

Carvalho, Carlos; Klagge, Nicholas; Moench, Emanuel

Working Paper

The persistent effects of a false news shock

Staff Report, No. 374

Provided in Cooperation with:

Federal Reserve Bank of New York

Suggested Citation: Carvalho, Carlos; Klagge, Nicholas; Moench, Emanuel (2009) : The persistent effects of a false news shock, Staff Report, No. 374, Federal Reserve Bank of New York, New York, NY

This Version is available at:

<https://hdl.handle.net/10419/60768>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Federal Reserve Bank of New York
Staff Reports

The Persistent Effects of a False News Shock

Carlos Carvalho
Nicholas Klagge
Emanuel Moench

Staff Report no. 374

May 2009

Revised March 2010

This paper presents preliminary findings and is being distributed to economists and other interested readers solely to stimulate discussion and elicit comments. The views expressed in the paper are those of the authors and are not necessarily reflective of views at the Federal Reserve Bank of New York or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

The Persistent Effects of a False News Shock

Carlos Carvalho, Nicholas Klagge, and Emanuel Moench

Federal Reserve Bank of New York Staff Reports, no. 374

May 2009; revised March 2010

JEL classification: G10, G14

Abstract

In September 2008, a six-year-old article about the 2002 bankruptcy of United Airlines' parent company resurfaced on the Internet and was mistakenly believed to be reporting a new bankruptcy filing by the company. This episode caused the company's stock price to drop by as much as 76 percent in just a few minutes, before NASDAQ halted trading. After the "news" had been identified as false, the stock price rebounded, but still ended the day 11.2 percent below the previous close. We explore this natural experiment by using a simple asset-pricing model to study the aftermath of this *false news shock*. We find that, after three trading sessions, the company's stock was still trading below the two-standard-deviation confidence band implied by the model and that it returned to within one standard deviation only during the sixth trading session. On the seventh day after the episode, the stock was trading at exactly the level predicted by the asset-pricing model. We also document that the false news shock had a persistent effect on the stock prices of other major airline companies.

Key words: false news, natural experiment, United Airlines, noise, market efficiency

Carvalho: Federal Reserve Bank of New York (e-mail: carlos.carvalho@ny.frb.org). Klagge: Federal Reserve Bank of New York (e-mail: nicholas.klagge@ny.frb.org). Moench: Federal Reserve Bank of New York (e-mail: emanuel.moench@ny.frb.org). The authors thank Tobias Adrian, Gara Afonso, Pérsio Arida, Markus Brunnermeier, Gauti Eggertson, Harrison Hong, Gur Hurlburt, Thomas Mertens, Ricardo Reis, Asani Sarkar, Dick Startz, and Paul Tetlock for comments and suggestions. Casidhe Horan provided excellent research assistance. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

1 Introduction

A central question of financial economics is whether markets are efficient. Among other things, market efficiency requires that asset prices react to news about fundamentals, as opposed to noise. However, in most circumstances relevant information and noise arise simultaneously, and cannot be easily separated. Agents have to make inference about fundamentals from possibly noisy pieces of information, and thus the noise component usually affects agent’s investment decisions. In this paper we explore a natural experiment that allows us to study a stock market’s reaction to an information release for which the noise component could be singled out very easily.

On September 8, 2008, an old article about the 2002 bankruptcy of United Airlines’ parent company (henceforth UA) resurfaced on the Internet and was mistakenly believed to be reporting a new bankruptcy filing by the company.¹ This caused the company’s stock price to drop by up to 76% in just a few minutes, before NASDAQ halted trading. After the false news had been identified as such, the stock price rebounded, but still ended the day 11.2% below the previous close. Trading volumes skyrocketed during these extreme price movements.

The episode can be thought of as comprising two pieces of information: the “news” that UA had filed for bankruptcy protection again, and the subsequent statements by UA and the media companies involved in the article’s release clarifying that it pertained to the 2002 bankruptcy filing. The clarification statements were widely circulated shortly after the large price drop, and were publicly available when trading resumed. Moreover, the false news appears to have made its way to the main sources of financial information by sheer accident. This justifies our assumption that the episode provides a natural experiment to study the effects of what we refer to as a *false news shock*: two pieces of information that cancel each other. Given this shock, we are left with the task of trying to make sense of the 11.2% drop of UA’s stock price on that day and its slow recovery on subsequent days.

In order to study the impact of the false news shock on UA’s stock price, we need a so-called “counterfactual”: the path that the stock price would likely have followed in the absence of the false news. In Section 3 we construct such a counterfactual path using a simple factor pricing model for UA’s stock return. In particular, we postulate that the excess return on UA stock depends linearly on the excess returns of three factors: the “market” (as proxied by the S&P 500), the “airline industry” (as proxied by Bloomberg’s World Airline Index), and crude oil. We estimate the asset-pricing model using data until the day before the false news impacted the market. The model captures the dynamics of UA excess returns quite well, explaining about 40% of its variation at both daily and intraday frequencies. We use our model to construct point estimates and confidence intervals for UA’s stock

¹Although the article referred to United Airlines’ parent company (UAL Corp.), throughout the paper we refer to the episode as pertaining to United Airlines. UAL Corp. is traded at NASDAQ under the ticker symbol “UAUA”.

price given the evolution of the three pricing factors on the day of the false news event, and over subsequent trading sessions.

We find that after three trading sessions UA shares were still trading below the two-standard-deviation confidence band implied by the model, and only returned to within one standard deviation of the model-implied price on the sixth trading session after the event. On the seventh day after the episode - and for quite a few days thereafter - the shares traded essentially at the level predicted by the asset-pricing model. These findings are robust to different specifications of the factor model.

In Section 5 we analyze the evolution of the stock prices of other major U.S. airlines during the episode (American Airlines, Continental Airlines, Delta Airlines and U.S. Airways). We find a very similar, although attenuated, pattern. On September 8, 2008, their share prices experienced maximum drops in the range of -25.6% to -31.8% , and ended the day between -2.5% to -5.3% relative to the previous closing price. The timing of the sharp price moves coincides with UA's. Employing the same type of factor pricing model as for UA, we construct a counterfactual path for the stock price of each of these four companies and find that the effects of the false news shock originated from the article on UA were also persistent. Finally, we document that intraday trading volumes spiked up considerably during the sharp price movements.

Throughout our analysis we maintain the assumption that the two pieces of information that comprise the false news shock cancel each other, in the sense of leaving investors with as much relevant information about the fundamentals of UA as they had before the old article reappeared on the Internet. However, it is possible that the episode revealed relevant information that agents would not otherwise have acquired. For example, agents might have learned about the depth of the market for the stock, or about the distribution of others' views on the company. We explore some of these possibilities in Section 6.

Our paper adds to the available evidence on systematic deviations from informationally frictionless and efficient markets. Huberman and Regev (2001) document that a front-page New York Times article about an old scientific discovery had a huge impact on the stock price of the company responsible for it (EntreMed), even though the scientific findings had been published in *Nature* and covered by a not-so-prominent New York Times article more than five months before. The prominent article also had spillover effects on the stock prices of other biotechnology companies. The authors conclude that "enthusiastic public attention" may induce important movements in stock prices in response to old news that may have been overlooked by a large fraction of market participants.

Like Huberman and Regev (2001), our paper provides very clean evidence on the importance of media vehicles in transmitting information to market participants and affecting how they perceive the world. While it is usually taken for granted that people receive and act on information transmitted by various media outlets, most models have no role for them - information is simply "received" (or

inferred) by agents without any reference to concrete communication channels. Apart from Huberman and Regev (2001), to the best of our knowledge there is no other paper that analyzes the effects of a news shock on stock prices by way of such a clean episode.

There is, however, a growing body of literature that aims at estimating the asset-pricing impact of news identified through application of linguistic tools to newspaper articles. Tetlock (2007) constructs a media-based measure of “sentiment” towards stock markets from a linguistic analysis of the Wall Street Journal’s “Abreast of the Market” column. He finds that high negative sentiment predicts lower returns for the Dow-Jones index over the next few days followed by a reversion, and that unusually high or low pessimism predicts high trading volume. Sinha (2009) uses a sentiment score from Thomson Reuters to measure the tone of news articles and constructs portfolios based on past sentiment. He finds that a portfolio long in positive- and short in negative-sentiment firms is positively correlated with a long-short momentum portfolio and generates positive future returns. Tetlock, Saar-Tsechansky and Mackassy (2008) use the fraction of negative words in firm-specific news articles to predict future earnings and stock returns. They find that negative language predicts negative earnings, even when they control for analysts’ forecasts and historical accounting data. They also document that stock prices respond with a one-day delay to negative language in the firm-specific news.

The papers summarized in the previous paragraph share the fact that they measure sentiment in news articles without assessing whether the information contained in the articles is actually novel. Accordingly, they are only loosely related to Huberman and Regev (2001), who document the asset-pricing impact of an article that contained old news about a company. Tetlock (2009) is somewhat closer to Huberman and Regev’s work. He constructs empirical proxies to capture the degree to which a news story about a company is stale, such as the presence of another news story in the prior week, the presence of an extreme abnormal stock return in the prior week, or high media coverage in the past month. Tetlock (2009) sorts individual stocks into calendar-time portfolios based on the firms’ recent public news. He documents that return reversals after news events are stronger when these events have a higher content of stale information according to his empirical proxies of news staleness.

While our paper is close to Huberman and Regev (2001) and Tetlock (2009) in some dimensions, there is a crucial distinction, associated with the difference between the concept of stale (or old) news, and that of false news. The old news in the EntreMed case studied by Huberman and Regev (2001), and likely in many of the stale news identified by Tetlock (2009), refer to factors which at the time were still potentially important for the future profitability of the firm. In contrast, in the United Airlines event that we study, the “news” that produced a significant and persistent deviation of the company’s stock price from its fundamental value was simply false: there was no new bankruptcy filing, and the reemergence of the six-year-old article should not matter for the company’s profitability going forward. To the best of our knowledge, the nature of the news event is unique to the episode

that we document.

The paper is organized as follows. In the next section we provide a description of the episode. Section 3 describes our pricing model and the data used in the estimation. Section 4 presents our results, and discusses robustness issues. Section 5 documents the impact of the false news shock on the stock prices and trading volumes of other major airlines. Section 6 provides a discussion of possible explanations for our findings, and the last section concludes.

2 The Episode

The episode took place on Monday, September 8, 2008, and was covered in detail in the main newspapers on subsequent days. Here we provide a summary based on articles published in the New York Times and Financial Times during the week of the event.

On Sunday, September 7, 2008, an article first published on December 10, 2002, in the Chicago Tribune about UA's bankruptcy filing allegedly made its way to the list of most viewed business-related articles on the webpage of the Sun-Sentinel, a Florida newspaper. Apparently the article contained no dates that explicitly tied it to the 2002 bankruptcy. It was found and scanned by a Google News program, which then indexed it so that it could be included in Google News' pages, and made available in Internet-search results. On Monday, September 8, the article caught the attention of an employee of a financial information company who was searching the Internet for news about recent bankruptcies. The employee sent a summary of the article with a link to the Sun-Sentinel webpage over Bloomberg's wire, and the "news" eventually appeared in a headline in Bloomberg's news service in the morning of September 8, with a reference to the Sun-Sentinel story.

At around 10:45 AM, UA's stock price was trading at \$12.05, after having traded as high as \$12.45 - from a close of \$12.30 on the previous session. Soon after 10:45 AM it started to drop sharply, and collapsed to as low as \$3.00 per share around the time that the headline citing the article appeared on Bloomberg News terminals (around 11:03 AM). By the time NASDAQ halted trading on UA stock (at 11:07 AM), a significant fraction of the drop had been reverted, and the stock was already trading at around \$9. Trading resumed at 12:30 PM, with the stock priced near \$11.25. It ended the session at \$10.92 - down 11.2% from the previous day's close. Trading volumes shot up significantly during the sharp price movements, and as a result the total volume on the day of the event was roughly three times larger than on either of the two adjacent trading days. In subsequent days, the stock price fell as low as \$9.12 (on September 11), and on September 15 it finally traded above the level of prices seen just prior to when the false news impacted the market.

At the close of the trading session prior to September 8, 2008, UA had a market capitalization of approximately \$1.6 billion. The figures for the other four airlines ranged from \$900 million for U.S.

Airways to \$2.8 billion for American Airlines. At the lowest price triggered by the false news shock the loss in market value was roughly \$1.2 billion for UA alone, and roughly \$3.5 billion for the five airlines combined. At the September 8 closing prices, the loss in market value for the five airlines exceeded \$500 million.

2.1 Two pieces of “news”

Two pieces of information arrived in connection with the reappearance of the 2002 article. The first one is the “news” that UA had filed for bankruptcy protection again. Without knowledge that the information was based on a six-year-old old article, this amounted to a substantial piece of (negative) fundamental information about UA that “justifies” the steep price drop observed right after the story was broadly circulated. The second piece of information is the clarification that the article was six years old, which “justifies” the rebound in the stock price.

A review of articles published on Bloomberg News and on the Internet on Monday, September 8, shows that the information that the Chicago Tribune article was six years old became available before trading in UA stock resumed. The second piece of news involved clarification statements from United Airlines, the Chicago Tribune, and Bloomberg. Even if one attaches some uncertainty as to whether investors had access to all the statements during that trading session, and took them at face value, the episode was widely covered by the main media on the subsequent day.

Thus, in our quantitative analysis we maintain the assumption that the two pieces of news cancel each other, in the sense of leaving investors with as much relevant information about the fundamentals of UA as they had before the old article reappeared, all else equal. We revisit this assumption in Section 6, where we discuss possible explanations for our findings.

3 Pricing model and data

We postulate a simple three-factor model for the excess return on UA’s stock:

$$r_{UA,t} - r_t = c + \beta_M (r_{M,t} - r_t) + \beta_A (r_{A,t} - r_t) + \beta_O (r_{O,t} - r_t) + e_t, \quad (1)$$

where $r_{UA,t}$, r_t , $r_{M,t}$, $r_{A,t}$, and $r_{O,t}$ denote the (logarithmic) returns between periods $t - 1$ and t on, respectively, UA’s stock, a one-period risk-free nominal bond, the market portfolio, the airline-industry portfolio, and crude oil. The loadings on the factors are given by the coefficients β_M , β_A and β_O , and c is a constant. We think it is natural to expect estimates of β_M and β_A to be positive, and estimates of β_O to be negative. Finally, e_t is an error term that captures the idiosyncratic component in UA’s stock return.

We estimate (1) by Ordinary Least Squares (OLS) using the effective federal funds rate as the risk-

free rate, the S&P 500 as the proxy for the market portfolio, Bloomberg’s World Airline Index as the measure of the airline-industry portfolio,² and the price of crude oil futures as reported in Bloomberg under the code “CL1 Comdty”. Our baseline specification uses 15-minute intraday data, but we also estimate the model with daily data. All daily data are from Bloomberg. Intraday data for UA’s stock price are from the Trade and Quote (TAQ) database. Intraday data for the S&P 500, Bloomberg’s World Airline Index and oil futures are from Bloomberg.³ In our estimations with intraday data, we use a sample from March 3, 2008 through September 5, 2008 (the day before the event). Our sample with daily data starts on August 1, 2007. Finally, we perform our counterfactual analysis on data for the period September 8, 2008 - September 17, 2008.

We use the estimated pricing model to ask the following counterfactual question: “what would have happened to the stock price of UA if the false news shock had not occurred?” Let t_0 denote the day of the event (or, in the case of intraday data, the first trading period on that day). Knowledge of the realizations of $r_{M,t}$, $r_{A,t}$, $r_{O,t}$, and r_t for $t \geq t_0$ allows us to construct an estimate of what the return on UA’s stock would have been in the absence of any variation in the firm-specific component of the stock return in (1):

$$\hat{r}_{UA,t} = r_t + \hat{c} + \hat{\beta}_M (r_{M,t} - r_t) + \hat{\beta}_A (r_{A,t} - r_t) + \hat{\beta}_O (r_{O,t} - r_t), \quad t \geq t_0, \quad (2)$$

where “hatted” variables denote estimates. Then we can add our estimates of $\hat{r}_{UA,t}$ for $t \geq t_0$ to UA’s (log) stock price at the close of the last trading session prior to the event, denoted s_{UA,t_0-1} , to obtain point estimates for what UA’s (log) stock price would have been in the absence of the false news shock (and of any other non-zero realization for e_t):

$$\hat{s}_{UA,t} = s_{UA,t_0-1} + \sum_{j=t_0}^t \hat{r}_{UA,j}, \quad t \geq t_0. \quad (3)$$

In addition, we can construct confidence intervals for the estimates $\hat{s}_{UA,t}$ using the standard error from the OLS regression, denoted $\hat{\sigma}_e$. Specifically, a k -standard-error confidence band for $\hat{s}_{UA,t}$ can be obtained as:⁴

$$\hat{s}_{UA,t} \pm \sqrt{t - (t_0 - 1)} k \hat{\sigma}_e. \quad (4)$$

²Although UA was included in the index during our sample period, its weight was small (of the order of 2%). Thus, for simplicity we do not incorporate that information in our estimation.

³We assume that r_t is constant during the day. For the intraday data on oil futures we use the ask price, stored under the code “CL1 Comdty BarTp=A”. Bloomberg only makes intraday data available for relatively short periods of time (up to approximately six months). We first downloaded intraday data shortly after the episode, in September 2008. Subsequently, around March 2009 we downloaded additional intraday data and found small discrepancies in the quotes for oil futures under the code “CL1 Comdty”. We spliced the two “vintages” of these data using the 9:30AM quotes for September 8, 2008.

⁴This construction does not adjust for the estimation uncertainty regarding the factor loadings. We also analyzed the results with this adjustment, and found the differences to be negligible. The reason is that, as we report in the next section, the regression coefficients are precisely estimated. To ease the exposition of how we construct the confidence intervals, we present the case without that adjustment.

4 Results

Table 1 presents the results for the estimation based on intraday data (15-minute intervals) for the period March 3, 2008 - September 5, 2008, and on daily data for the period August 1, 2007 - September 5, 2008. The factor loadings have the expected signs and are highly statistically significant. Moreover, the results are consistent across the two sampling frequencies.

Table 1: Three-factor model - United Airlines

Parameters and test statistics	Frequency	
	Intraday ^{a)}	Daily
$\hat{\beta}_M$	1.90*** (0.17)	1.53*** (0.31)
$\hat{\beta}_A$	1.36*** (0.21)	1.63*** (0.28)
$\hat{\beta}_O$	-1.24*** (0.10)	-1.33*** (0.22)
\hat{c}	-6 x 10 ⁻⁵ (0.0002)	7 x 10 ⁻⁴ (0.004)
R^2	0.40	0.42
# Obs	3394	278
$\hat{\sigma}_e$	0.013	0.06
$F : \beta_M = \beta_A = \beta_O = 0$	762***	66.15***
<i>Durbin-Watson</i>	1.93	2.02
<i>White heteroskedasticity F-test</i> (<i>p-value</i>)	13.53*** (0.00)	1.04 (0.41)
<i>Breusch-Godfrey F-test</i> <i>for serial correlation</i> (<i>p-value</i>)	1.63 ^{b)} (0.06)	0.87 ^{c)} (0.46)


Notes: Newey-West robust standard errors in parenthesis (8 lags for intraday, 5 lags for daily), unless indicated otherwise. a) 15-minute intervals. b) 15 lags. c) 3 lags. ***Denotes statistical significance at 1% level.

Figure 1 presents the results of our counterfactual analysis based on intraday data, as described by equations (2)-(4). The series are transformed back into levels, so that “Actual UA” corresponds to $\exp(s_{UA,t})$, “Counterfactual UA” corresponds to $\exp(\hat{s}_{UA,t})$, and so on. Low- k and Hi- k correspond to, respectively, lower and upper k -standard-error confidence bands. The results with daily data are very similar, and for completeness they are presented in Figure 2.

We focus on the intraday data (Figure 1), since they provide a more nuanced picture of the episode. After the sharp price movements that occurred before trading was halted in the morning of September 8, UA’s stock price remained below the lower two-standard-error confidence band implied by the model until the end of the September 10 trading session. During the September 11 and 12 trading sessions,

the stock price narrowed the gap relative to the counterfactual level implied by the model, and closed just below the lower one-standard-error confidence band. In the first hours of trading on Monday, September 15, UA’s stock price surpassed the lower one-standard-error band, and on the subsequent day moved almost exactly to the level implied by our counterfactual analysis. It is worth noting that UA’s stock price already opened the session on Monday September 8 trading just below the lower two-standard-error confidence band implied by the model. Given that the old Chicago Tribune article was already available in Internet-search results since the previous day, this might reflect the price impact of trading by investors that became aware of the article before it appeared in main financial news services later on Monday morning.

Figure 1: Counterfactual analysis, United Airlines (UA) - intraday data




4.1 Robustness

We performed the same analysis using counterfactuals and error bands based on factor models estimated with intraday data at other frequencies, and the findings are essentially unchanged. We also investigated a factor model for realized volatility of UA based on intraday data, and found that the shock had no effects on volatility beyond the day in which the false news first impacted UA’s stock price. Finally, we estimated an expanded version of the factor model that includes the Fama-French “Small Minus Big”, “High Minus Low” and “Momentum” factors, using daily data. Under this alternative counterfactual, UA’s stock price still traded below the two-standard-error confidence band implied by the model for the same three-day period after the event (September 8-10). It then traded below the one-standard-error band for one day, and on the sixth session after the event it traded almost at the exact level implied by the augmented factor model. For brevity we do not present these

results here, and report the estimates in the Appendix.

Figure 2: Counterfactual analysis, United Airlines (UA) - daily data



5 “Contagion”

In this section we study the impact of the false news shock on the stock prices of other major U.S. airlines. We estimate a factor model like the one in equation (1) for each of the following companies: American Airlines (AMR), Continental Airlines (CAL), Delta Airlines (DAL) and U.S. Airways (LCC), and construct a counterfactual stock price for each company using equations analogous to (2)-(4), where UA is replaced with the respective airline. Intraday data for the stock prices of those four companies are from the TAQ database, in 15-minute intervals. The results of OLS regressions are reported in Table 2. The bottom line is that they are similar to our findings for UA.

The counterfactual analyses are presented in Figures 3-6. The results for American Airlines and Continental Airlines are extremely similar to UA’s - with the clear exception that the price drop around 11:00 AM on September 8, 2008, is significantly smaller. The pattern for Delta Airlines and U.S. Airways is slightly different, in that on the day of the false news event their stock prices do return to a level that is somewhat closer to the counterfactual level implied by the model. However, it is still the case that it takes a few days for their stock prices to return to within one standard deviation of the model’s predicted price, and one week until they trade at the level implied by the factor models. Finally, Figure 7 shows intraday trading volumes in 15-minute intervals for all five airline companies. It is clear that volumes skyrocketed during the sharp price movements. Thus, overall it appears that the effects of the false news shock were “contagious”.

Table 2: Three-factor model for other airlines - intraday data

Parameters and test statistics	Company			
	AMR	CAL	DAL	LCC
$\hat{\beta}_M$	2.00*** (0.16)	1.94*** (0.15)	1.45*** (0.18)	2.17*** (0.17)
$\hat{\beta}_A$	1.15*** (0.19)	1.14*** (0.19)	1.33*** (0.20)	1.38*** (0.27)
$\hat{\beta}_O$	-1.06*** (0.07)	-1.10*** (0.08)	-0.93*** (0.08)	-1.32*** (0.12)
\hat{c}	1×10^{-4} (0.0002)	9×10^{-5} (0.0002)	7×10^{-5} (0.0002)	8×10^{-5} (0.0002)
R^2	0.49	0.49	0.40	0.42
# Obs	3394	3394	3394	3394
$\hat{\sigma}_e$	0.01	0.01	0.011	0.014
$F : \beta_M = \beta_A = \beta_O = 0$	1096***	1117***	754***	802***
<i>Durbin-Watson</i>	2.04	2.07	2.13	2.16
<i>White heteroskedasticity F-test</i> (p-value)	52.35*** (0.00)	28.17*** (0.00)	51.92*** (0.00)	47.80*** (0.00)
<i>Breusch-Godfrey F-test</i> <i>for serial correlation</i> ^{a)} (p-value)	0.62 (0.86)	1.80** (0.029)	2.11*** (0.007)	2.81*** (0.000)

Notes: Newey-West robust standard errors in parenthesis (8 lags), unless indicated otherwise. a) 15 lags.

***(**) Denotes statistical significance at a 1% (5%) level.

6 Discussion

One possible explanation for the failure of UA's share price to return to its pre-episode equilibrium value following the false news shock involves multiple equilibria for different levels of liquidity. Amihud and Mendelson (1986) present a model and supporting empirical evidence showing that, all else equal, investors demand a higher return on illiquid securities. This suggests that the price of a stock suffering an exogenous adverse liquidity shock should fall, in order to compensate buyers through higher expected future returns. Building on this work, Dow (2004) presents a model in which a security may exhibit multiple equilibria depending on its level of liquidity - in a simple case, a "normal" equilibrium and a "low-liquidity" equilibrium, which is characterized by a higher bid-ask spread and a lower price.

Thus, we consider the possibility that the false news event and subsequent halt of trading in UA stock acted as an exogenous negative shock to liquidity, moving the stock temporarily from a normal equilibrium to a low-liquidity equilibrium. In this interpretation, widened bid-ask spreads during the period of rapid price movement may have deterred some liquidity traders from trading in the stock. Even after the news was shown to be false, the lack of participating liquidity traders might have

Figure 3: Counterfactual analysis, American Airlines (AMR) - intraday data





Figure 4: Counterfactual analysis, Continental Airlines (CAL) - intraday data



resulted in a persistently higher bid-ask spread, yielding a new equilibrium with lower liquidity and a lower stock price. If this were the case, the slow return of UA's stock price to its model-predicted value could be interpreted as a drift back towards the normal-liquidity equilibrium, with the corresponding excess return compensating those traders who were willing to hold the stock during the low-liquidity post-event period.

We do not find evidence to support the liquidity story. To investigate the hypothesis of abnormally low liquidity following the event, we examined UA's daily bid-ask spreads from the database of the Center for Research in Security Prices (CRSP) for our full sample. We expressed them both in absolute (dollar-value) and proportional (percentage of mid-price) terms. On the day of the event (September 8), the closing absolute spread was \$0.06. This value, while high, is less than one standard deviation above the sample mean. More notably, on the following day, with the stock still trading at a low price relative to the level implied by the factor model, the closing absolute spread was \$0.01, well below the sample mean. Proportional bid-ask spreads show a similar pattern. This finding is inconsistent with the story of a persistently low stock price due to persistently low liquidity.

Other candidate explanations for UA's share price anomaly rely on the idea that the episode might have revealed additional relevant information - and thus that the two pieces of information discussed in Subsection 2.1 did not cancel each other. To spell out one such explanation, it is possible that trading activity triggered by the false news shock revealed relevant information that agents would not have otherwise acquired. For example, the presence of short-selling restrictions or limitations might have prevented the views of investors who were most negative about the prospects for UA from being fully reflected in the company's stock price prior to the episode. In that context, the price collapse due to the false news shock might have revealed that a non-negligible fraction of market participants believed that a sudden bankruptcy of UA was possible. This might have led other investors to reassess their views of the company's prospects, thus impacting its stock price even after the news were identified as false. One weakness of this explanation - shared by others that rely on the argument that additional relevant information was revealed during the episode - is that they need yet another story to rationalize the return of UA's stock price to the level predicted by the factor model. In other words, there needs to be a reason for the extra information that was revealed by the episode to gradually lose its importance over time after the day of the episode.

The last two explanations that we entertain involve uncertainty (or ambiguity) aversion. In the first one, the false news shock might have prompted some investors to sell all of their holdings of UA stock. This might have made them reluctant to buy back the shares after the false news had been clarified, due to the so-called portfolio inertia around a zero position that may be induced by uncertainty (or ambiguity) aversion (see Dow and Werlang 1992, and Mukerji and Tallon 2001). In that case the amount of UA shares outstanding would have ended up being held by a smaller number of investors. In the presence of limits to arbitrage - for instance, due to the existence of liquidity traders and limited risk-bearing capacity - restoring equilibrium would have required a lower price in order to compensate these investors for the risk of holding larger positions in UA stock. The second explanation builds on Epstein and Schneider (2008). The authors assume that, when faced with news of uncertain quality, ambiguity-averse investors act as if they take a worst-case assessment of quality. This leads them to react more strongly to bad news than to good news. The authors also show that,

for given fundamentals, shocks to information quality may have persistent negative effects on prices if ambiguity-averse agents are trying to learn about a possible change in fundamentals on the basis of news of uncertain quality. If one assumes that the false news shock amounted to a deterioration in the quality of information regarding UA, their results seem to provide a coherent qualitative explanation of the episode in its entirety.

7 Conclusion

We explore a natural experiment to study the impact of a false news shock on the stock price of United Airlines. We find that the shock had a persistent effect on the level of UA's stock price: it took six trading sessions for the stock to return to within one-standard error of the model-implied counterfactual level. On the seventh trading session after the episode, UA's stock price was trading at the exact counterfactual level implied by our factor model. We find a similar pattern for the stock prices of other major U.S. airlines.

It is difficult to find other episodes that could be similarly characterized as a false news shock. There are a number of at first seemingly related cases that were subsequently shown to involve a fraud or hoax. In such cases, false news are deliberately produced to impact the stock price. This changes the nature of the trading environment, since some market participants trade with knowledge of the false news. It is reasonable to assume that the hoaxer takes advantage of the induced price movements by trading in the stock or its derivatives. One should thus expect a more complete reversal of the price movements produced by the false news, as the hoaxer trades to his or her advantage.

Some prominent examples of false news due to fraud involved Pairgain Technologies (on April 7, 1999; the company later merged with ADC Telecommunications in 2000), and Emulex Corporation (on August 25, 2000). Curiously, despite reports that the frauds became apparent before the end of the respective trading sessions, in both cases the stock price still ended the day moving "in the direction" that the false information would have justified.⁵ Despite their different nature, we see these episodes as suggestive that our findings would generalize to other false news shocks.

Finally, one may reasonably argue that one week is not a long enough spell for the misvaluation of stocks to have relevant economic effects - beyond gains and losses by traders and investors. However, this is how long it took for the effects of a pure false news shock to dissipate. In most circumstances, relevant information ("signal") and noise arise simultaneously, and cannot be so easily separated.

⁵For the interested reader, the Pairgain episode involved fake news of a takeover, disseminated through an (illegitimate) Internet page that was set up to look like a Bloomberg News one. The stock price went up by more than 30% before the hoax was detected during the day. Pairgain's stock price still closed 10.3% higher. The Emulex case involved a false press release that the company was restating earnings results. The report made its way to Bloomberg News through an information firm that distributes press releases - the firm allegedly fell victim to a sophisticated fraud, which made the press release appear legitimate. Emulex's stock price fell by more than 58% before trading was halted. It ended the day 6.5% below the previous close.

Figure 5: Counterfactual analysis, Delta Airlines (DAL) - intraday data




Figure 6: Counterfactual analysis, U.S. Airways (LCC) - intraday data





Figure 7: Intraday trading volumes on September 8, 2008 - 15-minute intervals



References

- [1] Amihud, Y., and H. Mendelson (1986), "Asset Pricing and the Bid-Ask Spread," *Journal of Financial Economics* 17: 223-249.
- [2] Dow, J. (2004), "Is Liquidity Self-Fulfilling?," *Journal of Business* 77: 895-908.
- [3] Dow, J. and S. Werlang (1992), "Uncertainty Aversion, Risk Aversion, and the Optimal Choice of Portfolio," *Econometrica* 60: 197-204.
- [4] Epstein, L. and M. Schneider (2008), "Ambiguity, Information Quality, and Asset Pricing," *Journal of Finance* 63: 197-228.
- [5] Mukerji S. and J. Tallon (2001), "Ambiguity Aversion and Incompleteness of Financial Markets," *Review of Economic Studies* 68: 883-904.
- [6] Huberman G. and T. Regev (2001), "Contagious Speculation and a Cure for Cancer: A Non-Event that Made Stock Prices Soar," *Journal of Finance* 56: 387-396.
- [7] Sinha, N. (2009), "News Articles and Momentum," mimeo, University of Maryland.

- [8] Tetlock, P. (2007), “Giving Content to Investor Sentiment: The Role of Media in the Stock Market,” *Journal of Finance* 62: 1139-1168.
- [9] _____ (2009), “All the News That’s Fit to Reprint: Do Investors React to Stale Information?,” mimeo, Columbia University.
- [10] Tetlock, P., M. Saar-Tsechansky, and S. Macskassy (2008), “More Than Words: Quantifying Language to Measure Firms’ Fundamentals,” *Journal of Finance* 63: 1437-1467.

8 Appendix

The results for the three-factor model for United Airlines estimated on intraday data at 30- and 60-minute intervals are presented in Table 3. They are very close to our benchmark specification using intraday data at the 15-minute interval and do not affect our counterfactual analysis. For the sake of brevity, we therefore do not report the corresponding charts.

Table 3: Three-factor model - United Airlines, other frequencies

Parameters and test statistics	Frequency	
	30 min	60 min
$\hat{\beta}_M$	1.89*** (0.22)	2.15*** (0.24)
$\hat{\beta}_A$	1.39*** (0.24)	1.35*** (0.23)
$\hat{\beta}_O$	-1.36*** (0.11)	-1.45*** (0.13)
\hat{c}	-1×10^{-4} (0.0004)	2×10^{-4} (0.0009)
R^2	0.42	0.45
$\# \text{ Obs}$	1697	913
$\hat{\sigma}_e$	0.019	0.026
$F : \beta_M = \beta_A = \beta_O = 0$	405.9***	247.09***
<i>Durbin-Watson</i>	2.04	2.04

Notes: Newey-West robust standard errors in parenthesis (8 lags).

***Denotes significance at a less than 1% level.

In order to study the impact of the false news shock on volatility we estimate the following factor model for log realized volatility of United Airlines' intraday returns:

$$\sigma_{UA,t} = a + \gamma_M \sigma_{M,t} + \gamma_A \sigma_{A,t} + \gamma_{Lag} \sigma_{UA,t-1} + \varepsilon_t,$$

where $\sigma_{UA,t}$, $\sigma_{M,t}$, and $\sigma_{A,t}$ denote the log of the realized (annualized) volatility of 15-minutes returns for United Airlines, the S&P 500 index, and Bloomberg's World Airline Index.⁶ The results of this regression are given in Table 4. The results for the counterfactual analysis of log realized volatility for the period September 8 - September 17, 2008, based on the equation

$$\hat{\sigma}_{UA,t} = \hat{a} + \hat{\gamma}_M \sigma_{M,t} + \hat{\gamma}_A \sigma_{A,t} + \hat{\gamma}_{Lag} \hat{\sigma}_{UA,t-1},$$

are illustrated in Figure 8.


⁶We also considered versions that include the realized volatility of oil, but the estimated coefficients were always very small and statistically insignificant at the usual significance levels.

Table 4: Factor model for log realized volatility - United Airlines

Parameters and test statistics	Estimates
$\hat{\gamma}_M$	-0.13 (0.14)
$\hat{\gamma}_A$	0.86*** (0.11)
$\hat{\gamma}_{Lag}$	0.23** (0.10)
\hat{a}	1.70*** (0.28)
R^2	0.65
# Obs	130
$\hat{\sigma}_\varepsilon$	0.30
$F : \gamma_M = \gamma_A = \gamma_{Lag} = 0$	78.6***
<i>Durbin-Watson</i>	2.00

Notes: Newey-West robust standard errors in parenthesis (8 lags). ** denotes significance at a 5% level, and *** denotes significance at a less than 1% level.

Figure 8: Counterfactual analysis for log realized volatility - United Airlines



Finally, we estimated an expanded version of the factor model that includes the Fama-French “Small Minus Big” (SMB), “High Minus Low” (HML) and “Momentum” (UMD) factors, using daily data from August 1, 2007 through September 5, 2008. The results are reported in Table 5, and the counterfactual analysis is presented in Figure 9.

Table 5: Extended factor model - United Airlines

Parameters and test statistics	Estimates
$\hat{\beta}_M$	1.06*** (0.22)
$\hat{\beta}_A$	1.36*** (0.24)
$\hat{\beta}_O$	-0.86*** (0.17)
$\hat{\beta}_{SMB}$	2.28*** (0.88)
$\hat{\beta}_{HML}$	1.61 (1.33)
$\hat{\beta}_{UMD}$	-1.05*** (0.25)
\hat{c}	-5×10^{-4} (0.003)
R^2	0.51
# Obs	278
$\hat{\sigma}_e$	0.057
$F : \beta_M = \dots = \beta_{UMD} = 0$	46.5***
<i>Durbin-Watson</i>	2.12

Notes: Newey-West robust standard errors in parenthesis (5 lags). ***Denotes significance at a less than 1% level. $\hat{\beta}_{SMB}$, $\hat{\beta}_{HML}$ and $\hat{\beta}_{UMD}$ denote regression coefficients on Fama-French SMB, HML, and UMD factors.

Figure 9: Extended factor model - United Airlines

