on coincidences, then the semiconductor industry should display positive abnormal returns after a coincidence but not after a near miss.

To implement this test, I need to define an indicator variable for near misses that is analogous to the one for coincidences. Again, let $x_{n,i}$ be an indicator variable such that $x_{n,i} = 1$ if the n th stock is in the i th industry and $x_{n,i} = 0$ otherwise with where $\sum_{i=1}^I x_{n,i} = 1$. Let $top11:20_i$ demote the number of stocks in the i th industry that had somewhere between the 11th and 20th highest returns last quarter,

$$top11:20_i = \sum_{n=11}^{20} x_{(n:N),i}. \quad (18)$$

A near miss then occurs whenever there is exact one stock from an industry in the top-10 returns last quarter and one stock from that same industry with top-11:20 returns last quarter,

$$nearMiss_i = 1_{\{top10_i=1\}} \times 1_{\{top11:20_i=1\}}. \quad (19)$$

Top-10 Discontinuity. I test whether this 1.48%-per-month abnormal return is a result of traders' response to coincidences by regressing each industry's abnormal returns on this indicator variable for a near miss that is one if one stock from an industry realized 1:10 returns and another stock realized 11:20 returns:

$$abnormalReturns_{i,t+1} = \alpha + \beta \cdot nearMiss_{i,t} + \epsilon_{i,t+1}. \quad (20)$$

There are 82 such near misses in the data from Q1 1998 to Q3 2014. By comparing the estimates from the first and second columns in Table 4, we can see that an industry displaying coincidence will realize abnormal returns of 1.48% in the following month while an otherwise identical industry with a near miss will have statistically insignificant abnormal returns of only 0.49%. This result suggests that the post-coincidence abnormal returns are driven by traders' reactions to coincidences rather than some omitted variable.

4 Robustness Checks

Having outlined the main results, I now discuss a series of robustness checks that help reinforce these findings and allay concerns about potential confounds.

4.1 Momentum Controls

The post-coincidence abnormal returns in Table 2 are unlikely to be driven by momentum for three different reasons.

The Rest of the Industry. To begin with, I remove the stocks involved in any coincidences when computing each industry’s abnormal returns. For instance, in April 2011, I investigate the abnormal returns of all stocks in the semiconductor industry except for National Semiconductor and Sequans. Thus, the 1.48% per month estimate isn’t due to the continued good performance of National Semi and Sequans. Rather, it reflects how traders revise their beliefs about the rest of the semiconductor industry following a coincidence.

Firm-Specific Momentum. Second, the positive post-coincidence abnormal returns remain when I include controls for the fraction of each industry held in a Jegadeesh and Titman (1993)-style momentum strategy. I classify a stock as held in a momentum strategy if it’s returns over the previous quarter are above the 70th percentile for all stocks in that quarter. Obviously, the coincident stocks are going to be held in any momentum strategy.

Top-10 Discontinuity

Dependent Variable:	$abnormalReturns_{i,t+1}$		
<i>intercept</i>	0.34 (0.09)	0.36 (0.09)	
<i>coincidence_{i,t}</i>	1.48 (0.62)		[94]
<i>nearMiss_{i,t}</i>		0.49 (0.66)	[82]
Observations	4,393		
Sample	Q1/98-Q3/14		

Table 4: Each observation is an (industry, month) pair. $abnormalReturns_{i,t+1}$: Industry-average abnormal returns relative to the Fama and French (1993) 3-factor model, excluding any coincident or near-miss firms. $coincidence_{i,t}$: Indicator variable that is one if at least two stocks from the industry realized top-10 returns over the previous quarter and zero otherwise. $nearMiss_{i,t}$: Indicator variable that is one if one stock from an industry realizes top-10 returns in the previous quarter and one stock from that same industry realizes 11:20 returns in the previous quarter. Point estimates have units of percent per month. Numbers in parentheses are standard errors. Numbers in square brackets are event counts. Reads: “Although industries that display a top-10 coincidence in the previous quarter realize positive abnormal returns in the following month, industries that display a near miss don’t.”

But, an industry can display a coincidence even if relatively few of the other stocks have high returns.

I then regress each industry’s abnormal returns on both the coincidence indicator variable from Equation (15) as well as this variable representing the fraction of firms in the industry held in a firm-level momentum portfolio à la [Jegadeesh and Titman \(1993\)](#):

$$abnormalReturns_{i,t+1} = \alpha + \beta \cdot coincidence_{i,t} + \gamma \cdot inMomentumFrac_{i,t} + \epsilon_{i,t+1}. \quad (21)$$

If the post-coincidence abnormal returns were just the result of loading on industries with lots of firm-level momentum, that is, on industries with really good past performance, then adding this firm-level momentum control variable should compete away the effect. The second column of Table 5 shows that this is not the case. While the point does shrink a little bit, from 1.48% per month to 1.38% per month, the effect is still statistically significant and economically large. Some of the time the post-coincidence abnormal returns can be captured by firm-level momentum; indeed, it would be surprising if the effect didn’t shrink at all. Nevertheless, over $1.38/1.48 = 93\%$ of the original point estimate remains after controlling for firm-level momentum. Post-coincidence abnormal returns exist in industries where only two of the firms were superstars and the rest were average Joes in the previous quarter.

Momentum Controls

Dependent Variable:	$abnormalReturns_{i,t+1}$				
<i>intercept</i>	0.34 (0.09)	0.18 (0.17)	0.22 (0.11)	0.25 (0.17)	
<i>coincidence_{i,t}</i>	1.48 (0.62)	1.38 (0.62)	1.26 (0.63)	1.26 (0.63)	[80]
<i>inMomentumFrac_{i,t}</i>		0.51 (0.44)		-0.14 (0.59)	
<i>inIndustryMomentum_{i,t}</i>			0.40 (0.20)	0.44 (0.26)	
Observations	4,393				
Sample	Q1/98-Q3/14				

Table 5: Each observation is an (industry, month) pair. $abnormalReturns_{i,t+1}$: Industry-average abnormal returns relative to the [Fama and French \(1993\)](#) 3-factor model excluding any coincident firms. $coincidence_{i,t}$: Indicator variable that is one if at least two stocks from the same industry realize top-10 returns over the previous quarter. $inMomentumFrac_{i,t}$: Fraction of firms in an industry that would be held in a momentum portfolio based on their returns in the previous quarter. $inIndustryMomentum_{i,t}$: Indicator variable that is one if an industry would be held in an industry-momentum portfolio based on its returns in the previous quarter. Point estimates have units of percent per month. Numbers in parentheses are standard errors. Numbers in square brackets are event counts. Reads: “The 1.48% per month abnormal returns earned by the rest of the semiconductor industry in April 2011 after Sequans and National Semiconductor top-10 returns in Q1 2011 cannot be explained by momentum or industry momentum.”

Industry-Specific Momentum. Another potential confound for the results in Table 2 is the industry momentum analyzed in Moskowitz and Grinblatt (1999). I create an indicator variable that is one if a GICS industry would be held in a Moskowitz and Grinblatt (1999)-style industry-momentum portfolio and zero otherwise, where an industry is held in an industry-momentum strategy if its average returns over the previous quarter are above the 70th percentile among all GICS industries in that quarter. An industry can display a coincidence even if its average returns across all its constituent stocks isn't very high.

I then run a regression of each industry's abnormal returns on both the coincidence indicator variable from Equation (15) as well as this indicator variable that is one if the industry's ranking period returns were in the top 30% for the previous quarter:

$$abnormalReturns_{i,t+1} = \alpha + \beta \cdot coincidence_{i,t} + \gamma \cdot inIndustryMomentum_{i,t} + \epsilon_{i,t+1}. \quad (22)$$

If the post-coincidence abnormal returns were just the result of loading on industries with high past returns, then adding this industry momentum control variable should eliminate the effect. Again, Table 5 shows that this is not the case, with the third column revealing a point estimate of 1.26% per month even after adding the industry-momentum control. The fourth column in Table 5 shows that including both kinds of momentum controls doesn't kill the post-coincidence abnormal returns either.

At first, it might seem a bit odd that an industry would display a top-10 coincidence but not be held in an industry-momentum portfolio. How can this be? Let's consider a short example to see how trading on coincidences differs from industry momentum in more detail. Imagine you are a fund manager pursuing an industry-momentum strategy that is long the 30% of industries with the highest returns over the previous quarter, and think about the following question: how long would it take you to get exposure to a January 2011 semiconductor-industry shock? In order for you to be long the semiconductor industry in April 2011, traders have to bid up the price of every stock in the industry to boost its average return into the top 30%. If Taiwan Semi Corp. happened to have a bad month in February 2011, you might not gain exposure to the industry until May or June. By contrast, a top-10 industry coincidence could draw your attention in April in spite of Taiwan Semi. The portfolio holdings of a trading-on-coincidences strategy and an industry-momentum strategy will only agree when there has been an industry-specific shock to fundamental and *all* stocks in the industry have responded in the same way over the past quarter.

4.2 Information-Flow Controls

Next, I show that the post-coincidence abnormal returns in Table 2 aren't driven by existing patterns in how information flows between companies within an industry.

Analyst Coverage. One possible alternative explanation for the post-coincidence abnormal returns is that, instead of paying attention to coincidences per se, traders are just following important firms with lots of analyst coverage. If this were the case, then traders would be more likely to notice an industry-specific shock when these sorts of high-analyst-coverage firms perform well because there would be more analysts there to explain why they did well. So, I look at whether or not most of the post-coincidence abnormal returns are coming from coincidences involving a pair of stocks with lots of analyst coverage.

To do this, I count the number of analysts that made a recommendation (buy, sell, or hold) for each firm in the previous quarter using data from the Institutional Brokers Estimates System (I/B/E/S) database. This is the standard data used in other papers that study the link between analyst coverage and stock returns, like [Hong and Kacperczyk \(2010\)](#). Then every quarter I split the collection of stocks into two groups, high and low analyst coverage, based on whether or not the firm had more than the median number of

Information-Flow Controls

Dependent Variable:	$abnormalReturns_{i,t+1}$				$\log_2(vlm_{i,t+1})$	
<i>intercept</i>	0.34 (0.09)	0.34 (0.09)	0.34 (0.09)	0.34 (0.09)	0.06 (0.01)	
<i>coincidence_{i,t}</i>	1.48 (0.62)	1.26 (0.67)	1.77 (0.66)	1.48 (0.64)	0.06 (0.03)	[80]
<i>bothHighCoverage_{i,t}</i>		1.25 (1.58)				[17]
<i>bothBigFirms_{i,t}</i>			-2.11 (1.77)			[13]
<i>bothHighInstVlm_{i,t}</i>				-0.06 (2.19)		[8]
Observations	4,393					
Sample	Q1/98-Q3/14					

Table 6: Each observation is an (industry, month) pair. $abnormalReturns_{i,t+1}$: Industry-average abnormal returns relative to the [Fama and French \(1993\)](#) 3-factor model excluding any coincident firms. $\log_2(vlm_{i,t+1})$: Log of the industry-average trading volume, where each stock's volume has been normalized to be mean zero and have unit standard deviation. $coincidence_{i,t}$: Indicator variable that is one if at least two stocks from the same industry realize top-10 returns over the previous quarter. $bothHighCoverage_{i,t}$: Indicator variable that is one if both the coincident firms had an above-median number of analysts covering them in the previous quarter. $bothBigFirms_{i,t}$: Indicator variable that is one if both the coincident firms had above-median market capitalizations at the end of the previous quarter. $bothHighInstVlm_{i,t}$: Indicator variable that is one if both the coincident firms had an above-median amount of institutional trading in the previous quarter. Point estimates have units of percent per month when the left-hand-side variable is $abnormalReturns_{i,t+1}$ and are dimensionless when the left-hand-side variable is $\log_2(vlm_{i,t+1})$. Numbers in parentheses are standard errors. Numbers in square brackets are event counts. Reads: "The 1.48%-per-month abnormal returns realized by an industry after a coincidence persist when controlling for the analyst coverage, size, and institutional-trading volume of the coincident firms. Moreover, trading volume rises in industries after a coincidence."

analysts that quarter.

I run a regression of each industry's abnormal returns on both the coincidence indicator variable from Equation (15) as well as an indicator variable that is one if both the stocks involved in the coincidence had high analyst coverage in the previous quarter:

$$abnormalReturns_{i,t+1} = \alpha + \beta \cdot coincidence_{i,t} + \gamma \cdot bothHighCoverage_{i,t} + \epsilon_{i,t+1}. \quad (23)$$

If the post-coincidence abnormal returns were really just a proxy for traders learning from the good performance of stocks with lots of analyst coverage, then adding this high-analyst-coverage indicator variable should eliminate the effect. The first column of Table 6 shows that this is not the case. While coincidences involving a pair of stocks with high analyst coverage tend to have higher subsequent abnormal returns, this effect isn't statistically significant. More importantly, it doesn't drive out the baseline effect of realizing a coincidence. Put differently, even if both the companies involved in the coincidence have below-median analyst coverage, the industry will still realize post-coincidence abnormal returns of 1.26% per month.

Intra-Industry Cross-Autocorrelation. Another possible concern is the size of the firms involved in the coincidence. Hou (2007) documents that the smaller firms in an industry tend to realize positive abnormal returns in the months after a positive shock to one of the industry's larger firms. In principle, it could be that coincidences are just picking up this previously documented effect; however, there are two reasons why this potential confound is unlikely to explain this paper's results. First, large firms tend to have lower volatility, so they are unlikely to be involved in as many top-10 coincidences. Indeed, Table 6 reveals that there are only 13 instances where both the firms involved in a coincidence had above-industry-median market capitalizations.

Second, it turns out that most of the post-coincidence abnormal returns come from coincidences involving pairs of smaller stocks. To show this, I decompose the 94 top-10 coincidences realized in the sample period into two sets: those involving two big firms and all the rest. An industry displays a big firm coincidence if there are at least two stocks from that industry among the top-10 returns in the previous quarter and both of the firms had above-industry-median market capitalizations for that quarter. For example, that National Semiconductor and Sequans Communications top-10 coincidence in Q1 2011 wouldn't be a both-big-firms coincidence because only National Semiconductor had a market capitalization above the industry's median in Q1 2011.

I run a regression of each industry's abnormal returns on both the coincidence indicator variable from Equation (15) as well as an indicator variable that is one if both the stocks involved in the coincidence had above-median market capitalizations for their industry

in the previous quarter:

$$abnormalReturns_{i,t+1} = \alpha + \beta \cdot coincidence_{i,t} + \gamma \cdot bothBigFirms_{i,t} + \epsilon_{i,t+1}. \quad (24)$$

The third column in Table 6 shows that most of the post-coincidence abnormal returns occur after coincidences involving at least 1 smaller firm. The -2.11% -per-month point estimate on the both-birth-firms indicator variable suggests that the effect is negligible when the coincident firms are large. Thus, it's unlikely that trading on coincidences is really just intra-industry cross-autocorrelation in disguise.

Institutional Trading. A natural next question is: what sort of people trade on coincidences? Is it only retail investors? Or, is it a more general phenomenon? To answer these questions, I next look at the relationship between the post-coincidence abnormal returns and the institutional-trading activity in the coincident firms. If all the effect comes after coincidences involving two firms with very little institutional trading activity, then trading on coincidences is likely something that only retail investors do. By contrast, if there are post-coincidence abnormal returns after coincidences involving a pair of stocks with lots of institutional trading activity, then this is evidence that trading on coincidences is something that even sophisticated institutional investors do.

I use the Thompson-Reuters institutional holdings data to compute the change in holdings from one quarter to the next for each institution required to file a 13f. This is the standard data used in papers like [Gompers and Metrick \(2001\)](#), which look at the link between stock returns and institutional trading activity. Then, Each stock's institutional trading volume is then the sum of these quarterly changes for each stock. This sum represents a lower bound on the total institutional trading volume. An institution might, for example, buy 100 shares and then resell 100 shares of Apple stock during the month of February, which would leave their quarter-to-quarter portfolio holdings unchanged.

I run a regression of each industry's abnormal returns on both the coincidence indicator variable from Equation (15) as well as an indicator variable that is one if both the stocks involved in the coincidence had an above-median fraction of their trading volume coming from institutional investors in the previous quarter:

$$abnormalReturns_{i,t+1} = \alpha + \beta \cdot coincidence_{i,t} + \gamma \cdot bothHighInstVlm_{i,t} + \epsilon_{i,t+1}. \quad (25)$$

The fourth column in Table 6 shows that controlling for the fraction of trading that came from institutional investors in the previous quarter has very little impact on the result. The coefficient on the both-firms-had-high-institutional-trading-volume indicator variable is nearly zero. Thus, it seems like trading on coincidences isn't just something that retail investors do.

Trading Volume. Finally, I give evidence that an industry’s trading volume rises after a coincidence. I compute the log of each stock’s average daily trading volume in the first month of each quarter. Because, some industries have more trading volume on average than others, I then standardize these estimates so that each stock’s log trading volume has a mean of zero and unit standard deviation. I do this subtracting off each stock’s mean trading volume over the previous quarter and dividing through by its standard deviation. After all, if coincidences tend to happen in industries with lots of trading volume, I wouldn’t want to conclude that the trading volume rises after a coincidence.

To test whether or not trading volume rises after a coincidence, I run a variant of the baseline regression in Equation (15) with each industry’s (normalized) log trading volume as the left-hand-side variable rather than its abnormal returns.

$$\log_2(vlm_{i,t+1}) = \alpha + \beta \cdot coincidence_{i,t} + \epsilon_{i,t+1}. \quad (26)$$

The fifth column of Table 6 reveals that the log trading volume rises by an average of 0.06 standard deviations after a coincidence. Comparing this estimate to the summary statistics in Table 1, we can see that this point estimate of 0.06 standard deviations is economically large, amounting to somewhere between 1/5th and 1/6th of the cross-sectional variation we see in the data.

4.3 Other Coincidences

I conclude this section by analyzing the post-coincidence abnormal returns when the coincidences are defined in two alternative ways.

3-Digit SIC-Code Coincidences. First, I re-estimate Equation (15) using the Standard Industrial Classification (SIC) to assign an industry to each company. Using SIC codes rather than GICS codes has the advantage of increasing the sample size. The Global Industry Classification Standard (GICS) was developed in the late 1990’s as a joint project by Standard & Poor’s and Morgan Stanley Capital International. Prior to the advent of the GICS, people used the Standard Industrial Classification (SIC) system developed by the U.S. Office of Management and Budget (OMB). Thus, using SIC codes, I can extend the sample back to Q1 1978. However, traders don’t generally use SIC codes, so they are less likely to notice these SIC code coincidences.

Panel (a) of Table 7 shows the results of trading on SIC-code coincidences. In this specification, two firms are in the same industry if they have the same 3-digit SIC code. The point estimates are similar to those found in the main specification using GICS codes in Table 2: industries realize positive abnormal returns after a coincidence, and these returns are larger when the industry is smaller and when the coincidence is fresh. But,

Other Coincidences

(a) SIC-Code Industry

Dependent Variable:	<i>abnormalReturns</i> _{<i>i,t+1</i>}			
<i>intercept</i>	0.24 (0.04)	0.26 (0.10)	0.24 (0.04)	
<i>coincidence</i> _{<i>i,t</i>}	1.16 (0.74)	8.17 (3.15)		[126]
$\log_2(\#firms_{i,t})$		-0.00 (0.03)		
<i>coincidence</i> _{<i>i,t</i>} × $\log_2(\#firms_{i,t})$		-1.38 (0.61)		
<i>freshCoincidence</i> _{<i>i,t</i>}			1.58 (0.78)	[113]
<i>staleCoincidence</i> _{<i>i,t</i>}			-1.77 (2.30)	[13]
Observations	40,108			
Sample	Q1/78-Q3/14			

(b) Headquarters Location

Dependent Variable:	<i>abnormalReturns</i> _{<i>i,t+1</i>}			
<i>intercept</i>	0.32 (0.17)	0.18 (0.65)	0.32 (0.17)	
<i>coincidence</i> _{<i>i,t</i>}	0.93 (0.43)	5.54 (2.70)		[117]
$\log_2(\#firms_{i,t})$		0.04 (0.17)		
<i>coincidence</i> _{<i>i,t</i>} × $\log_2(\#firms_{i,t})$		-0.89 (0.51)		
<i>freshCoincidence</i> _{<i>i,t</i>}			1.47 (0.53)	[65]
<i>staleCoincidence</i> _{<i>i,t</i>}			0.25 (0.93)	[52]
Observations	1,547			
Sample	Q1/98-Q3/14			

Table 7: *AbnormalReturns*_{*i,t+1*}: Industry-average abnormal returns relative to the *Fama and French (1993)* 3 factor model excluding any coincident firms. *coincidence*_{*i,t*}, Panel (a): Indicator variable that is one if at least two stocks from the same 3-digit SIC-code industry realized top-10 returns over the previous quarter and zero otherwise. *coincidence*_{*i,t*}, Panel (b): Indicator variable that is one if at least two stocks headquartered in the same country realized top-10 returns over the previous quarter and zero otherwise. $\log_2(\#firms_{i,t})$: Log of the number of stocks with that feature on the last day of the previous quarter. *freshCoincidence*_{*i,t*}: Indicator variable that is one if a feature displayed a coincidence in the previous quarter but not the quarter before that. Point estimates have units of percent per month. Numbers in parentheses are standard errors. Numbers in square brackets are event counts. Reads: “Because Sequans and National Semiconductor realized top-10 returns in Q1 2011, the rest of the firms with their 3-digit SIC-code industry, 367, will earn 1.16%-higher abnormal returns in April 2011. The other NYSE-listed companies headquartered in the Netherlands tend to realize abnormal returns of 0.93% in the months after at least two Dutch companies earn top-10 returns.”

these point estimates are less precisely estimated. This is exactly what you'd expect to find if people are trading on coincidences, but SIC codes are really just a noisy proxy for the industry classification system that they're actually using.

Headquarters-Country Coincidences. I primarily study traders' reactions to industry-specific coincidence because they receive a lot of attention and are easy to collect data on. But, there are many different kinds of coincidences, and people shouldn't just pay attention to industry-specific coincidences. For example, traders might want to study a particular country if two stocks headquartered in that country realize top-10 returns. This country-specific coincidence also suggests an unexpected connection worthy of further investigation. To make this point clear, I now give evidence of trading on country-specific coincidences.

I pull the location of each company's headquarters from Compustat. I classify two companies as having the same location if their headquarters are located in the same country. Because the vast majority of NYSE-listed companies are located in United States, I ignore this category in the analysis. Panel (b) of Table 7 gives evidence of trading on country-specific coincidences. The first column in Panel (b) says that the other NYSE-listed companies headquartered in the Netherlands tend to realize abnormal returns of 0.93% in the months after at least two Dutch companies earn top-10 returns. This effect is stronger when there is a coincidence in countries with fewer NYSE-listed firms and following a fresh coincidence.

5 Related Literature

This paper gives evidence of a particular cognitive-control device that people use when trading, namely, top-10 coincidences. But, the problem of how boundedly rational agents learn and allocate their attention has been studied in a variety of other contexts.

Models of Costly Learning. There is a long history of models where people have to pay a cost, either cognitive or monetary, to learn about some feature. In the asset-pricing literature, this setup dates back to at least [Grossman and Stiglitz \(1980\)](#). For some more recent examples, see [Sims \(2003\)](#), [Corwin and Coughenour \(2008\)](#), [Mondria \(2010\)](#), [Van Nieuwerburgh and Veldkamp \(2010\)](#), and [Gabaix \(2014\)](#). In the broader economics and statistics literature, this sort of problem goes back much further. For instance, see [DeGroot \(1969\)](#). But, all of these models take the list of things worth learning about as given. How do people decide which features to learn about?

There are many practical situations where it just isn't clear which things are worth

learning about. Let's take a look at a short example to clarify this point. Suppose a cancer drug passes a vital phase of testing as in [Huberman and Regev \(2001\)](#). It doesn't take much bandwidth to convey this news. People immediately recognized its significance when the New York Times wrote about it. Yet, it clearly took too much work to continuously mine the entire body of new scientific publications each month and select only the relevant journal articles. After all, the journal Nature reported this exact same news five months earlier and no one batted an eye!

Cognitive-Control Problem. The challenge of figuring out which inference problems to solve is known as the cognitive-control problem in the psychology literature and is an active field of research. As [Botvinick, Braver, Barch, Carter, and Cohen \(2001\)](#) writes, "very little is yet known about how the intervention of control processes is itself brought about. [...] The lack of such an account is problematic, for without it control remains a sort of homunculus that 'just knows' when to intercede."

In the past, the economics and finance literature has thought about this cognitive-control problem in one of two ways. First, behavioral economists have studied framing effects in great detail. Here, there people use mental frames or reference points which distort their choices as initially highlighted in [Tversky and Kahneman \(1981\)](#). For some more recent examples, see [Thaler \(1999\)](#), [Peng and Xiong \(2006\)](#), [Rabin and Weizsäcker \(2009\)](#), and [Krüger, Landier, and Thesmar \(2012\)](#). In all these papers, people would be better off if they didn't use frames. Second, and more recently, researchers have developed a collection of models on strategic attention allocation. In these models, boundedly rational traders not only choose how to form beliefs from a set of facts but also choose which facts to use. For some examples, see [Mullainathan \(2002\)](#), [Mullainathan, Shleifer, and Schwartzstein \(2008\)](#), [Gennaioli and Shleifer \(2010\)](#), [Bordalo, Gennaioli, and Shleifer \(2012\)](#), [Köszegi and Szeidl \(2013\)](#), and [Schwartzstein \(2014\)](#). While these papers do an excellent job of capturing the layered nature of traders' problem—that is, traders first have to decide which features to learn about and then do the actual learning—there is generally no discussion of how traders come up with their list of features that might be worth investigating.

6 Conclusion

This paper proposes that traders use coincidences—events where two or more stocks with the same feature realize top-10 returns—to spot promising trading opportunities and then tests this hypothesis empirically. I show that, if people trade on coincidences, then they will only notice an industry-specific shock, for example, once a pair of stocks from that

industry realizes top-10 returns. As a result, the industry as a whole will experience positive abnormal returns the month after a coincidence.

Main Results. The data confirm this prediction. Industries with a coincidence in the previous quarter have abnormal returns of 1.48% in the following month, even after excluding the coincident firms. The effect is strongest in smaller industries and in industries that didn't display a coincidence two quarters ago. Crucially, industries with one top-10 stock and one 11:20 stock don't exhibit any abnormal returns. Industries with a coincidence and those with a near miss have the same fundamentals, so the results are not driven by unobservable differences between industries with a coincidence and those without a coincidence.

The results are robust to controlling for a variety of confounding effects. The positive post-coincidence abnormal returns remain when I include controls for the fraction of each industry held in a momentum strategy. Post-coincidence abnormal returns exist even in industries that aren't held in an industry-momentum strategy. If an industry displays a top-10 coincidence, then it will have positive abnormal returns in the subsequent month, even if it didn't have particularly stellar returns in the previous quarter. The findings aren't due to intra-industry cross-autocorrelation or analyst coverage. In fact, most of the post-coincidence effect comes from coincidences involving pairs of smaller stocks. The results aren't specific to stocks with lots of retail trading and the effect operates via increased trading volume. Finally, there is evidence of trading on geography-based coincidences.

Further Implications. More generally, coincidences are just one kind of salient and informative pattern that traders might use to solve their cognitive-control problem. The human brain is built to recognize and generalize patterns. We engage in this type of learning on a daily basis. To illustrate this point, Ripley (1996) gives the following list of everyday decision problems where pattern recognition plays a key role: "Name the species of a flowering plant. Grade bacon rashers from a visual image. Classify an X-ray image of a tumor as cancerous or benign. Decide to buy or sell a stock option. Give or refuse credit to a shopper." In fact, humans are often much better than machines at these sorts of tasks. The innate pattern recognition skills that make people good doctors, lawyers, and engineers from 9-to-5 are the same skills that alert them to subtle changes in the market when they sit down in front of a terminal. There is much scope for studying other cognitive-control devices.

References

- Asness, C., T. Moskowitz, and L. Pedersen (2012). Value and momentum everywhere. *The Journal of Finance*, Forthcoming.
- Bordalo, P., N. Gennaioli, and A. Shleifer (2012). Salience theory of choice under risk. *The Quarterly Journal of Economics* 127(4), 1243–1285.
- Botvinick, M., T. Braver, D. Barch, C. Carter, and J. Cohen (2001). Conflict monitoring and cognitive control. *Psychological Review* 108(3), 624–652.
- Corwin, S. and J. Coughenour (2008). Limited attention and the allocation of effort in securities trading. *The Journal of Finance* 63(4), 3031–3067.
- DeGroot, M. (1969). *Optimal Statistical Decisions* (1 ed.). McGraw-hill.
- Fama, E. and K. French (1993). Common risk factors in the returns on stocks and bonds. *The Journal of Financial Economics* 33(1), 3–56.
- Gabaix, X. (2014). A sparsity-based model of bounded rationality. *The Quarterly Journal of Economics* 129(4), 1661–1710.
- Gennaioli, N. and A. Shleifer (2010). What comes to mind. *The Quarterly Journal of Economics* 125(4), 1399–1433.
- Gompers, P. and A. Metrick (2001). Institutional investors and equity prices. *The Quarterly Journal of Economics* 116(1), 229–259.
- Griffiths, T. and J. Tenenbaum (2007). From mere coincidences to meaningful discoveries. *Cognition* 103(2), 180–226.
- Grossman, S. and J. Stiglitz (1980). On the impossibility of informationally efficient markets. *The American Economic Review* 70(3), 393–408.
- Hong, H. and M. Kacperczyk (2010). Competition and bias. *The Quarterly Journal of Economics* 125(4), 1683–1725.
- Hong, H., J. Stein, and J. Yu (2007). Simple forecasts and paradigm shifts. *The Journal of Finance* 62(3), 1207–1242.
- Hou, K. (2007). Industry information diffusion and the lead-lag effect in stock returns. *The Review of Financial Studies* 20(4), 1113–1138.
- Huberman, G. and T. Regev (2001). Contagious speculation and a cure for cancer: A nonevent that made stock prices soar. *The Journal of Finance* 56(1), 387–396.
- Jegadeesh, N. and S. Titman (1993). Returns to buying winners and selling losers: Implications for stock market efficiency. *The Journal of Finance* 48(1), 65–91.

- Kőszegi, B. and A. Szeidl (2013). A model of focusing in economic choice. *The Quarterly Journal of Economics* 128(1), 53–104.
- Krüger, P., A. Landier, and D. Thesmar (2012). Categorization bias in the stock market. SSRN Working Paper #2034204.
- Mondria, J. (2010). Portfolio choice, attention allocation, and price comovement. *The Journal of Economic Theory* 145(5), 1837–1864.
- Moskowitz, T. and M. Grinblatt (1999). Do industries explain momentum? *The Journal of Finance* 54(4), 1249–1290.
- Mullainathan, S. (2002). A memory-based model of bounded rationality. *The Quarterly Journal of Economics* 117(2), 735–774.
- Mullainathan, S., A. Shleifer, and J. Schwartzstein (2008). Coarse thinking and persuasion. *The Quarterly Journal of Economics* 123(2), 577–619.
- Natarajan, B. (1995). Sparse approximate solutions to linear systems. *The SIAM Journal of Computing* 24(2), 227–234.
- Peng, L. and W. Xiong (2006). Investor attention, overconfidence, and category learning. *The Journal of Financial Economics* 80(1), 563–602.
- Rabin, M. and G. Weizsäcker (2009). Narrow bracketing and dominated choices. *The American Economic Review* 99(4), 1508–1543.
- Ripley, B. (1996). *Pattern Recognition and Neural Networks* (1 ed.). Cambridge University Press.
- Schwartzstein, J. (2014). Selective attention and learning. *The Journal of the European Economic Association* 12(6), 1423–1452.
- Sims, C. (2003). Implications of rational inattention. *The Journal of Monetary Economics* 50(3), 665–690.
- Thaler, R. (1999). Mental accounting matters. *The Journal of Behavioral Decision Making* 12(3), 183–206.
- Tirole, J. (2008). Cognition and incomplete contracts. *The American Economic Review* 99(1), 265–294.
- Tversky, A. and D. Kahneman (1981). The framing of decisions and the psychology of choice. *Science* 221(4481), 453–458.
- Van Nieuwerburgh, S. and L. Veldkamp (2010). Information acquisition and underdiversification. *The Review of Economic Studies* 77(2), 779–805.