2020

Short and Distort

Joshua Mitts
Columbia Law School, jmitts@law.columbia.edu

Short and Distort

Joshua Mitts

ABSTRACT

Pseudonymous attacks on public companies are followed by stock price declines and sharp reversals. These patterns are likely driven by manipulative stock options trading by pseudonymous authors. Among 1,720 pseudonymous attacks on mid- and large-cap firms from 2010 to 2017, I identify over $20.1 billion in mispricing. Reputation theory suggests these reversals persist because pseudonymity allows manipulators to switch identities without accountability.

1. INTRODUCTION

Anonymous political speech has a celebrated history (including Publius in *The Federalist;* Hamilton, Madison, and Jay [1788] 2013) and has long enjoyed strong protections under the US Constitution.[1] But there is a dark side to pseudonymity: fictitious identities can wreak havoc in financial markets. A large literature in economics examines why markets are vulnerable to rumors and information-based manipulation (Benabou and Laroque 1992; Van Bommel 2003; Vila 1989). In a review of this body of work, Putniņš (2012, p. 957) emphasizes the importance of reputation:

JOSHUA MITTS is Associate Professor of Law and Milton Handler Fellow at Columbia Law School. I am grateful to the editor, two referees, workshop participants at Columbia University and Harvard University, Ian Ayres, Lucian Bebchuk, Bernie Black, Alon Brav, Anthony Casey, Emiliano Catan, Albert Choi, John C. Coffee Jr., Martijn Cremers, Ofer Eldar, Jill Fisch, Elisabeth de Fontenay, Slava Fos, Ed Fox, Merritt Fox, Jens Frankenreiter, Jesse Fried, Jeffrey Gordon, Zohar Goshen, Arpit Gupta, Ernan Haruvy, Scott Hirst, Colleen Honigsberg, Gur Huberman, Robert J. Jackson Jr., Wei Jiang, Donald Langevoort, Jonathan Macey, Justin McCrary, Tamar Mitts, Ed Morrison, Darius Palia, Jim Park, David Pozen, Alex Raskolnikov, Gabriel Rauterberg, Adriana Robertson, Roberta Romano, Natasha Sarin, Robert Scott, Michael Simkovic, Holger Spamann, Eric Talley, Robert Weber, and Perrie Weiner for very insightful comments. I consult on regulatory and litigation matters related to short-and-distort campaigns, but I have no financial interest in any of the cases studied in this paper.

1. In the words of Justice John Paul Stevens, "Anonymity is a shield from the tyranny of the majority" (*McIntyre v. Ohio Elections Commission,* 514 U.S. 334, 357 [1995]), quoting Mill (1946, pp. 3–4).

[*Journal of Legal Studies,* vol. 49 (June 2020)]
 0047-2530/2020/4902-0023$10.00

"if market participants are able to deduce that false information originated from a manipulator, the manipulator will quickly be discredited and the manipulation strategy will cease to be profitable." Pseudonymity undermines these reputational sanctions, allowing manipulators to exploit investors' trust, profitably distort stock prices, and switch fictitious personas with impunity.

In this paper, I show how pseudonymity undermines reputational accountability in financial markets. I examine 2,900 articles attacking mid- and large-cap firms published on the website Seeking Alpha and show that pseudonymous articles are followed by stock price declines and sharp reversals, leading to over $20.1 billion in mispricing. I employ propensity-score matching of pseudonymous and real-name attacks and use a triple-difference design to show abnormal put options trading with publication. On the day of publication, two measures of options trading—so-called open interest (the number of outstanding contracts) and total trading volume—rise for put options written on the target of a pseudonymous article as compared with call options. While I cannot prove that the pseudonymous author is trading, the universe of potential traders is small: only that author, his or her tippees, and the Seeking Alpha editorial staff know an article is forthcoming.

During the second to fifth day following publication, the trading activity in call options written on the target of a pseudonymous article increases, relative to put options, as measured by open interest and trading volume. In the absence of informed or manipulative trading, one would expect that the open interest and trading volume of these contracts would remain similar, on average, across these periods. Indeed, for both measures, call options and put options follow parallel trends in the preceding days, which strengthens the inference that the divergence in open interest and volume is causally attributable to informed or manipulative trading accompanying the article. Following Cremers and Weinbaum (2010), I show that these periods are indicative of informed trading as measured by deviations from put-call parity (a measure of option pricing). A textual analysis suggests that provocative article content is unlikely to be driving these price reversals. The words and phrases correlated with pseudonymous authorship do not refer to fraud or similar evocative improprieties.

Fox, Glosten, and Rauterberg (2018, p. 113) show that "liquidity suppliers will increase their spreads to compensate for the prospect of losing money to misstatement manipulators." I test this proposition by examining how market makers adjust bid-ask spreads in anticipation of informed buying during the reversal period. While the publication of the article comes as a surprise to market makers, they can anticipate the possibility of selling to an informed buyer who purchases in anticipation of a postpublication price correction. I show that a strong negative abnormal return on the publication day is linked to an increase in bid-ask spreads until 2 days after publication, when call options trading is expected to commence.

A central contribution of this paper is to test the predictions of the theoretical literature on reputation (Benabou and Laroque 1992; Van Bommel 2003; Vila 1989). I show that pseudonymous authors may be manipulating markets when they are perceived as nonliars, that is, when they have nonreversals in the past, on average, or have no history. First-time authors must be perceived as nonliars in a Bayesian model. Pseudonymous authors disappear after the market realizes fraud is taking place, which enables them to switch to a new identity. In the Online Appendix, I show that switching pseudonymous identities leaves subtle traces of writing style detectable using stylometry, a method of authorship attribution in computational linguistics.

These price reversals persist because investors learn which authors are liars. In an equilibrium without learning, investors presciently anticipate the probability of fraud, and the price does not reverse on average even when fraud occurs. However, in a learning model like Benabou and Laroque (1992), rational investors can erroneously estimate the chance of fraud, incorrectly concluding that a pseudonymous author might be telling the truth as long as he or she has not been revealed to be a liar. A large literature on dynamic asset pricing shows how current equilibrium prices reflect investors' beliefs conditional on available information, and these beliefs update as new information arrives (Detemple 1986; Pastor and Veronesi 2009). In effect, a fraudster teaches the market to believe the wrong probabilities.

Some may wonder if pseudonymous activists are simply making honest mistakes instead of engaging in intentional manipulation. To be sure, a study of aggregate data does not establish the mens rea required to hold a given defendant liable for securities fraud. But I present evidence of abnormal trading prior to the article's publication date. I show that put options open interest rises on the day of the attack, but because open interest is reported as of 9 a.m. and thus lagged by 1 day, this is direct evidence of abnormal options trading the day before publication of the article. Similarly, I show that this sort of options trading occurs when pseudonymous authors establish credibility under fictitious identities, as suggested by reputation theory. It is difficult to understand why pseudonymous authors would be especially prone to mistakes after establishing credibility under fictitious identities—and precisely when they are trading in advance of publication.

This article contributes to an emerging literature on the link between media and markets. Kogan, Moskowitz, and Niessner (2018) show that the publication of "fake news" on social media, blogs, and similar outlets is followed by temporary price impact and reversals for small firms but not for mid- or large-cap firms. Unlike Kogan, Moskowitz, and Niessner (2018), I focus on analysis pieces by short sellers rather than factually false articles. And while they do not identify pseudonymous authors or consider options trading, Kogan, Moskowitz, and Niessner (2018) show that managers of small- and mid-cap firms may be engaging in market manipulation by issuing press releases, filing form 8-K disclosures, and engaging in insider trading.

Prior work has studied enforcement actions brought against manipulators of over-the-counter and small-cap stocks via spam and message boards (Aggarwal and Wu 2006; Antweiler and Frank 2004; Frieder and Zittrain 2007). But these forums are characterized by anonymity rather than pseudonymity; that is, they do not provide a way to establish a reputation under an assumed name. Options trading is often nonexistent for over-the-counter stocks and thinly traded small-cap stocks. Nonetheless, consistent with my findings, Renault (2018) examines over 7 million posts on Twitter and finds that a burst of social media activity about small-cap stocks is followed by a price increase and subsequent reversal over the next week.

My findings also speak to the large body of work on informed trading in options markets (Augustin et al. 2016; Augustin, Brenner, and Subrahmanyam 2016; Chakravarty, Gulen, and Mayhew 2004; Easley, O'Hara, and Srinivas 1998). Cremers and Weinbaum (2010) find that deviations from put-call parity predict future stock returns: stocks with more expensive call options outperform stocks with more expensive put options. An et al. (2014) examine the joint cross section of stocks and options and find that implied volatility predicts future stock returns, as suggested by a rational model of informed trading. Consistent with this literature, I show that targets of pseudonymous attacks undergo similar deviations from put-call parity in the days accompanying the attack.

Finally, this article relates to the growing literature on activist short selling. Appel, Bulka, and Fos (2018) find that the increasing disclosure of short positions by activist hedge funds is linked to sharp stock price declines. Ljungqvist and Qian (2016) explain the revelation of research by short sellers as a consequence of limits to arbitrage on the short side, which motivates the study of short campaigns in particular. Zhao (2018) identifies a correlation between being targeted by activist short sellers and firms' characteristics like overvaluation and uncertainty. Unlike Zhao (2018), I consider the link between pseudonymous short attacks and market manipulation and do not consider why activist short sellers might target certain firms and not others. Wong and Zhao (2017) show that the targets of short activism experience a subsequent decline in investment, financing, and payouts. Campbell, DeAngelis, and Moon (2017) find that position disclosures by activist short sellers are linked to differences in short-run returns and are not interpreted by investors as evidence of bias.

Legal constraints on short selling have long been known to reduce price efficiency (Beber and Pagano 2013; Boehmer, Jones, and Zhang 2013; Comerton-Forde and Putniņš 2009; Saffi and Sigurdsson 2011), and Fox, Glosten, and Tetlock (2009–10) show that short selling on one day predicts negative news the next. The Securities and Exchange Commission (SEC 2015) has justified short-selling restrictions out of a concern that some shorting could be "manipulative or abusive" of market regulation. My findings suggest that short attacks carried out by pseudonymous authors may indeed be manipulative, which justifies greater regulatory scrutiny. And as I explain in the Online Appendix, pseudonymous attacks pose unique challenges for the law of securities fraud.

2. ANECDOTAL EXAMPLE

As of May 2018, Insulet Corporation (NASDAQ: PODD) was a publicly traded medical device manufacturer with a market value of $5.8 billion. Insulet manufactures the Omnipod insulin pump, which gives diabetics an alternative to multiple daily insulin injections. On November 29, 2016, an article about Insulet was published on Seeking Alpha with a salacious title—"Insulet Investors Being Kept in the Dark, CEO Alleged to Encourage Questionable Sales Techniques"—and its author asserted that he or she had "obtained evidence of yet another whistleblower payoff," that the "CEO allegedly directed employees to bribe physicians," and that "multiple sell-side analysts claimed [the] CEO deceived investors by not fully disclosing the extent of Omnipod product defects and prior management's fraudulent acts" (SkyTides 2016).

There is no immediate indication that the article contained false statements of material facts.[2] The article was written by SkyTides, a pen name for a pseudonymous blogger on Seeking Alpha. The platform proudly encourages pseudonymity, pointing out that "regulations at their workplace or other factors" make "some contributors not able to reveal their real names. In addition, many well-known, veteran stock market bloggers (some of the finest, in fact) write under a pseudonym."[3] The profile page for SkyTides reveals nothing about who this author is.[4]

One might assume that markets would pay little attention to a pseudonymous author like SkyTides. After all, unlike an identifiable author posting under a real name, it is hard to hold SkyTides accountable for authoring misleading or inaccurate information. These kinds of pseudonymous postings seem like a quintessential example of cheap talk lacking credibility (Farrell and Rabin 1996): pseudonymity makes it virtually costless for SkyTides to lie, so rational investors should ascribe little, if any, weight to what SkyTides says.

Immediately following the posting of SkyTides' article, Insulet's stock price fell by over 7 percent from $35.21 on November 28 (the day before the article's publication) to $32.77 on December 1 (2 days after it). Perhaps SkyTides was right—Insulet had serious problems, and the market recognized this by bidding down the price of Insulet's stock. But then Insulet's price climbed right back up on December 5 and ultimately closed higher than before the article was published. Figure 1 shows the stock price of Insulet Corporation from November 21, 2016, to December 14, 2016, and reveals a V pattern centered on the publication of SkyTides' article.[5]

A decline of over 7 percent is highly unlikely to have been caused by

2. The article draws its factual claims narrowly. For example, it claims that there is "evidence" of a whistleblower payoff without characterizing the quality or reliability of such evidence. The article states that one Mr. Oliva "met with a member of PODD's compliance team . . . to voice his objections to [the CEO's] repeated instructions to conduct unlawful acts" and then states that Mr. Oliva settled his claim. This is listed under the heading "Another PODD cover-up and apparent pay-off of a whistleblower" and is the only evidence given as to the existence of any whistleblower payoff. While the author is thus clearly speculating as to the nature of the settlement, this sort of speculation conveys a loose, albeit nonexistent, evidentiary foundation. Similarly, the article refers to "allegations" of bribery and "claims" by analysts. The factual statement that there were allegations may be literally true, even if those allegations are false.

3. Seeking Alpha, Policy on Pseudonymous Contributors (https://seekingalpha.com/page/policy_anonymous_contributors).

4. See Seeking Alpha, SkyTides (https://seekingalpha.com/author/skytides#regular_articles).

5. The stock price graph is from TradingView (https://www.tradingview.com).

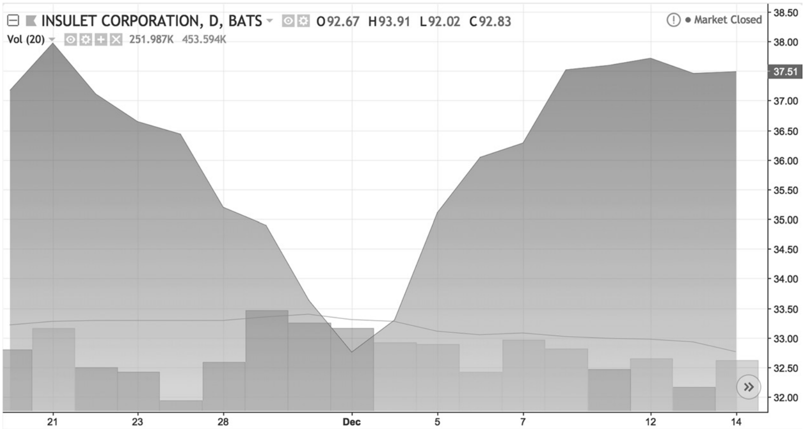

Figure 1. Insulet Corporation's stock price

random chance. And there is evidence that Insulet's stock price was subjected to manipulative options trading alongside the publication of the article. Put and call options are contracts that allow investors to make bets that a company's stock price will rise or fall. These bets are high risk, high reward: if the stock price goes up, the value of a call option increases by a lot, but if it goes down, the call option becomes virtually worthless. Buying options suggests information about where the stock price will go and has featured in SEC insider-trading cases (Augustin et al. 2016; Augustin, Brenner, and Subrahmanyam 2016; Chakravarty, Gulen, and Mayhew 2004; Meulbroek 1992).

Figure 2 plots the number of outstanding put and call options contracts on Insulet's stock in the days before and after the publication of SkyTides' article.[6] Figure 2 shows a large purchase of put options the day before the article's publication,[7] which pay off if the stock price declines (which it did), and a sale of those options thereafter, which would cause the stock price to rise (which it did). That kind of well-timed options trading suggests that someone knew the article was about to be published and that the price would revert to its prior level thereafter because the

6. Figure 2 is limited to options that are nearly at the money; that is, they have an absolute delta between .45 and .55. A delta is a measure of the likelihood that an option will close in the money, and an absolute delta of .5 implies that the option is exactly at the money; that is, it is just as likely to expire in the money as out of the money. Augustin et al. (2016) explain why it is prohibitively expensive for informed investors to trade options that are deeply out of the money.

7. Open interest is lagged by 1 day, so options reported on November 29 were purchased on November 28.

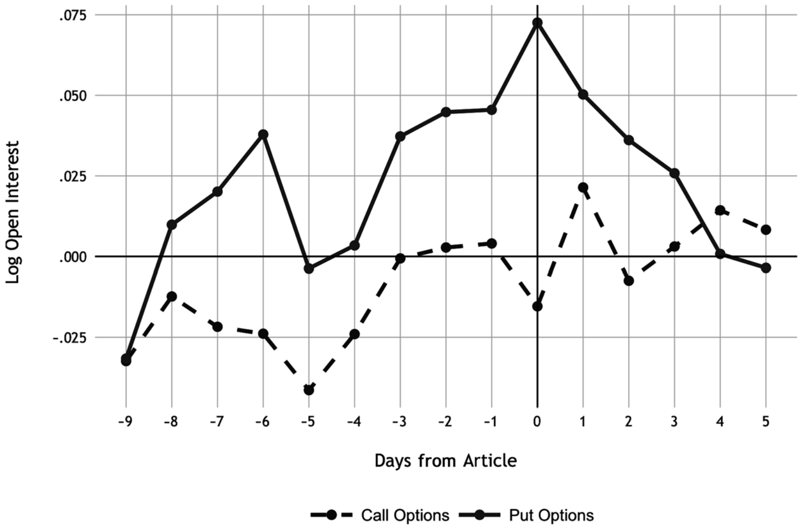

Figure 2. Insulet Corporation's put and call option contracts

article did not contain sufficient information to bring about a downward revision in the price of the magnitude observed on the day of publication. While the subsequent rise in the price could have been driven by public arbitrageurs, nobody knew the article was forthcoming except the author, his or her tippees, and the Seeking Alpha editorial staff,[8] so the put options are especially suspicious.

SkyTides' article on Insulet is hardly an isolated case. Short sellers have increasingly embraced this kind of pseudonymous online activism. Two lawyers at Ropes and Gray LLP wrote that "anonymous online hit pieces against public companies have become an increasingly common and effective form of short activism" (Katz and Hancock 2017) and point to several recent substantial price declines in the wake of pseudonymous attacks. And three lawyers at DLA Piper discussed the case of Chromadex Inc., which was attacked by a pseudonymous short seller and lost $100 million of market capitalization in a single day (Weiner, Weber, and Hsu 2017).

8. As a matter of policy, Seeking Alpha prohibits editors from trading ahead of a forthcoming article (Seeking Alpha, Seeking Alpha Conduct and Investment Policy [https://seekingalpha.com/page/seeking-alpha-conduct-and-investment-policy]).

3. WHY LISTEN TO PSEUDONYMOUS AUTHORS?

Why does pseudonymous market manipulation persist in an efficient market with sophisticated investors who have billions of dollars at stake? A useful starting point is the canonical model of market manipulation articulated in Benabou and Laroque (1992). In that model, a false message leads the market to disregard future messages from a recipient. But this result turns on the market's ability to continue to attribute future messages to the same author. If an author is able to reset the market's prior belief, the market will respond to a manipulative message going forward.

Pseudonymity serves this function by allowing those authors who lack credibility to reset the market's prior belief as to their credibility. The net effect is a kind of pooling equilibrium: absent a history of manipulation, market participants cannot separate a manipulative article from a non-manipulative article when a new pseudonym emerges. As I will show, there is no difference in observable options characteristics that might signal to market participants that manipulation is taking place.

That said, a puzzle remains: why is pseudonymity tolerated by markets? Why not suppose all pseudonymous authors are lying? The classical unraveling result in the disclosure literature shows that a seller who possesses private information that a good is high quality has an incentive to fully disclose this information so as to induce buyers to pay for the quality; absent disclosure, buyers will assume that the seller has something to hide (Grossman 1981). Pseudonymity invites this kind of adverse inference: market participants might rationally conclude that authors with truthful information should have no trouble risking their reputation by making claims using their real names; after all, if the information is true, no harm to their reputation will result. The use of pseudonymity suggests that an author has something to hide. Rational investors should simply ignore pseudonymous articles, inferring that if a claim is truthful, it will be made by an author using a real name.

But that kind of inference breaks down when authors may have a reason other than false information to prefer pseudonymity. Indeed, as I will show, trading on pseudonymous attacks is profitable on average—just not as profitable as it would be absent the price reversals I document. Consistent with this finding, Seeking Alpha justifies its pseudonymity policy by pointing out various reasons contributors would hide their identities. The willingness of market participants to sell the stock of targets of pseudonymous attacks can be rationalized by pointing to the ambiguity underlying the use of a pen name: authors may prefer pseudonymity precisely because they are conveying truthful information and fear the adverse consequences that may result from being identified as the author of such truthful analysis.

In addition to workplace prohibitions on social media commentary, authors may also fear litigation. Attackers may worry that target firms will pursue defamation or securities fraud claims if the publication of an attack piece leads to a decline in the price of the stock. Even if the author can fully establish the truth of every claim made in the piece, doing so would involve protracted, time-consuming litigation, which imposes nontrivial costs. Pseudonymity allows these authors to make damaging but truthful claims without worrying that target firms will bring an unfounded lawsuit that is costly to defend against.

Yet another example is legal uncertainty: the precise contours of securities fraud liability are not always clear. Consider, for example, the Supreme Court's decision in *Omnicare, Inc. v. Laborers District Council Construction Industry Pension Fund* (135 S. Ct. 1318 [2015]), which overturned the Sixth Circuit's holding that statements of opinion that ultimately turn out to be incorrect may constitute an "untrue statement of a material fact" under section 11 of the Securities Act of 1933. Instead, the Court held that opinions may constitute misstatements when they are not sincerely held. This kind of uncertainty in the doctrinal landscape may lead authors to prefer pseudonymity to make it more difficult to be held accountable for violating a legal rule whose interpretation is shifting and subject to judicial clarifications ex post.

For these reasons, market participants may be hesitant to conclude that pseudonymity necessarily implies a lack of credibility. But by allowing authors to effectively switch names, pseudonymity undermines the effectiveness of the reputation mechanism envisioned in Benabou and Laroque (1992). Their model seems to implicitly assume a lack of pseudonymity: "[i]f the insiders' information was perfect, one could easily tell ex post whether or not they had been truthful. In this case they could lie at most once, and sanctioning fraud would eliminate the problem" (p. 924). They further argue that "in reality even private information is not fully reliable, so that the possibility of honest mistakes makes it very difficult to establish fraud conclusively" (p. 924). Yet the mere possibility of honest mistakes is not a roadblock to establishing fraud: there is often evidence as to whether a given misstatement was driven by deceptive intent. On the other hand, it is difficult to sanction pseudonymous authors for fraud measured by ex post price reversals, as discussed in the Online Appendix.

This theoretical framework yields several predictions. First, pseudonymous authors should focus informed trading and manipulation on times when they are perceived by the market as nonliars, such as when they have no history or their history has had few mistakes. I explain below why it must be the case in a Bayesian model that first-time authors are not rejected as implausible. Second, pseudonymous authors should disappear after the market realizes they have been misleading, so that they can switch to a new identity. Finally, identity switching by pseudonymous authors should leave traces of underlying authorship, which may be detectable using techniques of linguistic stylometry.

4. EMPIRICAL ANALYSIS

I now proceed to empirically test these hypotheses. Section 4.1 describes the data, the construction of the sample, and the research design using propensity-score matching. Section 4.2 presents the main results, including comparing cumulative abnormal returns between pseudonymous and nonpseudonymous attacks on public companies, testing whether pseudonymous articles are followed by greater stock price reversals, and evaluating the theoretical prediction that pseudonymous identities are likely to engage in informed or manipulative trading so long as the market cannot conclude that they are lying. Finally, Section 4.3 presents extensions, including examining whether these price reversals are driven by manipulative trading in the options market, variation in bid-ask spreads associated with pseudonymous attacks, and aggregate trading losses due to the mispricing caused by the publication of pseudonymous articles.

4.1. Data

4.1.1. Data and Sample Construction. I begin by collecting all articles published on Seeking Alpha under the category Short Ideas from January 1, 2010, to December 31, 2017. That category contains all articles that advocate taking a short position in one or more firms. Seeking Alpha provides the date and time that the article was published and the ticker symbol of the subject firm(s). This yields an initial sample of 14,730 articles.

To determine which authors are pseudonymous, I hired workers from the crowdsourcing website Figure Eight. I asked workers to look up the name of the author on Seeking Alpha and determine whether he or she is pseudonymous on the basis of the absence of personally identifiable biographical information in the Seeking Alpha profile. For each author, I had three workers evaluate his or her profile, and I coded an author as pseudonymous only if all three authors agreed that the profile did not refer to an identifiable individual. In addition, I manually verified and corrected a few sporadic errors in the coding. Table OA3 in the Online Appendix shows 10 examples of authors from the pseudonymous and nonpseudonymous groups.

To accurately measure trading behavior around the publication of attacks, I remove any article published about the same firm within 7 calendar days of a prior article. A few firms (like Tesla) are the subject of near-daily attacks by short sellers. In that case, it is difficult to view the publication of each additional article as a new informative attack rather than a reiteration of what is already known. Moreover, it is important to verify that the results are not driven by these arguably pathological cases of incessant publications about the same firm rather than publications that bring new information to the market. This yields 9,121 articles about 2,311 publicly traded firms.

In addition, because this study depends heavily on market participants rapidly responding to and trading on the basis of information publicly disclosed in these articles, I limit my primary analysis to mid-cap and large-cap firms with at least $2 billion in market capitalization. The inclusion of small- and micro-cap stocks is problematic, as prices are often much slower to respond, and their relative illiquidity and lower nominal prices lead to much greater return volatility. For this reason, it is difficult to detect price reversals with the same power in this group.[9] This yields 4,785 articles about 837 publicly traded firms.

For each of these firm-article pairs, I obtain standard characteristics from Compustat like market value of equity, total assets, total liabilities, and net income for the year preceding the article, and I derive the Amihud (2002) illiquidity measure and firm-specific volatility using daily returns over the period $[t_0 - 120, t_0 - 7)$, where t_0 is the date of the article's publication. Summary statistics on my primary data set are presented in Table 1. Which firms are targeted by pseudonymous authors? Table 2 considers predictors of pseudonymous authorship among the full sample of 4,785 firm-articles. As Table 2 shows, pseudonymous targets tend to

9. Put differently, for small-cap firms it is difficult to assume relatively strong market efficiency, which is necessary to form a null hypothesis of no expected reversals within a short time window, which thereby facilitates a comparison between reversals and nonreversals. Small-cap firms are very much worth studying but may require different methods that are not as dependent on rapid price efficiency, and for that reason I plan on addressing these in a subsequent project.

Table 1. Summary Statistics

Variable	N	Mean	SD	Min	Max	Percentile 25th	Percentile 50th	Percentile 75th
$\text{CAR}_{i,j,t_0+2,t_0+5}$	4,784	.001	.048	−.722	.4	−.018	.001	.021
$\text{CAR}_{i,t_0-1,t_0+1}$	4,785	−.006	.069	−1.94	1.229	−.022	−.001	.017
$\text{Rev}_{i,j,t}$	4,784	.007	.083	−1.107	1.923	−.025	.004	.033
$\text{Rev}_{i,j,t} > 0$	4,785	.535	.499	0	1	0	1	1
$\text{Rev}_{i,j,t} > .02$	4,785	.346	.476	0	1	0	0	1
$\Delta\text{Spread}_{i,t_0,t_0+2}$	4,772	.335	2.313	−1	80.919	0	0	0
$\widehat{\Delta\text{Spread}}_{i,t_0,t_0+2}$	4,772	.336	2.31	−1	79.229	−.034	−.001	.031
Market value ($millions)	4,785	58,518	104,994	2,013	626,550	4,702	14,078	55,930
Total assets	4,785	96,405	329,839	1.001	2,807,491	3,091	12,548	51,839
Total liabilities	4,736	77,085	301,654	.128	2,736,580	1,807	7,584	33,269
Net income	4,785	3,183	7,685	−14,685	53,394	23.767	373	2,856
Amihud (2002) illiquidity	4,785	.001	.018	0	1.139	0	0	0
Firm-specific volatility	4,785	.024	.017	.005	.309	.014	.02	.028
Article hour	4,785	10.502	4.75	0	23	7	10	14
Year	4,785	2014	2001	2010	2017	2012	2014	2016

Note. Summary statistics are for the continuous variables in the primary sample of 4,785 firm-articles by firms with $2 billion in market capitalization or more.

Table 2. Predictors of Pseudonymous Authorship: Full Sample

Variable	Pseudonym	Real Name	Difference		
			% Bias	*t*-Statistic	$p > \|t\|$
Market value ($millions)	51,195	62,692	−11.3	−3.65	0
Total assets	93,556	98,028	−1.4	−.45	.652
Total liabilities	75,682	77,885	−.7	−.24	.809
Net income	2,740	3,436.1	−9.3	−3.02	.003
Amihud (2002) illiquidity	.00093	.00037	2.7	1.04	.296
Firm-specific volatility	.02404	.02344	3.5	1.15	.25
Article hour	10.503	10.501	0	.01	.993
Year	2,013.9	2,013.9	−3.7	−1.24	.214
Industry:					
Materials	.04147	.03483	3.5	1.17	.244
Capital goods	.04896	.03911	4.8	1.62	.105
Commercial and professional services	.00403	.0069	−3.9	−1.25	.211
Transportation	.02477	.01906	3.9	1.32	.187
Automobiles and components	.02362	.03385	−6.1	−1.99	.047
Consumer durables and apparel	.04205	.03089	6	2.02	.043
Consumer services	.07604	.06901	2.7	.91	.365

Media	.01325	.0253	−8.8	−2.8	.005
Retailing	.10253	.14558	−13.1	−4.26	0
Food and staples retailing	.01037	.01249	−2	−.65	.513
Food, beverage, and tobacco	.04032	.03943	.5	.15	.88
Household and personal products	.01671	.02596	−6.4	−2.07	.038
Health care equipment and services	.02247	.01249	7.6	2.64	.008
Pharmaceuticals, biotechnology, and life sciences	.07028	.03023	18.4	6.46	0
Banks	.02995	.03319	−1.9	−.61	.541
Diversified financials	.0265	.02202	2.9	.98	.327
Insurance	.00634	.00296	5	1.74	.082
Software and services	.1394	.1791	−10.9	−3.56	0
Technology hardware and equipment	.06509	.07756	−4.8	−1.59	.111
Semiconductors and semiconductor equipment	.06336	.05587	3.2	1.06	.289
Telecommunication services	.02247	.03418	−7.1	−2.29	.022
Utilities	.01325	.00526	8.4	2.96	.003
Real estate	.01152	.00789	3.7	1.26	.206

Note. Results are means for univariate predictors of pseudonymous authorship on firm- and article-level covariates prior to propensity-score matching. Industry groups are from the Global Industry Classification Standard.

be slightly smaller and less profitable than real-name targets but indistinguishable in terms of assets and liabilities. There are sector-specific differences; for example, consumer durables and apparel are more likely to be targeted by pseudonymous authors, whereas retailing and software and services are more likely to be targeted by real-name authors. Section 4.1.2 details the use of propensity-score matching to obtain a sample that is balanced on these observable characteristics.

4.1.2. Propensity-Score Matching. A naive comparison of market reactions to pseudonymous and nonpseudonymous articles is subject to the critique that the reactions may be driven by unobserved differences between firms that are the targets of the articles. To be sure, this concern is less compelling in this kind of event-study setting involving high-frequency outcomes like price changes in the days following the publication of a blog post attacking a publicly traded company. To further mitigate selection concerns, I employ a matched design to ensure that I compare firms that are as similar as possible on observable characteristics. I match pseudonymous and nonpseudonymous articles on the following characteristics of firms and articles: market value of equity; total assets; total liabilities; net income; Amihud (2002) illiquidity; the volatility of the firm's stock; Global Industry Classification Standard industry group code; publication hour of the article, which adjusts for time-varying market liquidity conditions; and the year of publication.

I present my results using nearest-neighbor matching, which yields a weighted sample of 2,900 article-firms. A balance test on these covariates is given in Table 3, which shows that the treatment and control groups are balanced across all of these characteristics. A t-test of each of the variables yields p-values that are all above 5 percent, which indicates that the differences in means are not statistically significant. As additional evidence that the two samples are balanced on these characteristics, Figure OA1 in the Online Appendix presents the density of the propensity score between the treatment and control groups for the single-neighbor matching. As Figure OA1 shows, the two groups have very similar densities.

4.2. Main Results

4.2.1. Abnormal Returns and Articles' Publication. I begin my analysis by comparing cumulative abnormal returns related to pseudonymous and nonpseudonymous attacks on public companies. I fit a standard four-factor model of expected returns by estimating for each article in my data set the following ordinary least squares (OLS) regression on daily returns over the interval $[t_0 - 120, t_0 - 7)$ in calendar days (approximately $[t_0 - 85, t_0 - 5)$ in trading days):

$$r_{i,t} - r_{f,t} = \beta_{i,0} + \beta_{i,1} m_t + \beta_{i,2}\mathrm{SMB}_t + \beta_{i,3}\mathrm{HML}_t + \beta_{i,4}\mathrm{UMD}_t + \varepsilon_{i,t},$$

where $r_{i,t}$ is the log return on the common stock of firm i on day t, $r_{f,t}$ is the log risk-free rate on day t, m_t is the log return on the market on day t, SMB_t is the log return on the Fama-French small-minus-big portfolio factor on day t, HML_t is the log return on the Fama-French high-minus-low portfolio factor on day t, UMD_t is the log return on the winners-minus-losers momentum portfolio (Carhart 1997) on day t, and $\varepsilon_{i,t}$ is a random error term.

Next, I obtain daily abnormal log returns by subtracting the predicted values given by this model from the actual returns for each day in the interval $[t_0 - 5, t_0 + 5]$ in trading days:

$$\alpha_{i,t} = r_{i,t} - r_{f,t} - (\beta_{i,0} + \beta_{i,1} m_t + \beta_{i,2}\mathrm{SMB}_t + \beta_{i,3}\mathrm{HML}_t + \beta_{i,4}\mathrm{UMD}_t).$$

Finally, I derive the cumulative abnormal log return from day t to day τ for firm-article i written by author j by summing the daily log abnormal returns:

$$\mathrm{CAR}_{i,j,t,\tau} = \sum_{k=t}^{\tau} \alpha_{i,k}.$$

My results hold with simple returns as well, as shown in the Online Appendix.

Figure 3 plots $\mathrm{CAR}_{i,j,t_0-5,\tau}$ for pseudonymous and nonpseudonymous articles with $\tau \in (t_0 - 5, t_0 + 5]$ in trading days. As Figure 3 shows, articles published by pseudonymous authors are followed by price reversals, which are indicated by the shaded region. Both pseudonymous and nonpseudonymous articles are accompanied by negative cumulative abnormal returns on the order of .01 log point, that is, approximately 1 percentage point. It is clear that most of the decline is concentrated around the publication of an article. In the full sample, the hypothesis that $\mathrm{CAR}_{i,t_0-3,t_0-1}$ is equal to $\mathrm{CAR}_{i,t_0-1,t_0+1}$ around nonpseudonymous articles is rejected at the 1 percent level ($t = -3.77$). And the value for $\mathrm{CAR}_{i,t_0-1,t_0+1}$ is indeed statistically significant at the 1 percent level ($t = -6.24$). In terms of magnitudes, nearly 80 percent of the total $[t_0 - 5, t_0 + 1]$ CAR occurs in the $[t_0 - 1, t_0 + 1]$ period. The results are similar when limited to the matched sample, although significance is only at the 5 percent level (owing to the smaller sample). Thus, the evidence is clear

Table 3. Matched Sample: Balance Test

Variable	Pseudonym	Real Name	Difference % Bias	Difference *t*-Statistic	Difference $p > \|t\|$
Market value ($millions)	51,545	51,576	0	−.01	.992
Total assets	94,452	85,140	2.8	.89	.371
Total liabilities	75,726	66,491	3.1	.98	.329
Net income	2,765.1	2,891.1	−1.7	−.53	.595
Amihud (2002) illiquidity	.00094	.00026	3.3	.98	.326
Firm-specific volatility	.02408	.02476	−4	−1.06	.291
Article hour	10.515	10.49	.5	.15	.878
Year	2013.9	2013.8	3.9	1.13	.258
Fraud-related text	.2936	.33953	−2.3	−.82	.412
Industry:					
Materials	.04186	.04942	−3.9	−1.06	.288
Capital goods	.04942	.03779	5.6	1.67	.095
Commercial and professional services	.00407	.0064	−3.1	−.95	.345
Transportation	.025	.03663	−7.9	−1.97	.048
Automobiles and components	.02384	.02616	−1.4	−.44	.662
Consumer durables and apparel	.04244	.04477	−1.2	−.33	.738
Consumer services	.06919	.07151	−.9	−.27	.79

Media	.01337	.01337	0	0	1
Retailing	.10349	.10465	−.4	−.11	.911
Food and staples retailing	.01047	.00988	.5	.17	.865
Food, beverage, and tobacco	.0407	.0407	0	0	1
Household and personal products	.01686	.01163	3.6	1.29	.195
Health care equipment and services	.02267	.01977	2.2	.59	.554
Pharmaceuticals, biotechnology, and life sciences	.07093	.07151	−.3	−.07	.947
Banks	.03023	.02442	3.3	1.05	.296
Diversified financials	.02674	.03023	−2.3	−.61	.539
Insurance	.0064	.01047	−6	−1.31	.192
Software and services	.13895	.13721	.5	.15	.882
Technology hardware and equipment	.0657	.06686	−.4	−.14	.891
Semiconductors and semiconductor equipment	.06395	.06628	−1	−.28	.782
Telecommunication services	.02267	.02093	1	.35	.726
Utilities	.01337	.0093	4.2	1.13	.26
Real estate	.01163	.01105	.6	.16	.872

Note. Industry groups are from the Global Industry Classification Standard.

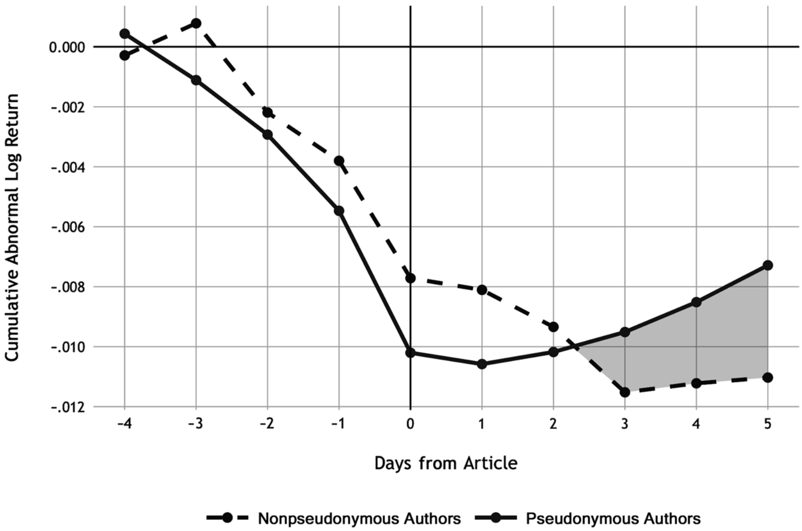

Figure 3. Daily cumulative abnormal log returns

that the magnitude of the price decline over $[t_0 - 1, t_0 + 1]$ far exceeded the decline in the $[t_0 - 5, t_0 - 1]$ period.

There is little difference in the cumulative abnormal log returns between pseudonymous and nonpseudonymous articles in the $[t_0 - 5, t_0 - 1]$ window: returns of both groups experience a roughly parallel minor decline prior to publication of an article.[10] However, returns of firms targeted by pseudonymous articles decline further on the day of publication (t_0) and display a sharp pattern of reversal over the $[t_0 + 2, t_0 + 5]$ window, with returns increasing from $-.0106$ to $-.0073$ from $t_0 + 2$ to $t_0 + 5$, a difference of .33 log point or approximately 31.1 percent in relative terms, from day 1 to day 5 following publication.

There are three possibilities underlying the (relatively minor) price decline prior to publication. One is that the market becomes aware of an attack a few days before the corresponding publication date. As a matter of policy, Seeking Alpha prohibits editors from trading ahead of a forthcoming article. In confidential conversations with Seeking Alpha staff, they insist that news of a future publication does not leak, and there

10. Short attacks are often written in response to prior negative news or other negative sentiment, which is why it is difficult to identify the causal effects of these attacks (Zhao 2018). My design compares pseudonymous with nonpseudonymous articles, where, as Figure 3 shows, the two groups follow roughly parallel, albeit declining, pretrends.

is no evidence that they are mistaken. Furthermore, I have verified with third parties that Seeking Alpha publications generally served as an important source of news for algorithmic trading over these years, and the relevant date for the algorithms was the publication date on Seeking Alpha.[11] A second possibility is that short selling by the author of an article causes the price declines. In my view, this is the most likely explanation of the slight price decline leading up to the article's publication. It is clear that the accumulation of a short position will put downward pressure on the price. While I cannot conclude with certainty that this is taking place without confidential deanonymized data identifying every trade, it is most consistent with the available evidence. Third, it could be that such attacks are timed to follow a period of negative stock price returns. But there is no evidence of that. The $[t_0 - 5, t_0 - 4]$ CAR is 1 basis point and statistically insignificant. The $[t_0 - 5, t_0 - 3]$ CAR is -1 basis point and statistically insignificant as well.

4.2.2. Stock Price Reversals. I test whether pseudonymous articles are followed by greater stock price reversals by estimating several different regression models on my data. I begin by implementing an overreaction measure following Tetlock (2011), regressing the cumulative abnormal return over the $[t_0 + 2, t_0 + 5]$ interval on the cumulative abnormal return over the $[t_0 - 1, t_0 + 1]$ interval. Unlike the average differences in Figure 3, a significant negative coefficient indicates a pair-wise negative correlation between the abnormal returns in the two periods: the farther prices fall after publication of an article, the higher they rise afterward.

This initial specification tests whether price reversals differ between pseudonymous and nonpseudonymous articles. In particular, if the price declines from $t_0 - 1$ to $t_0 + 1$, is the subsequent increase from $t_0 + 2$ to $t_0 + 5$ greater for pseudonymous as opposed to nonpseudonymous articles? A stronger price reversal indicates a higher degree of mispricing—while mispricing does not necessarily prove that manipulation was occurring, it is a necessary condition for manipulation to have occurred. I estimate the following model by OLS regression, employing propensity-score matching to weight matched pairs and exclude unmatched pairs:

$$\begin{aligned}\mathrm{CAR}_{i,j,t_0+2,t_0+5} = \beta_0 + \beta_1 \mathrm{CAR}_{i,t_0-1,t_0+1} + \beta_2 \mathrm{Pseudo}_j \\ + \beta_3(\mathrm{Pseudo}_j \times \mathrm{CAR}_{i,t_0-1,t_0+1}) + \varepsilon_{i,j,t_0+2,t_0+5},\end{aligned}$$

where Pseudo_j equals one if author j is pseudonymous and $\varepsilon_{i,j,t_0+2,t_0+5}$ is a

11. More recently, Seeking Alpha has begun to offer day-before access to upcoming articles to subscribers, but that was not available during the years studied in my sample.

random error term. As an additional measure of price reversal, I consider the simple difference between the cumulative abnormal return over the $[t_0 + 2, t_0 + 5]$ interval and the cumulative abnormal return over the $[t_0 - 1, t_0 + 1]$ interval:

$$\text{Rev}_{i,j,t} = \text{CAR}_{i,j,t_0+2,t_0+5} - \text{CAR}_{i,j,t_0-1,t_0+1}.$$

This measure increases with the divergence between $\text{CAR}_{i,j,t_0+2,t_0+5}$ and $\text{CAR}_{i,j,t_0-1,t_0+1}$. For example, if $\text{CAR}_{i,j,t_0-1,t_0+1} = -.02$ but $\text{CAR}_{i,j,t_0+2,t_0+5} = .04$, then $\text{Rev}_{i,j,t} = .06$. Note that this does not incorporate positive reversals. If $\text{CAR}_{i,j,t_0-1,t_0+1} = .02$ and $\text{CAR}_{i,j,t_0+2,t_0+5} = -.04$, then $\text{Rev}_{i,j,t} = -.06$. However, if $\text{CAR}_{i,j,t_0-1,t_0+1} = .02$ and $\text{CAR}_{i,j,t_0+2,t_0+5} = .8$, then $\text{Rev}_{i,j,t} = .06$. The expression $\text{Rev}_{i,j,t} > 0$ thus corresponds to either a negative reversal (that is, a decline in price followed by a subsequent increase) or a larger increase in price over $[t_0 + 2, t_0 + 5]$ than the increase over $[t_0 - 1, t_0 + 1]$. This latter case is a kind of positive correction in the sense that the increase over $[t_0 - 1, t_0 + 1]$ may have been depressed.

I regress $\text{Rev}_{i,j,t}$ on Pseudo and compare to an indicator equal to 1 if $\text{Rev}_{i,j,t} > 0$. To verify that the results are not driven by price increases, as a robustness check I estimate the first specification but limit the sample to cases in which $\text{CAR}_{i,t_0-1,t_0+1} < 0$.[12] I further compare the sample to cases in which $\text{CAR}_{i,t_0-1,t_0+1} > 0$, where no effect is expected. The results are shown in Table 4.

Table 4 shows that pseudonymous articles are linked to a negative correlation between the postpublication price and the price over the following days: a 1 log point increase in cumulative abnormal log returns in the window $[t_0 - 1, t_0 + 1]$ is followed by a decline of .11 log point of cumulative abnormal returns, on average, in the window $[t_0 + 2, t_0 + 5]$. This coefficient estimate is significant at the 5 percent level. The noninteracted coefficient on $\text{CAR}_{i,t_0-1,t_0+1}$ is positive, which indicates that nonpseudonymous articles are not followed by price reversals using the measure in Tetlock (2011). Table OA4 in the Online Appendix shows that the results are very similar when using simple returns.

Moreover, Table 4 shows that when limiting to negative news where $\text{CAR}_{i,t_0-1,t_0+1} < 0$, the negative correlation is stronger in magnitude and statistical significance: a 1 log point increase in cumulative abnormal log returns in the window $[t_0 - 1, t_0 + 1]$ is followed by a decline of .20 log point of cumulative abnormal returns, on average, in the window $[t_0 +$

12. The standalone reversal measure $\text{Rev}_{i,j,t}$ is negatively correlated with $\text{CAR}_{i,t_0-1,t_0+1}$ by construction, so conditioning on $\text{CAR}_{i,t_0-1,t_0+1} < 0$ is problematic in those specifications.

Table 4. Stock Price Reversals: Pseudonymous versus Nonpseudonymous Articles

		$\text{CAR}_{i,j,t_0+2,t_0+5}$			
	$\text{CAR}_{i,j,t_0+2,t_0+5}$	$\text{CAR}_{i,t_0-1,t_0+1} < 0$	$\text{CAR}_{i,t_0-1,t_0+1} > 0$	$\text{Rev}_{i,j,t}$	$\text{Rev}_{i,j,t} > 0$
$\text{Pseudo} \times \text{CAR}_{i,t_0-1,t_0+1}$	−.1139*	−.1957**	−.0559		
	(−2.43)	(−2.69)	(−.65)		
Pseudo	.0055**	−.0011	.0074*	.0080*	.0465*
	(2.73)	(−.32)	(2.27)	(2.51)	(2.28)
$\text{CAR}_{i,t_0-1,t_0+1}$	.1024**	.1420*	.0795		
	(2.87)	(2.48)	(1.56)		
Intercept	−.0023	.0012	−.0035	.0030	.5047**
	(−1.45)	(.47)	(−1.47)	(1.18)	(30.50)
N	2,899	1,523	1,376	2,899	2,900

Note. All regressions employ propensity-score matching with treatment-control pairs as ordinary least squares weights and robust standard errors; t-statistics are in parentheses.

* $p < .05$.

** $p < .01$.

2, $t_0 + 5$], and this estimate is significant at the 1 percent level. Similarly, Table 4 shows that there is no effect when limiting to positive news where $\mathrm{CAR}_{i,t_0-1,t_0+1} > 0$, which confirms that the effect is not driven by the positive news subsample. Similarly, $\mathrm{Rev}_{i,j,t}$ is .0080 higher for pseudonymous articles on average, a difference that is significant at the 5 percent level, whereas $\mathrm{Rev}_{i,j,t}$ is indistinguishable from 0 for nonpseudonymous articles. Pseudonymous articles are 9.2 percent more likely to be followed by a positive $\mathrm{Rev}_{i,j,t}$ of any magnitude (this is obtained by dividing the coefficient .0465 by the intercept term .5047) and nearly 13.2 percent more likely to be followed by a positive $\mathrm{Rev}_{i,j,t}$ exceeding 2 log points in magnitude (.0430/.3262 $\approx$ .132). This statistical evidence is consistent with the pattern displayed in Figure 3.

To be sure, Figure 3 shows that prices do not fully reverse. But prices need not fully reverse for there to be market manipulation. As discussed in the Online Appendix, the courts have held that short selling may constitute market manipulation in violation of the Securities Exchange Act of 1934 when it is "willfully combined with something more to create a false impression of how market participants value a security" (*ATSI Communications Inc. v. Shaar Fund Ltd.,* 493 F.3d 87, 101 [2d Cir. 2007]). A partial reversal is expected when an author releases some truthful information alongside giving a false impression, as appears to be the case with pseudonymous short attacks on Seeking Alpha.

4.2.3. Pseudonymity and Reputation. A straightforward prediction of the theoretical model is that pseudonymous identities are likely to engage in informed or manipulative trading so long as the market cannot conclude that they are lying. For an anecdotal example of a pseudonym proudly emphasizing a history of successes, consider again the case of SkyTides and Insulet. Figure 4 is from SkyTides' website and shows the short seller's history of success prior to Insulet. It is clear that the market was justified in listening to SkyTides, as the pseudonymous author had established a track record as a nonliar and a history of successful nonreversals prior to attacking Insulet while purchasing put options prior to publication of the article.

It is clear that the market is likely to believe a pseudonym in two types of cases: cases in which an author's prior predictions have historically yielded nonreversals, on average, and cases in which the author has no history. The latter is less intuitive but can be seen directly in the theoretical model. Suppose that, on encountering a pseudonymous author for the first time, the market were to disbelieve the author. Then by definition, in

Company	Ticker	Rating	Date Report Released	$ Stock Price at Report Release	$ Stock Low Price Since Report Release	Stock Price % Decline	Officer or Board Departures
Insulet Corporation	PODD	Sell Short	11/15/16	$ 34.63	$ 32.12	-7.2%	None
Vocera Communications, Inc.	VCRA	Sell Short	4/5/16	$ 12.23	$ 10.46	-14.5%	Board Members
Conformis, Inc.	CFMS	Sell Short	1/19/16	$ 13.39	$ 3.89	-70.9%	CEO
Tantech Holdings Ltd	TANH	Sell Short	9/24/15	$ 21.39	$ 1.00	-95.3%	None
HII Technologies, Inc.	HIIT	Sell Short	12/30/13	$ 0.52	$ 0.00	-99.8%	CEO, CFO
Green Automotive Company	GACR	Sell Short	12/19/13	$ 0.25	$ 0.00	-99.1%	CEO, CFO

Figure 4. SkyTides' history of nonliar trading

a Bayesian model, the market will not believe any future message by that author.

Intuitively, in a Bayesian model, posterior beliefs are a weighted product of prior beliefs and new information. If my prior belief is that there is 0 probability that a never-before-seen pseudonym is telling the truth, that will be my posterior belief as well. Therefore, if markets react to pseudonymous authors, they must believe that it is possible that never-before-seen authors are telling the truth. There is also ample anecdotal evidence that markets respond to first-time authors—in my sample, 182 articles by first-time authors had an abnormal return below -3 percent from $t_0 - 1$ to $t_0 + 1$. An author who has never appeared before clearly cannot be a liar in the view of the market, but once having been proved as such, it is rational to ignore him or her going forward—forever trapping the author in the curse of disbelief. Being proved wrong, on average, breaks the pooling equilibrium, which allows the market to conclude that an author is likely a fraudster.

I systematically test the hypothesis that pseudonymous authors exploit the market's inferences as to the prior truth or falsity of their statements by defining, for each firm-article, a nonliar as a prior cumulative sum of nonreversals or the absence of any reputational history (that is, the author's first article on Seeking Alpha). I define the following variable for the article published by author j about firm i at time t:

$$\text{Nonliar}_{i,j,t} = \begin{cases} 1 & \text{if } t = 0 \quad \text{or} \quad \sum_{\tau=0}^{t} \text{Rev}_{k,\tau} < 0 \;\; \forall \; k \in J_{k,t} \\ 0 & \text{otherwise} \end{cases},$$

where $J_{k,t}$ is the set of articles written by author j prior to time t, with each article indexed by k. I estimate the same regression model as in Section 4.2.2, comparing the samples in which $\text{Nonliar}_{i,j,t}$ equals one or zero. The results are presented in Table 5.

Table 5. Stock Price Reversals: Pseudonymous Authors by Period Type

	Nonliar Periods				Liar Periods			
	CAR_{i,j,t_0+2,t_0+5}	$Rev_{i,j,t}$	$Rev_{i,j,t} > 0$	$Rev_{i,j,t} > .02$	CAR_{i,j,t_0+2,t_0+5}	$Rev_{i,j,t}$	$Rev_{i,j,t} > 0$	$Rev_{i,j,t} > .02$
Pseudo	.0047*	.0077*	.0567*	.0483*	.0025	.0038	−.0118	−.0006
	(2.00)	(2.07)	(2.38)	(2.02)	(.62)	(.75)	(−.34)	(−.02)
Pseudo × CAR_{i,t_0-1,t_0+1}	−.1351*				.0252			
	(−2.40)				(.26)			
CAR_{i,t_0-1,t_0+1}	.1426**				.0782$^+$			
	(3.17)				(1.83)			
Intercept	.0032$^+$	.0173**	.5887**	.4080**	−.0137**	−.0323**	.2978**	.1247**
	(1.73)	(5.80)	(30.03)	(21.10)	(−4.62)	(−8.05)	(11.10)	(6.61)
N	2,107	2,107	2,107	2,107	792	792	793	793

Note. All regressions employ robust standard errors; *t*-statistics are in parentheses.
$^+ p < .10$.
$^* p < .05$.
$^{**} p < .01$.

Table 5 shows that negative reversals occur when authors are perceived as nonliars by the market. I next examine whether differences in informed trading are driven by trading at nonliar times.[13] I estimate the same model as in Section 4.3.1, comparing the samples in which $\text{Nonliar}_{i,j,t}$ equals one or zero. The results are given in Table 6.

As Table 6 shows, informed trading in options markets is occurring when pseudonymous authors are perceived as nonliars by the market. It is possible that the lack of statistical significance in the liar subsample is driven by insufficient power, but notice that these samples have 226,940 and 86,755 observations. Taken together, this evidence is consistent with the theoretical prediction in Benabou and Laroque (1992) that informed trading will be concentrated where the market perceives an author to be a nonliar.

A second implication of this theoretical framework is that pseudonymous authors should switch identities once the market realizes they are promulgating misleading articles. I examine whether pseudonymous authors are more likely to disappear after it is apparent that the market is no longer listening to what they have to say. I test three distinct propositions.

First, I examine whether pseudonymous authors are more likely to disappear, that is, whether a given article is likely to be the last one written by an author. I estimate the following regression on the matched sample:

$$\text{Last}_{i,j,t} = \beta_0 + \beta_1 \text{Pseudo}_j + \varepsilon_{i,j,t},$$

where $\text{Last}_{i,t}$ equals one if article i written at time t is the last one by author j, Pseudo_j equals one if author j is pseudonymous, and $\varepsilon_{i,j,t}$ is a random error term.

Second, I test whether the market response to an article is linked to the presence or absence of prior reversals. For each author I derive the mean of prior negative reversals, which is based on the same metric used to determine nonliar periods:

$$\text{Prior}_{i,j,t} = \frac{1}{N_{J_{k,t}}} \sum_{\tau=0}^{t} \text{Rev}_{k,\tau} < 0 \quad \forall\ k \in J_{k,t},$$

where $J_{k,t}$ is the set of articles written by author j prior to time t, with each article indexed by k, and $N_{J_{k,t}}$ denotes the length of $J_{k,t}$. I define the market response to an article as $|\text{CAR}_{i,j,t_0-1,t_0+1}|$—that is, the absolute abnor-

13. Volume yields qualitatively similar but noisier estimates, as expected with many days having no volume.

Table 6. Pseudonymous Attacks and Options Trading by Period Type

	Nonliar Periods		Liar Periods	
	(1)	(2)	(3)	(4)
Pseudo × Pub × Call	−.0782*		−.0427	
	(−2.12)		(−.70)	
Pseudo × Correction Period × Call		.0945**		.0548
		(3.09)		(1.19)
Pseudo × Pub	.0368		.0655	
	(1.24)		(1.29)	
Pseudo × Call	−.0401	−.0524	.0422	.0228
	(−1.02)	(−1.09)	(.69)	(.32)
Pub × Call	−.0148		.1150*	
	(−.52)		(2.47)	
Pseudo × Correction Period		−.0337		−.1001**
		(−1.40)		(−2.88)
Correction Period × Call		.0013		−.1101**
		(.05)		(−3.18)
Pub	.0070		−.1048**	
	(.29)		(−2.61)	
Correction Period		.0070		.1374**
		(.37)		(5.30)
Call	.4786**	.4151**	.5335**	.6679**
	(15.87)	(10.81)	(12.38)	(11.98)
Intercept	18.6586**	22.0316**	17.4739**	19.5176**
	(7.85)	(6.02)	(4.06)	(3.04)
N	418,650	226,940	159,731	86,755

Note. Results are from a difference-in-difference-in-differences model that employs firm-article fixed effects and controls for strike price, absolute delta, gamma, and time to expiration. The outcome is the log open interest for a nearly at-the-money option with an absolute delta between .45 and .55. Robust standard errors are clustered by firm-article; *t*-statistics are in parentheses.

* $p < .05$.

** $p < .01$.

mal return over the interval $[t_0 - 1, \ t_0 + 1]$[14]—and estimate the following OLS regression on the matched sample:

$$|\mathrm{CAR}_{i,j,t_0-1,t_0+1}| = \beta_0 + \beta_1 \mathrm{Last}_{i,j,t} + \varepsilon_{i,j,t},$$

where $\varepsilon_{i,j,t}$ is a random error term.

Finally, I link the two prior tests and consider whether pseudonymous authors are more likely to disappear when the market has ceased to respond to the publication of an article. I estimate the following OLS regression on the matched sample:

$$\begin{aligned}\mathrm{Last}_{i,j,t} = {} & \beta_0 + \beta_1 \mathrm{Pseudo}_j + \beta_2 \,|\mathrm{CAR}_{i,j,t_0-1,t_0+1}| \\ & + \beta_3 (|\mathrm{CAR}_{i,j,t_0-1,t_0+1}| \times \mathrm{Pseudo}_j) + \varepsilon_{i,j,t}.\end{aligned}$$

In addition, I consider a robustness check in which I define a variable that reflects a lack of credibility for a pseudonymous author using arbitrary cutoffs:

$$\text{Low-Credibility Pseudo}_{i,j,t} = \begin{cases} 1 & \text{if } |\mathrm{CAR}_{i,j,t_0-1,t_0+1}| < .01, \\ & \quad \mathrm{Prior}_{i,j,t} > .05, \text{ and } \mathrm{Pseudo}_j = 1. \\ 0 & \text{otherwise} \end{cases}$$

These cutoffs are simply another way to measure the broader patterns identified in prior specifications, and the results are not sensitive to the choice of this particular cutoff. The results of these estimations are given in Table 7.

Column 1 of Table 7 shows that the last article for an author is more likely to be written by a pseudonymous author than a real-name author, which is consistent with pseudonymous authors switching identities. Column 2 shows that the market response to a given article decreases as the author accumulates a history of negative reversals. Column 3 shows that the last article for an author is especially likely to have been written by a pseudonymous author with a low market response—the negative coefficient indicates that for pseudonymous authors, the probability of the last article increases as the market response to the article decreases. Column 4 shows that low-credibility articles by pseudonymous authors are extremely likely to be the last articles written by those authors. To ensure that these results are not driven by the choice of a cutoff, I repeated the analysis with other cutoffs, and the results are statistically significant and

14. The closer $|\mathrm{CAR}_{i,j,t_0-1,t_0+1}|$ is to 0, the less stock prices change in response to the publication of the article.

Table 7. Pseudonymous Attacks and Disappearing Authors

	$\text{Last}_{i,j,t}$ (1)	$\lvert\text{CAR}_{i,j,t_0-1,t_0+1}\rvert$ (2)	$\text{Last}_{i,j,t}$ (3)	$\text{Last}_{i,j,t}$ (4)
Pseudo	.0390*		.0242	
	(2.42)		(1.41)	
$\text{Prior}_{i,j,t}$		−.2497**		
		(−3.96)		
Pseudo × $\lvert\text{CAR}_{i,j,t_0-1,t_0+1}\rvert$			−.0838*	
			(−2.50)	
$\lvert\text{CAR}_{i,j,t_0-1,t_0+1}\rvert$			.0396	
			(1.37)	
Low-Credibility Pseudo				.7606**
				(27.54)
Intercept	.1872**	.0336**	.1936**	.1925**
	(14.93)	(30.59)	(14.07)	(24.07)
N	2,900	2,601	2,900	2,900

Note. Regressions are estimated with robust standard errors on the matched sample, with treatment-control pairs using ordinary least squares regression weights; *t*-statistics are in parentheses.

* $p < .05$.

** $p < .01$.

similar in economic magnitude.[15] Taken together, this evidence is consistent with the theoretical prediction that pseudonymous authors disappear when they lose credibility.[16]

4.3. Extensions

4.3.1. Informed Trading in Options Markets. Fox, Glosten, and Rauterberg (2018, p. 112) point out that a "misstatement manipulator makes his purchases on the basis of something that he knows and the market does not: the falsity of the price-depressing misstatement for which he is responsible." A large literature finds that informed traders exploit their informational advantages in options markets. Chakravarty, Gulen, and Mayhew (2004) show that options markets contribute 17 percent to price discovery. Future stock returns can be predicted both by options volume (Pan and Poteshman 2006) and by deviations from put-call parity

15. For example, defining $\text{Pseudo}_{i,j,t} = 1$ if $\lvert\text{CAR}_{i,j,t_0-1,t_0+1}\rvert < .005$, $\text{Prior}_{i,j,t} > .025$, and $\text{Pseudo}_j = 1$ yields a coefficient estimate of .6500, which is significant at the 1 percent level.

16. In unreported estimations, I reran the analysis in column 4 on nonpseudonymous articles and found a similar result. While columns 1–3 are significantly different for pseudonymous articles, this suggests that when authors have truly lost credibility, they cease posting regardless of whether they are pseudonymous.

(Cremers and Weinbaum 2010). And Mitts and Talley (2019) find that put-options-trading volume and open interest rise in the months preceding the disclosure of a cybersecurity breach.

Are these price reversals driven by this kind of manipulative informed buying at prices that have been artificially depressed by the publication of a pseudonymous attack article? A measure of bullish or bearish sentiment in options markets is the relative demand for put versus call options, which has been found to predict informed trading (for example, Pan and Poteshman 2006). Augustin et al. (2016, p. 3) show that trading costs make it very expensive for informed traders to trade options that are far out of the money, which leads them to conclude that informed investors will "trade options that are only slightly" out of the money. Similar to Augustin, Brenner, and Subrahmanyam (2016), I examine trading behavior over these windows using individual quotes for options that are nearly at the money (delta between .45 and .55) in the OptionMetrics IvyDB for each of the firm-articles in the single-neighbor matched sample ($n = 992{,}946$). My results are qualitatively similar when options that are deeply out of the money are included.

4.3.2. Open Interest and Volume. It can be difficult to measure informed trading in options markets, so I consider multiple approaches. Prior literature shows the ratio of demand for put options to call options predicts future stock returns (Pan and Poteshman 2006), and studies measure this demand with abnormal open interest, the number of outstanding open put or call contracts, and transaction volume (Cao, Chen, and Griffin 2005; Chakravarty, Gulen, and Mayhew 2004; Jayaraman, Frye, and Sabherwal 2001). Augustin, Brenner, and Subrahmanyam (2016) identify informed options trading prior to takeover announcements using abnormal volume in call options written on the target's stock. Accordingly, I employ a difference-in-difference-in-differences design that compares the difference over time in the open interest and volume of put versus call options between pseudonymous and nonpseudonymous articles prior to and following two periods: the date of disclosure (t_0) and the reversal period $[t_0 + 2, t_0 + 5]$. A total of 92.3 percent of the firm-events in the matched sample have options-trading data for these periods.

I begin by plotting time trends on the average difference in log open interest between pseudonymous and nonpseudonymous articles in Figure 5 (that is, $y_{i,t} = a_{i,t} - n_{i,t}$, where $a_{i,t}$ is log open interest for pseudonymous articles and $n_{i,t}$ is log open interest for nonpseudonymous articles)

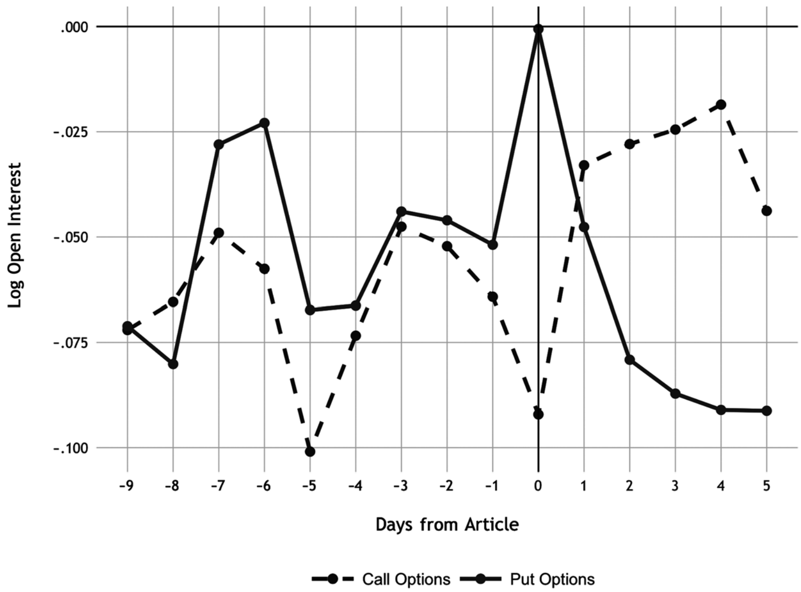

Figure 5. Difference in log open interest

for put options and call options, separately, after subtracting the average log open interest for call options and put options written on each firm-article in the interval $[t_0 - 9, t_0 + 5]$ (that is, a fixed-effect specification). To examine whether the parallel-trends assumption holds, I begin Figure 5 at $t_0 - 9$.

As Figure 5 shows, the trends are roughly parallel over the interval $[t_0 - 9, t_0 - 1]$. At t_0, the demand for put options skyrockets, which suggests that some of the price decline on the day of publication may be driven by highly leveraged option trades on that day. This trend flips direction immediately thereafter: from $t_0 + 2$ to $t_0 + 5$, the period during which prices reverse direction, the demand for call options exceeds the demand for put options. Figure OA2 in the Online Appendix shows a similar pattern for log volume.

This evidence suggests a manipulative pattern of informed buying exploiting the negative reaction to an article, which causes prices to reverse direction during the interval $[t_0 + 2, t_0 + 5]$. To examine this statistically, I estimate two different models by OLS regression:

$$\begin{aligned} y_{i,j,t} = {} & \beta_0 + \beta_1 \mathrm{Pub}_t + \beta_2 \mathrm{Call}_j + \beta_3(\mathrm{Pseudo}_j \times \mathrm{Pub}_t) + \beta_4(\mathrm{Pseudo}_j \times \mathrm{Call}_j) \\ & + \beta_5(\mathrm{Pub}_t \times \mathrm{Call}_j) + \beta_6(\mathrm{Pseudo}_j \times \mathrm{Pub}_t \times \mathrm{Call}_j) + \alpha_i + \varepsilon_{i,j,t} \end{aligned}$$

and

$$y_{i,j,t} = \beta_0 + \beta_1 \text{Post}_t + \beta_2 \text{Call}_j + \beta_3(\text{Pseudo}_j \times \text{Post}_t) + \beta_4(\text{Pseudo}_j \times \text{Call}_j) + \beta_5(\text{Post}_t \times \text{Call}_j) + \beta_6(\text{Pseudo}_j \times \text{Post}_t \times \text{Call}_j) + \alpha_i + \varepsilon_{i,j,t},$$

where $y_{i,j,t}$ is log open interest or volume on day t for option j written on the stock of the firm that is the subject of article i, Pseudo_j equals one if the article was published by a pseudonymous author, Pub_t equals one if day $t = t_0$, Post_t equals one if day t lies within the correction period $[t_0 + 2, t_0 + 5]$, Call_j equals one if the option is a call option, α_i is a fixed effect for firm-article i, and $\varepsilon_{i,j,t}$ is a random error term. The coefficient of interest is β_6, which captures the difference in open interest or volume between call and put options for pseudonymous articles on publication day t_0 or during the correction period $[t_0 + 2, t_0 + 5]$. Standard errors are clustered by firm-article, and the results are presented in Table 8.

As Table 8 shows, the triple-difference coefficient β_6 is negative and statistically significant in the publication-day specification (Pseudo $\times$ Pub $\times$ Call). Columns 1 and 3 show that the open interest and volume of a call option written on the target of a pseudonymous article are 7.66 and 7.75 log points lower, respectively, than put options on the day of publication. Similarly, the triple-difference coefficient β_6 is positive and statistically significant in the correction-period specification (Pseudo $\times$ Correction Period $\times$ Call). Columns 2 and 4 show that the open interest and volume of a call option written on the target of a pseudonymous article are 8.92 and 6.20 log points higher, respectively, than put options during the correction period compared with the day of disclosure.

Can market participants detect manipulative options trading during this period? Table 9 reports the results of estimating this triple-difference model on observable characteristics of these options: the time to expiration, strike price, absolute delta, and gamma. Table 9 shows that the triple-difference coefficient is statistically indistinguishable from 0 on both the day of the article's publication and the reversal day. This indicates that informed trading on the knowledge of a forthcoming manipulative short attack is occurring among options that are observationally similar; pseudonymous authors are not tipping their hand by trading in options that expire more quickly or are otherwise unusual.

4.3.3. Put-Call Parity. Cremers and Weinbaum (2010) find that deviations from put-call parity predict stock returns, which suggests the presence of informed trading. They estimate these deviations by measuring

Table 8. Pseudonymous Attacks and Options Trading

	Open Interest		Volume	
	(1)	(2)	(3)	(4)
Pseudo × Pub × Call	−.0766*		−.0775*	
	(−2.41)		(−2.23)	
Pseudo × Correction Period × Call		.0892**		.0620*
		(3.49)		(2.04)
Pseudo × Pub	.0489+		−.0006	
	(1.90)		(−.02)	
Pseudo × Call	−.0210	−.0446	.0130	−.0277
	(−.64)	(−1.10)	(.54)	(−.86)
Pub × Call	.0251		.0548*	
	(1.02)		(1.97)	
Pseudo × Correction Period		−.0560**		−.0077
		(−2.80)		(−.25)
Correction Period × Call		−.0316		−.0171
		(−1.59)		(−.72)
Pub	−.0254		.0834**	
	(−1.22)		(2.72)	

Correction Period		.0451**		−.0887**
		(2.85)		(−3.77)
Call	.4954**	.4918**	.4763**	.5026**
	(19.97)	(15.24)	(24.89)	(19.49)
| Delta |	−2.1201**	−2.4260**	−2.4397**	−2.4949**
	(−23.42)	(−21.50)	(−26.71)	(−21.17)
Time to expiration	−.0006**	−.0006**	−.0031**	−.0031**
	(−8.03)	(−6.92)	(−37.97)	(−34.20)
Log of strike price	−1.0745**	−1.3208**	−1.3427**	−1.4128**
	(−5.72)	(−4.69)	(−8.51)	(−5.33)
Gamma	−1.1147**	−1.2278**	2.8839**	2.9693**
	(−9.88)	(−10.50)	(15.57)	(16.90)
Intercept	18.3620**	21.2865**	19.9068**	20.7935**
	(8.72)	(6.74)	(11.21)	(6.96)
N	578,381	313,695	403,609	219,923

Note. Results are from a difference-in-difference-in-differences model with firm-article fixed effects. The data set consists of options that are nearly at the money written on firms in the matched sample with an absolute delta between .45 and .55. Robust standard errors are clustered by firm-article; *t*-statistics are in parentheses.

$^{+}p < .10$.

$^{*}p < .05$.

$^{**}p < .01$.

Table 9. Pseudonymous Attacks and Options Characteristics

	Time to Expiration	Strike Price	| Delta |	Gamma
Publication day ($N = 686{,}809$):				
Pseudo × Pub × Call	−.0470	−.0008	.0001	.0007^{+}
	(−.05)	(−1.16)	(.11)	(1.66)
Pseudo × Pub	−2.3974^{+}	.0072^{+}	−.0056*	.0001
	(−1.71)	(1.91)	(−1.97)	(.51)
Pseudo × Call	.1597	.0023*	−.0000	−.0005
	(.26)	(2.40)	(−.24)	(−1.64)
Pub × Call	.1860	.0006	−.0001	−.0004
	(.23)	(1.20)	(−.20)	(−1.41)
Pub	1.6128	−.0061^{+}	−.0000	.0004
	(1.32)	(−1.86)	(−.20)	(.64)
Call	−.0207	−.0061**	.0003^{+}	.0023**
	(−.04)	(−7.62)	(1.89)	(10.06)
Intercept	186.8216**	11.2191**	.4999**	.0599**
	(1,192.03)	(39,049.44)	(10,458.51)	(715.55)

Correction period ($N = 369{,}825$):				
Pseudo × Pub × Call	−.7423	.0003	.0001	.0001
	(−.85)	(.55)	(.20)	(.43)
Pseudo × Correction Period	2.2532^{+}	.0058**	−.0000	−.0003
	(1.83)	(2.77)	(−.07)	(−.38)
Pseudo × Call	.8266	.0019^{+}	.0001	−.0002
	(.94)	(1.66)	(.21)	(−.61)
Correction Period × Call	.4804	−.0001	−.0004	−.0000
	(.72)	(−.18)	(−1.09)	(−.08)
Correction Period	−.7897	−.0022	.0002	−.0001
	(−.82)	(−1.49)	(1.15)	(−.17)
Call	−.4832	−.0060**	.0004	.0022**
	(−.69)	(−6.73)	(1.46)	(7.31)
Intercept	189.0077**	11.2178**	.4997**	.0612**
	(445.73)	(15,235.01)	(5,333.84)	(239.18)

Note. Results are from a difference-in-difference-in-differences model with firm-article fixed effects. Robust standard errors are clustered by firm-article; t-statistics in parentheses.
$^{+}$ $p < .10$.
* $p < .05$.
** $p < .01$.

the difference in implied volatility between put and call options with the same strike price and expiration date. An et al. (2014) find that changes in implied volatility predict future stock returns. I examine whether deviations from put-call parity predict informed trading during the period of a pseudonymous attack by matching put and call options for a given security on expiration date and strike price, and I also consider whether implied volatility differs between the matched put options and call options, as in Cremers and Weinbaum (2010).

I estimate the same triple-difference specification but replace the firm-event fixed effect α_i with a fixed effect corresponding to the unique combination of the underlying security, expiration date, and strike price. First, I compare the period $[t_0, t_0 + 2]$, when put-call parity should reflect informed trading in the direction of put options, with the baseline period $[t_0 - 9, t_0 - 1]$. Second, I compare the reversal period $[t_0 + 3, t_0 + 5]$, when put-call parity should return to the baseline, with the elevated period $[t_0, t_0 + 2]$. The estimations are lagged by 1 day in contrast to Table 8 to account for options markets updating in response to informed order flow. The results are presented in Table 10.

As Table 10 shows, implied volatility is higher for put options relative to call options written on the targets of pseudonymous attacks over the window $[t_0, t_0 + 2]$. Similarly, implied volatility is higher for call options relative to put options over the lagged reversal period $[t_0 + 3, t_0 + 5]$. Like Cremers and Weinbaum (2010) and An et al. (2014), this deviation from put-call parity indicates the presence of informed trading in options markets over these windows. In Sections 4.3.4 and OA1 in the Online Appendix, I measure the resulting increase in bid-ask spreads and show that these price reversals are not driven by provocative content.

4.3.4 Bid-Ask Spreads. Glosten and Milgrom (1985) show that the presence of informed trading causes market makers to enlarge bid-ask spreads to compensate for expected trading losses. I examine whether spreads widen for the targets of pseudonymous attacks. Both the acquisition of put options on the day of the attack and the accumulation of long positions during the correction period constitute a kind of informed trading on the fact that an article does not have implications for the fundamental value of the firm. However, market makers are able to anticipate only the latter, because the publication of the article comes as a surprise to the market. In anticipation of the accumulation of call options during the correction period (which will be hedged by counterparties opening long positions in the underlying stock), market makers are likely to widen the spread. Figure 3 suggests that this occurs at $t_0 + 2$, so a natural starting point is to ask whether bid-ask spreads increase from the day of publication to 2 days after, when informed traders will aggressively begin to purchase the shares of target firms whose stock prices were artificially depressed by the pseudonymous attack.

I measure bid-ask spreads using daily pricing data reported by the Center for Research on Securities Prices. These data are rough approximations but useful for daily analysis of this kind. Bid-ask spreads are highly persistent, so over-time variations in spreads tend to be multiplicative in nature. A firm with a small spread of \$.01 is extremely unlikely to see its spread double to \$.02, even with substantially increased informed trading, but a firm with a spread of \$.20 could easily see that spread increase to \$.21. This motivates the following percentage definition of the change in the spread:

$$\Delta\text{Spread}_{i,t,\tau} = \frac{\text{Spread}_{i,t}}{\text{Spread}_{i,\tau}} - 1.$$

An alternative normalizes the spread by the price of the underlying stock:

$$\Delta\widehat{\text{Spread}}_{i,t,\tau} = \frac{\text{Spread}_{i,t}/p_{i,t}}{\text{Spread}_{i,\tau}/p_{i,\tau}} - 1.$$

For the reasons described above, I focus on $\Delta\text{Spread}_{i,t_0,t_0+2}$ and $\Delta\widehat{\text{Spread}}_{i,t_0,t_0+2}$, that is, the percentage change in the spread from the day of publication to 2 days thereafter.

In a competitive market among liquidity providers, market makers will increase the spread commensurately with the risk of informed trading. As of day $t_0 + 2$, market makers observe the extent of the price decline over the interval $[t_0 - 1, t_0 + 1]$, and the analysis in Table 4 indicates that this price decline is a key proxy for the expected reversal. For this reason, I estimate the following model using OLS regression, employing propensity-score matching at the firm-article level to weight matched pairs and exclude unmatched pairs:

$$\begin{aligned}\Delta\text{Spread}_{i,t_0,t_0+2} = {} & \beta_0 + \beta_1\text{Pseudo}_j + \beta_2\text{CAR}_{i,t_0-1,t_0+1} \\ & + \beta_3(\text{Pseudo}_j \times \text{CAR}_{i,t_0-1,t_0+1}) + \varepsilon_{i,t,\tau}.\end{aligned}$$

The key coefficient of interest is β_3, which reflects the percentage-point change in the spread with the change in the cumulative abnormal log return over the interval $[t_0 - 1, t_0 + 1]$. The prediction is that $\beta_3 < 0$; that is,

Table 10. Pseudonymous Attacks and Put-Call Parity

	(1)	(2)	(3)	(4)
Pseudo × [t_0, t_0 + 2] × Call	−.0018**	−.0017**		
	(−3.68)	(−3.48)		
Pseudo × [t_0 + 3, t_0 + 5] × Call			.0023**	.0022**
			(4.60)	(4.67)
Pseudo × [t_0, t_0 + 2]	.0028**	.0023**		
	(3.94)	(3.42)		
Pseudo × Call	.0047**	.0045**	.0030**	.0029**
	(11.68)	(11.40)	(5.41)	(5.45)
[t_0, t_0 + 2] × Call	.0007*	.0006+		
	(1.96)	(1.73)		
Pseudo × [t_0 + 3, t_0 + 5]			.0002	−.0000
			(.30)	(−.04)
[t_0 + 3, t_0 + 5] × Call			−.0007+	−.0006+
			(−1.74)	(−1.68)
[t_0, t_0 + 2]	.0003	.0007		
	(.50)	(1.52)		
[t_0 + 3, t_0 + 5]			−.0024**	−.0010*
			(−5.07)	(−2.19)

Call	−.0129**	−.0125**	−.0116**	−.0112**
	(−33.39)	(−33.27)	(−23.96)	(−23.73)
| Delta |		.0020		−.0021
		(.88)		(−.70)
Gamma		−.3565**		−.4393**
		(−14.31)		(−9.90)
Vega		.0003**		.0003**
		(7.99)		(5.78)
Theta		−.0005**		−.0004**
		(−14.93)		(−9.42)
Intercept	.4238**	.4205**	.4171**	.4248**
	(633.47)	(146.05)	(438.11)	(104.66)
N	813,599	813,599	369,825	369,825

Note. Results are from a difference-in-difference-in-differences model with fixed effects for put and call options on the same underlying security, expiration date, and strike price. The data set consists of options that are nearly at the money written on firms in the matched sample with an absolute delta between .45 and .55. Robust standard errors are clustered by the unique combination of security, expiration date, and strike price; *t*-statistics are in parentheses.

$^{+}p < .10$.

$^{*}p < .05$.

$^{**}p < .01$.

as $\text{CAR}_{i,t_0-1,t_0+1}$ declines (becomes more negative), the spread increases. I also consider two alternative regressors: an indicator that equals one if $\text{CAR}_{i,t_0-1,t_0+1} < 0$, which indicates a negative market reaction to article publication, and an indicator that equals one if $\text{CAR}_{i,t_0-1,t_0+1} < -.05$, which is a strongly negative market reaction. The results are given in Table 11.

Columns 1 and 4 of Table 11 show that, on average, a 1 log point decrease in the cumulative abnormal log return over the interval $[t_0 - 1, t_0 + 1]$ is linked to an increase of 2.18 to 2.24 percentage points in the bid-ask spread from the day of publication to 2 days after publication, when the informed call options trading is expected to commence. Similarly, spreads increase by approximately 43–45 percentage points for pseudonymous articles, with a decline in the cumulative abnormal log return over the interval $[t_0 - 1, t_0 + 1]$ and an increase of 50–52 percentage points for $\text{CAR}_{i,t_0-1,t_0+1} < -.05$, that is, a strongly negative market reaction. This evidence is consistent with the concerns raised in Fox, Glosten, and Rauterberg (2018) that this sort of market manipulation constitutes a form of informed trading that imposes social welfare costs by widening the bid-ask spread.

4.3.5. Aggregate Trading Losses. What are the aggregate trading losses due to the mispricing caused by the publication of pseudonymous articles? It is important not to confuse these trading losses, which are merely ex post transfers between traders, with the welfare costs of informed trading. Those welfare losses are driven by the reduction in liquidity and increase in the bid-ask spread as a result of pseudonymous market manipulation (Glosten and Putniņš 2016); here I simply compute the extent to which trades were executed at an incorrect price ex post.

I consider solely the 1,720 firm-articles written by pseudonymous authors and calculate the aggregate dollar volume of trading on each of the trading days from $[t_0 , t_0 + 4]$ and exclude day $t_0 + 5$ because that is used to calculate the counterfactual price. I then calculate the counterfactual dollar volume by multiplying the number of shares that were traded for each firm by the price of the firm on $t_0 + 5$. This is the price that sellers would have received in the absence of any price distortion, that is, if the shares had been sold at their price on day $t_0 + 5$. To calculate net mispricing, I subtract the actual dollar volume from the counterfactual dollar volume, which measures the price sellers would have received if the counterfactual price at $t_0 + 5$ had prevailed over those days. The price

Table 11. Pseudonymous Attacks and Bid-Ask Spreads

	$\Delta \text{Spread}_{i,t,T}$			$\Delta \widehat{\text{Spread}}_{i,t,\tau}$		
	(1)	(2)	(3)	(4)	(5)	(6)
Pseudo $\times$ $\text{CAR}_{i,t_0-1,t_0+1}$	-2.1793^{**}			-2.2366^{**}		
	(2.89)			(2.92)		
Pseudo $\times$ $\text{CAR}_{i,t_0-1,t_0+1} < 0$		$.4352^{**}$			$.4479^{**}$	
		(2.48)			(2.55)	
Pseudo $\times$ $\text{CAR}_{i,t_0-1,t_0+1} < -.05$			$.5037^{*}$			$.5180^{*}$
			(2.46)			(2.39)
Pseudo	$-.0789$	$-.1821$	$-.1242$	$-.0824$	$-.1877^{+}$	$-.1286$
	$(-.91)$	(-1.62)	(-1.31)	$(-.96)$	(-1.69)	(-1.38)
$\text{CAR}_{i,t_0-1,t_0+1}$	1.3319^{**}			1.1158^{*}		
	(2.65)			(2.22)		
$\text{CAR}_{i,t_0-1,t_0+1} < 0$		$-.2897^{*}$			$-.2626^{*}$	
		(-2.56)			(-2.33)	
$\text{CAR}_{i,t_0-1,t_0+1} < -.05$			$-.2210^{+}$			$-.1828$
			(-1.96)			(-1.62)
Intercept	$.3354^{**}$	$.4074^{**}$	$.3539^{**}$	$.3373^{**}$	$.4031^{**}$	$.3525^{**}$
	(4.29)	(3.89)	(4.08)	(4.36)	(3.89)	(4.11)

Note. All regressions employ propensity-score matching with treatment-control pairs as ordinary least squares regression weights with standard errors robust to heteroskedasticity; *t*-statistics are in parentheses. $N = 2{,}893$.

$^{+}$ $p < .10$.
* $p < .05$.
** $p < .01$.

Table 12. Pseudonymous Attacks and Net Trading Losses to Sellers

	t_0	$t_0 + 1$	$t_0 + 2$	$t_0 + 3$	$t_0 + 4$
Actual dollar volume	1,044	976	977	950	933
Counterfactual dollar volume	1,050	981	982	954	935
Net losses (counterfactual − actual)	5.40	4.62	4.22	3.96	1.93
Total net losses	20.1				

Note. Values are in billions of dollars.

at $t_0 + 5$ may be greater than or less than the price on the days $[t_0, t_0 + 4]$ but is higher on average. These calculations are given in Table 12. As Table 12 shows, sellers would have received a total of $20.1 billion more during the interval $[t_0, t_0 + 4]$ if trades had been executed at the price on $t_0 + 5$.

5. CONCLUSION

This paper shows that in financial markets, pseudonymity facilitates profitable manipulation of stock prices. Pseudonymous authors publish negative rumors about public companies that lead to significant short-term trading profits—and sharp reversals of the stock price decline. When markets realize that the pseudonymous author is spreading baseless rumors, the author switches to a new pseudonym and repeats the pattern. Pseudonymity thus undermines reputational sanctions and allows manipulators to exploit investors' trust.

In the Online Appendix, I discuss the legal issues implicated by these empirical findings in greater detail. One of the challenges with addressing the sort of market manipulation documented here is that pseudonymous attacks are not easily captured by either the antimanipulation or antifraud provisions of securities laws. Aggressive options trading accompanying the publication of an article may be insufficient to establish intent to artificially depress the price of a security.

One possibility, which I have written on elsewhere with a colleague, is that "traders who anticipate a market rebound and buy ahead of it (after selling short heavily only a day or two earlier) are conceding that they did not believe their earlier purchases were truly establishing a new price equilibrium. We do not suggest that this reversal in trading supplies irrebuttable evidence of manipulation, but it could be given presumptive weight. . . . One way to justify our position is to look to *Omnicare v. La-*

borers District Council Construction Industry Pension Fund, which held that an expression of opinion can contain 'embedded' factual assertions, both that the speaker sincerely holds the view stated and did some minimal research. The sudden reversal in position by the trader in the new 'V' pattern strongly suggests it never believed in the adverse news or rumors that it cited" (Coffee and Mitts 2018).

In addition, on February 12, 2020, a group of 12 securities law professors, including me, submitted a rulemaking petition to the SEC on the topic of manipulative short selling (Robinson and Bain 2020). The petition urges the SEC to enact two rules: one that would impose a duty to update a voluntary short-position disclosure that no longer reflects current holdings or trading intention and one that would clarify that rapidly closing a short position after publishing (or commissioning) a report, without having specifically disclosed an intent to do so, can constitute fraudulent scalping in violation of Rule 10b-5.

This project raises a number of additional questions that are worthy of future study. For example, the role of intermediaries like Seeking Alpha is not fully understood. It would be interesting to understand better how Seeking Alpha detects pseudonymous authors who repeatedly switch identities and whether the administrators apply any sort of sanction to suspected cases of repeated manipulation. Moreover, it would be fascinating to see how readers react to potentially manipulative historical trading patterns. These are questions that might be fruitfully explored in future work.

REFERENCES

Aggarwal, Rajesh K., and Guojun Wu. 2006. Stock Market Manipulations. *Journal of Business* 79:1915–53.

Amihud, Yakov. 2002. Illiquidity and Stock Returns: Cross-Section and Time-Series Events. *Journal of Financial Markets* 5:31–56.

An, Byeong-Je, Andrew Ang, Turan G. Bali, and Nusret Cakici. 2014. The Joint Cross Section of Stocks and Options. *Journal of Finance* 69:2279–2337.

Antweiler, Werner, and Murray Z. Frank. 2004. Is All That Talk Just Noise? The Information Content of Internet Stock Message Boards. *Journal of Finance* 59:1259–94.

Appel, Ian, Jordan Bulka, and Vyacheslav Fos. 2018. Public Short Selling by Activist Hedge Funds. Working paper. Boston College, Carroll School of Management, Boston.

Augustin, Patrick, Menachem Brenner, Gunnar Grass, and Marti G. Subrahmanyam. 2016. How Do Informed Investors Trade in the Options Market? Working Paper No. WP 18-09. McGill University, Desautels Faculty of Management, Montreal.

Augustin, Patrick, Menachem Brenner, and Marti G. Subrahmanyam. 2016. Informed Options Trading prior to Takeover Announcements: Insider Trading? Working paper. McGill University, Desautels Faculty of Management, Montreal.

Beber, Alessandro, and Marco Pagano. 2013. Short-Selling Bans around the World: Evidence from the 2007–09 Crisis. *Journal of Finance* 68:343–81.

Benabou, Roland, and Guy Laroque. 1992. Using Privileged Information to Manipulate Markets: Insiders, Gurus, and Credibility. *Quarterly Journal of Economics* 107:921–58.

Boehmer, Ekkehart, Charles M. Jones, and Xiaoyan Zhang. 2013. Shackling Short Sellers: The 2008 Shorting Ban. *Review of Financial Studies* 26:1363–1400.

Campbell, John L., Matthew D. DeAngelis, and James Moon. 2017. The Credibility of the Crowds: Personal Stock Holdings and Investors' Response to Stock Opinions on Social Media. Working paper. University of Georgia, J. M. Tull School of Accounting, Athens.

Cao, Charles, Zhiwu Chen, and John M. Griffin. 2005. Informational Content of Option Volume prior to Takeovers. *Journal of Business* 78:1073–1109.

Carhart, Mark M. 1997. On Persistence in Mutual Fund Performance. *Journal of Finance* 52:57–82.

Chakravarty, Sugato, Huseyin Gulen, and Stewart Mayhew. 2004. Informed Trading in Stock and Option Markets. *Journal of Finance* 59:1235–57.

Coffee, John C., Jr., and Joshua Mitts. 2018. Short Selling and the New Market Manipulation. *CLS Blue Sky Blog,* March 3. https://clsbluesky.law.columbia.edu/2019/03/18/short-selling-and-the-new-market-manipulation/.

Comerton-Forde, Carole, and Tālis J. Putniņš. 2009. Are Short Sellers Manipulating the Market? Working paper. University of Sydney, School of Business, Sydney.

Cremers, Martijn, and David Weinbaum. 2010. Deviations from Put-Call Parity and Stock Return Predictability. *Journal of Financial and Quantitative Analysis* 45:335–67.

Detemple, Jérôme B. 1986. Asset Pricing in a Production Economy with Incomplete Information. *Journal of Finance* 41:383–91.

Easley, David, Maureen O'Hara, and P. S. Srinivas. 1998. Option Volume and Stock Prices: Evidence on Where Informed Traders Trade. *Journal of Finance* 53:431–65.

Farrell, Joseph, and Matthew Rabin. 1996. Cheap Talk. *Journal of Economic Perspectives* 10(3):103–18.

Fox, Merritt B., Lawrence R. Glosten, and Gabriel V. Rauterberg. 2018. Stock

Market Manipulation and Its Regulation. *Yale Journal on Regulation* 35:67–126.

Fox, Merritt B., Lawrence R. Glosten, and Paul C. Tetlock. 2009–10. Short Selling and the News: A Preliminary Report on Empirical Study. *New York Law School Law Review* 54:645–86.

Frieder, Laura, and Jonathan Zittrain. 2007. Spam Works: Evidence from Stock Touts and Corresponding Market Activity. *Hastings Communications and Entertainment Law Journal* 30:479–520.

Glosten, Lawrence R., and Paul R. Milgrom. 1985. Bid, Ask, and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders. *Journal of Financial Economics* 14:71–100.

Glosten, Lawrence R., and Tālis J. Putniņš. 2016. Welfare Costs of Informed Trade. Paper presented at the annual meeting of the American Finance Association, San Francisco, January 3–5.

Grossman, Sanford J. 1981. The Informational Role of Warranties and Private Disclosure about Product Quality. *Journal of Law and Economics* 24:461–83.

Hamilton, Alexander, James Madison, and John Jay. [1788] 2013. *The Federalist,* edited by Benjamin Fletcher Wright. Cambridge, MA: Belknap Press.

Jayaraman, Narayanan, Melissa B. Frye, and Sanjiv Sabherwal. 2001. Informed Trading around Merger Announcements: An Empirical Test Using Transaction Volume and Open Interest in Options Market. *Financial Review* 36(2):45–74.

Katz, Jeff, and Annie Hancock. 2017. Short Activism: The Rise in Anonymous Online Short Attacks. *Harvard Law School Forum on Corporate Governance,* November 27. https://corpgov.law.harvard.edu/2017/11/27/short-activism-the-rise-in-anonymous-online-short-attacks/.

Kogan, Shimon, Tobias J. Moskowitz, and Marina Niessner. 2018. Fake News: Evidence from Financial Markets. Working paper. Massachusetts Institute of Technology, Sloan School of Management, Cambridge, MA.

Ljungqvist, Alexander, and Wenlan Qian. 2016. How Constraining Are Limits to Arbitrage? *Review of Financial Studies* 29:1975–2028.

Meulbroek, Lisa K. 1992. An Empirical Analysis of Illegal Insider Trading. *Journal of Finance* 47:1661–99.

Mill, J. S. 1946. *On Liberty and Considerations on Representative Government,* edited by R. B. McCallum. Oxford: Basil Blackwell.

Mitts, Joshua, and Eric Talley. 2019. Informed Trading and Cybersecurity Breaches. *Harvard Business Law Review* 9:1–51.

Pan, Jun, and Allen M. Poteshman. 2006. The Information in Option Volume for Future Stock Prices. *Review of Financial Studies* 19:871–908.

Pastor, Lubos, and Pietro Veronesi. 2009. Learning in Financial Markets. *Annual Review of Financial Economics* 1:361–81.

Putniņš, Tālis J. 2012. Market Manipulation: A Survey. *Journal of Economic Surveys* 26:952–67.

Renault, Thomas. 2018. Pump-and-Dump or News? Stock Market Manipulation on Social Media. Working paper. Institute of Scientific Economy and Management, School of Management, Lille.

Robinson, Matt, and Benjamin Bain. 2020. SEC Urged to Seek More Disclosure When Investors Tout Short Bets. *Bloomberg,* February 13. https://www.bloomberg.com/news/articles/2020-02-13/sec-urged-to-seek-more-disclosure-when-investors-tout-short-bets.

Saffi, Pedro A. C., and Kari Sigurdsson. 2011. Price Efficiency and Short Selling. *Review of Financial Studies* 24:821–52.

SEC (Securities and Exchange Commission). Division of Market Regulation. 2015. Responses to Frequently Asked Questions concerning Regulation SHO. SEC.gov, October 15. https://www.sec.gov/divisions/marketreg/mrfaqregsho1204.htm.

SkyTides. 2016. Insulet Investors Being Kept in the Dark, CEO Alleged to Encourage Questionable Sales Techniques: Significant Downside Remains. Seeking Alpha, November 29. https://seekingalpha.com/article/4026931-insulet-investors-kept-dark-ceo-alleged-encourage-questionable-sales-techniques-significant.

Tetlock, Paul C. 2011. All the News That's Fit to Reprint: Do Investors React to Stale Information? *Review of Financial Studies* 24:1481–1512.

Van Bommel, Jos. 2003. Rumors. *Journal of Finance* 58:1499–1520.

Vila, Jean-Luc. 1989. Simple Games of Market Manipulation. *Economics Letters* 29:21–26.

Weiner, Perrie M., Robert D. Weber, and Kirby Hsu. 2017. The Growing Menace of "Short and Distort" Campaigns. *Thomson Reuters Expert Analysis,* August 31.

Wong, Yu Ting Forester, and Wuyang Zhao. 2017. Post-apocalyptic: The Real Consequences of Activist Short-Selling. Working Paper No. 17-25. University of Southern California, Marshall School of Business, Los Angeles.

Zhao, Wuyang. 2018. Activist Short-Selling. PhD diss., University of Toronto, Rottman School of Management.