Does Social Media Accelerate Product Recalls? Evidence from the Pharmaceutical Industry

(Authors' names blinded for peer review)

Social media has become a vital platform for voicing product-related experiences that may not only reveal product defects but also impose pressure on firms to act more promptly than before. This study scrutinizes the rarely-studied relationship between these voices and the speed of product recalls in the context of the pharmaceutical industry where social media pharmacovigilance is becoming increasingly important for the detection of drug safety signals. Using Federal Drug Administration (FDA) drug enforcement reports and social media data crawled from online forums and Twitter, we investigate whether social media can *accelerate* the product recall process in the context of drug recalls. Results based on discrete-time survival analyses suggest that more adverse drug reaction (ADR) discussions on social media lead to a higher hazard rate of the drug being recalled and, thus, a shorter time to recall. To better understand the underlying mechanism, we propose the *information effect*, which captures how extracting information from social media helps detect more signals and mine signals faster to accelerate product recalls, and the *publicity effect*, which captures how firms and government agencies are pressured by public concerns to initiate speedy recalls. Estimation results from two mechanism tests support the existence of these conceptualized channels underlying the acceleration hypothesis of social media. This study offers new insights for firms and policymakers concerning the power of social media and its influence on product recalls.

Key words: product recall; drug recall; social media; pharmacovigilance; discrete-time survival analysis

1. Introduction

On March 10, 2019, Ethiopian Airlines Flight 302 crashed approximately six minutes after takeoff, killing all 149 passengers and eight crew members. Ethiopian Airlines grounded its fleet of Boeing 737

MAX aircraft the same day, followed by the grounding decision of the Civil Aviation Administration of China the next day. The issue was intensely discussed on social media, as is evidenced by over 870,000, mostly negative, tweets in just a few days (Stoll 2019). On March 13, 2019, the United States (US) Federal Aviation Administration (FAA) grounded the MAX aircraft. While the acting FAA chief, Daniel Elwell, said the agency was not bowing to public concerns, social media users, such as Lee Conner, an attorney in Washington, were convinced that the pressure mattered (Stoll 2019). Agnes Grossman, an attorney in Chicago, also thought otherwise, commenting, "*even a single voice is part of a bigger chorus, and that can result in positive change.*"

Clearly, there was a disagreement on whether social media played a role in the grounding of the Boeing 737 MAX. While we are unlikely to resolve the debate in that particular case, the broader question of whether social media can *accelerate* the product recall process is worth careful empirical investigation. Moreover, the answer to this important research question can shed light on the role of social media in accelerating positive changes in business organizations and society at large.

This paper addresses the research question in the context of the pharmaceutical industry. Although the Food and Drug Administration (FDA) requires clinical trials by pharmaceutical companies to ensure the safety and effectiveness of their drugs, adverse drug reactions (ADRs) which are harmful or unpleasant reactions to medicines, are nevertheless common after FDA approval. For example, Downing et al. (2017) showed that a third of the drugs approved by the FDA between 2001 and 2010 were involved in safety events *after* they became available on the market. Leaving unsafe drugs on the market endangers patient health and leads to significant financial and reputation costs for pharmaceutical companies (Lazarou et al. 1998, Ahmad 2003, Xu and Wang 2014). It is thus crucial to detect ADRs and take appropriate actions in a timely manner. Therefore, understanding the role of social media in accelerating the product recall process in this industry is of particular value, with important implications for both regulators and firms. Empirically, the relatively frequent occurrence of drug recalls allows us to go beyond anecdotal evidence, as in the Boeing 737 MAX example, by statistically analyzing the connection between ADR discussions on social media and the time it takes for pharmaceutical companies to recall a drug.

To understand why social media may *accelerate* product recalls, we conceptualize two distinct channels for this process. The first channel, which we refer to as the *information effect*, is particularly relevant in the context of a food or drug recall. Unlike the FDA Adverse Event Reporting System (FAERS), which is the formal reporting system, social media provides consumers extremely convenient channels through which they can discuss potential ADRs and often receive prompt feedback from peers or firms. For example, instead of complaining through the FDA, many consumers of Soylent's food bars took to social media sites to voice concerns after experiencing gastrointestinal distress, including nausea, vomiting, and diarrhea.[1] The prevalent and important nature of social media sites as ADR discussion platforms is clearly evidenced by the fact that many companies and government agencies have started monitoring ADR discussions on social media in the hope of discovering potential product defects faster than through traditional methods. One success story of this approach is Soylent's voluntary recall of its food bars on October 12, 2016, only two months after the product launch, without the FDA ever being involved.[2]

The second channel, which we refer to as the *publicity effect*, is based on the belief that the reach and intensity of social media have drastically increased consumer expectations for immediate accountability and fast responses by firms. Unlike the pre-social media era, when traditional communication methods (e.g., phone calls, letters, emails) were mostly solitary and brands were largely in control, consumer complaints on social media are public by nature, enabling a positive feedback loop that can quickly turn a small itch into a big rash. No longer in control, brands have little room to decide if or when to respond. Almost like direct democracy, social media complaints now force brands to respond and to respond fast.

Despite our hypothesized mechanisms through which social media accelerates the product recall process (henceforth referred to as the **acceleration hypothesis**), neither of the mechanisms has to be true in reality. Regarding the information effect, critics of social media pharmacovigilance question

---

[1] See https://www.engadget.com/2016/10/13/soylent-recalls-its-food-bars-after-making-some-customers-sick/

[2] See https://www.fooddive.com/news/is-soylents-social-media-fueled-recall-a-vision-of-the-future/428270/

the *validity* of ADR discussions on social media. Consumers may not accurately report ADRs on social media, resulting in useless or even misleading data for companies and government agencies to review. The inclusion of inaccurate self-reports from social media can flood the existing ADR monitoring system with noise, thereby diluting the true safety signals. Even if social media is a credible source of ADR information with high validity, such information may add little incremental value to companies if competing information sources are no less prompt and no less reliable. Extracting information from social media can overburden pharmaceutical companies and the FDA with non-life-threatening or redundant ADRs, causing new ADRs that actually bring harm to patients to be overlooked. Therefore, social media may not necessarily accelerate the recall process and can even cause delays if social media pharmacovigilance distracts pharmaceutical companies and consumes too much of their resources. Regarding the publicity effect, many doubt that the process involved in making high-stakes decisions, such as product recall decisions, can be influenced by social media at all. As Daniel Elwell publicly stated, "*we don't make decisions about grounding aircraft or regulating or even shutdowns of aircraft without actionable data*"(Stoll 2019). Indeed, initiating a hasty product recall under social media pressure or based on faulty data may do more harm than good to the brand and to customers.

The goal of the current paper is to empirically evaluate, in the context of drug recalls, whether social media can accelerate the product recall process. We collect drug-related social media data from *MedHelp, DailyStrength, SteadyHealth, MedsChat, Drugs.com, DrugBuyersGuide,* and *Twitter*, and obtain drug recall data from the FDA. To model the linkage between social media mentions and the speed of the drug recall process, we conduct discrete-time survival analysis. We control for a wide range of potential confounding factors, such as ADR reports on FAERS and traditional newspaper mentions of drugs, as well as drug characteristics. Various estimation results consistently show that more ADR discussions on social media are associated with a higher hazard rate. In other words, we find empirical evidence suggesting that social media accelerates the product recall process, at least in the context of drug recalls. Moreover, the results from two mechanism tests support the existence of the information effect and the publicity effect, which explain why social media accelerates product

recalls. We also find that the acceleration effect is stronger for drugs that are only distributed within multiple states than those that are more widely distributed.

The rest of the paper is organized as follows. In Section 2, we provide the research background by explaining drug recalls, ADRs, and pharmacovigilance, followed by a review of relevant literature. We develop the acceleration hypothesis in Section 3 by elaborating on the information effect and the publicity effect. We describe our data, variables, and the econometric model in Section 4 and report various estimation results in Section 5. We conclude the paper in Section 6 by discussing its implications and limitations.

2. Research Background

In this section, we briefly describe the institutional background regarding drug recalls, ADRs, and pharmacovigilance, and then review two streams of relevant literature: product-harm crisis and social media. While both streams of literature are vast, there is little overlap between the two, with a few exceptions. In particular, no study has examined whether social media can accelerate the product recall process.

2.1. Institutional Background

A drug recall removes a prescription or over-the-counter (OTC) drug from the market.[3] The financial impact of recalling a drug can be substantial. For example, drug manufacturer Merck lost nearly $2.5 billion by recalling Vioxx, which increased the risk of heart attacks (Appleby and Krantz 2004). However, a timely drug recall is the most effective way to protect patients from a defective or potentially dangerous product. In addition, promptly taking responsibility is in the best interest of the pharmaceutical company because it minimizes the long-term damage to its brand image. In the US, nearly all recalls are voluntarily initiated by firms.[4] As a government agency, the FDA is responsible for monitoring drug safety and regulating pharmaceutical companies. It also reserves the right to issue warnings to firms and mandates that firms recall their products when necessary.

---

[3] See https://www.fda.gov/drugs/drug-recalls/fdas-role-drug-recalls

[4] See https://www.fda.gov/consumers/consumer-updates/fda-101-product-recalls

ADRs, defined as any harmful or unpleasant reactions (Edwards and Aronson 2000), are among the main factors that cause a drug recall. ADRs result in massive costs for consumers, companies, and government agencies, including up to 5% of hospital admissions, 28% of emergency visits, 5% of hospital deaths, and associated financial costs of about $75 billion per year (Lazarou et al. 1998, Ahmad 2003, Xu and Wang 2014). Early detection of ADRs and prompt actions undoubtedly improve patients' health status and reduce financial costs. Pharmaceutical companies and government agencies thus have strong incentives to detect ADRs after drugs are on the market, part of the practice known in the industry as *pharmacovigilance.*

The World Health Organization (2002) defines pharmacovigilance as activities related to the detection, assessment, understanding, and prevention of ADRs attributable to prescription drugs. The FDA and pharmaceutical firms start practicing pharmacovigilance during a drug's pre-approval stage and continue throughout its life on the market. In the pre-approval stage, firms and the FDA collect information, including ADRs, during Phases I - III of clinical trials. After pharmaceutical companies release drugs into the market, they monitor the safety of the drugs mainly through Phase IV clinical trials and the so-called spontaneous adverse event reports. Clinical trials undoubtedly provide complex safety analyses, such as drug-drug interaction analyses, but are limited with respect to the number and characteristics of the patients exposed, the duration, and the type of data collected (Harpaz et al. 2012). As a result, government agencies have expanded their pharmacovigilance efforts to spontaneous reporting systems, such as FAERS. Due to the voluntary nature and complexity of reporting to FAERS, the system has many limitations, including possible underreporting, delayed data, incomplete data, and duplicative reporting. In addition, patients' perspectives could be selectively filtered through healthcare professionals and regulatory agencies (Forster et al. 2012, Harpaz et al. 2012). To overcome these limitations, pharmaceutical companies and academic researchers have explored additional sources for monitoring ADRs. In particular, social media has emerged as an important alternative source for detecting potential ADR signals.

Researchers from different disciplines have proposed models, which are commonly referred to as *social media pharmacovigilance*, to monitor ADRs using social media data. Pharmaceutical companies are also investing heavily in social media pharmacovigilance. For example, companies provide

guidance to facilitate effective safety reporting via social media.[5] According to a survey by Aris-Global,[6] 78% of companies are monitoring or developing artificial intelligence (AI) technology to monitor social media for drug safety signals.

2.2. Product-Harm Crisis Literature

Our work draws references from the product-harm crisis literature (summarized in Table 1), which can be broadly categorized into two streams. The first stream of literature focuses on firms' performance outcomes following product recalls. It has been well established in previous studies that product recalls are negatively associated with stock market returns and cause significant shareholder wealth loss (Jarrell and Peltzman 1985, Marcus 1989). Recent studies have advanced our understanding of this negative impact by investigating how firms' strategies can moderate the relationship between product recalls and financial outcomes. Chen et al. (2009) showed that, regardless of the firm and product characteristics, proactive strategies have a more negative effect on firm value than more passive strategies. Lee et al. (2015) examined how corporate social media affects the capital market consequences of firms' disclosure in the context of consumer product recalls. They found that, although corporate social media attenuates the negative price reaction to product recall announcements, the attenuation benefits, while still significant, lessened with the arrival of Facebook and Twitter. Using data on automobile recalls and advertising expenditures from 2005 to 2012, Gao et al. (2015) showed that firms' adjustment on advertising expenditure before initiating a recall can either mitigate or amplify the negative effect of the recall on stock market value, depending on the direction of advertising adjustment and the recall characteristic. Hsu and Lawrence (2016) found that the volume, valence, and growth rate of online word-of-mouth (WOM) exacerbate the negative effect of a product recall on firm value.

Meanwhile, product recalls also hurt product sales and the effectiveness of marketing instruments (Van Heerde et al. 2007). In fact, the negative impact of recalls on sales and advertising effectiveness

---

[5] See https://www.roche.com/guidelines/socialmedia.htm

[6] See https://www.arisglobal.com/extracting-safety-signals-social-media-artificial-intelligence/

is stronger when the recall is severe, is widely publicized, and involves higher-reputation brands (Liu and Shankar 2015). Borah and Tellis (2016) demonstrated the negative spillover effect in social media, whereby negative mentions about one product increase negative social media discussions of competing products. They also showed that social media significantly amplifies the negative effect of recalls on sales by about 4.5 times. Given the substantial costs induced by product recalls, researchers are strongly motivated to propose effective strategies to manage these recalls. To that end, Mukherjee and Chauhan (2019) built a game-theoretical model to derive the equilibrium advertising strategies of competing manufacturers when one or both firms can issue a product recall at a random time. Wang et al. (2020) explored the joint effect of competitive actions and social media related interactions on offline car sales after automobile recalls. Collectively, these studies show the critical impact of product recalls on firm performance outcomes.

The other stream of literature explores how firms' characteristics affect product recall decisions. Existing studies have shown that the likelihood of recall is affected by: 1) *intrinsic factors*, such as the research and development focus (Thirumalai and Sinha 2011), product variety, plant utilization (Shah et al. 2017), and the addition of female directors to a firms' board (Wowak et al. 2020); and 2) *extrinsic factors*, such as product competition (Ball et al. 2018b) and supply chain performance (Steven et al. 2014, Steven and Britto 2016). Several studies examined which factors affect the time to recall. By analyzing over 500 toy recalls within 15 years, Hora et al. (2011) concluded that recall timing is associated with the recall strategy, the type of product defect, and the supply chain entity that issued the recall. Eilert et al. (2017) found that brand characteristics, including a reputation for reliability and previous recall experiences, moderated the effect of problem severity on the time to recall. Wowak et al. (2020) demonstrated that firms that added female directors to the board make faster recall decisions for the most serious defects, which are high in severity and dangerous for customers.

Consumer voices, which contain valuable first-hand experience for detecting defects, are important information sources that can influence recall decisions. The exponential growth of social media usage

has made it easier for consumers to share unpleasant experiences. Meanwhile, firms and government agencies have been empowered by the development of AI technologies to monitor consumer voices on social media. More importantly, social media has become a platform for consumers to syndicate their individual voices to pressure firms and government agencies. Hence, the impact of consumer voices is considerably amplified by social media. However, to the best of our knowledge, no prior research has investigated how consumer voices on social media may influence the product recall process, and the current paper is the first to address this important question.

2.3. Social Media Literature

The existing research on firms' social media strategies can be categorized into three broad streams: brand marketing, customer service, and social media intelligence. The first stream empirically studies the impact of social media marketing strategies on consumer purchasing behaviors (e.g., Goh et al. 2013, Xie and Lee 2015, Gong et al. 2017) and customer engagement on social media (e.g., Miller and Tucker 2013, Lee et al. 2018). The second stream examines firms' social media customer service strategies (e.g., Gunarathne et al. 2018, Hu et al. 2019a) and their influence on customers' relationships with firms (e.g., Ma et al. 2015, Gunarathne et al. 2017). The third stream develops models and algorithms to understand and extract information on social media (e.g., Bhattacharya and Ram 2012, Dow et al. 2013, Culotta and Cutler 2016, Abbasi et al. 2018, Kanuri et al. 2018, Hu et al. 2019b, Abbasi et al. 2019, Velichety and Ram 2020). Our study is related to the social media intelligence literature in general and the social media pharmacovigilance literature in particular but differs significantly from the social media pharmacovigilance literature by linking social media pharmacovigilance to drug recall decisions.

Previous work on social media pharmacovigilance has predominately focused on ADR detection through the development of two types of methods: entity extraction and text classification. Entity extraction relies on lexicons and expert knowledge to extract drug names and corresponding ADRs from social media posts. Leaman and Wojtulewicz (2010) were the first to extract ADRs using a lexicon-based method. Sarker and Gonzalez (2015) created a comprehensive lexicon with 16,183

ADR mentions (including standard and non-standard terms) for over 10,000 drugs. In addition to extracting ADR mentions, several studies focused on identifying the relations between drugs and ADRs using association rule mining (Nikfarjam and Gonzalez 2011, Benton et al. 2011, Yeleswarapu et al. 2014, Yang et al. 2014).

Text classification aims to build supervised classifiers to predict whether a social media post mentions ADRs. Bian et al. (2012) collected two million tweets from Twitter and built two classifiers: one to predict if a Twitter user had used a drug and another to classify if a tweet contained an ADR. Yang et al. (2013) trained support vector machines (SVM) and naïve Bayes classifiers using a mixture of syntactic and semantic features. Sarker and Gonzalez (2015) incorporated features from distinct research areas and combined multiple corpora to boost classification performance.

Previous studies on social media pharmacovigilance have several limitations. First, ADR detection research has focused on a small number of drugs. Most studies investigated less than 100 drugs and specific domains of medicines, such as breast cancer or diabetes (Benton et al. 2011, Liu and Chen 2013). Second, there are a limited number of annotated instances in existing data sets. The most comprehensive data set, which was discussed by Ginn et al. (2014) and Sarker and Gonzalez (2015), only included 10,822 tweets with a binary feature indicating whether a tweet contained ADRs. In contrast, the present study investigates over 600 drugs and use annotated data that includes more than 30,000 tweets and 5,000 forum posts.

3. Hypothesis Development

Figure 1 illustrates the timeline from a drug's release to its final recall. After a drug is released to the market, the FDA and the pharmaceutical firm start monitoring ADR signals through pharmacovigilance. Once such signals are detected, it is the obligation of the FDA and the firm to verify the direct relationship between the drug and the ADR and to determine whether to issue a recall (Hauben and Aronson 2009). Therefore, conditional on an eventual drug recall, the duration between drug release and drug recall depends on the speed of the ADR signal detection as well as the speed of investigation and deliberation processes. The conceptual distinction between the two stages and the different roles social media might play in each stage lead to two hypothesized effects of social media on the overall speed of drug recalls: the *information effect* and the *publicity effect.*

3.1.   Information Effect

The information effect works mainly through social media pharmacovigilance, which detects signals reported by patients on social media; such signals are often unavailable through conventional channels. Table 2 contains several examples of social media posts related to ADRs.

There are two reasons why social media pharmacovigilance might accelerate signal detection. First, it may reduce the delay in ADR signal detection because social media is more convenient than FAERS for ADR reporting. This allows firms and the FDA to detect ADRs much earlier and faster than before (Ransohoff et al. 2018, Abbasi et al. 2019). In addition, there are no restrictions on the discussion of personal experiences with medications on social media. Almost every consumer has access to social media and can easily share first-hand experiences of drugs. In contrast, to report an ADR to FAERS, professionals and consumers must complete a mandatory official form. Even though the FDA introduced a more consumer-friendly reporting form, FDA3500B, in 2013, reporting an ADR to FAERS still takes significantly more time than posting on social media (Rose and Fritsch 2013, Toki and Ono 2018). Abbasi et al. (2019) suggested that ADR signals on social media are one to three years ahead of ADR signals through traditional methods, such as official event databases (e.g., FAERS).

Second, social media pharmacovigilance may detect more safety signals because the better accessibility of safety signals from social media than FAERS can, in turn, encourage more patients to report safety signals that would otherwise not surface. A full understanding of clinical reports requires domain knowledge, which consumers usually lack, and adverse event reports on FAERS are not readily accessible to consumers or physicians because accessing reports requires complicated actions, including data collection and data extraction. In contrast, social media can easily reveal drug-related information to consumers. A growing number of patients are actively sharing and posting health-related information in online forums and on social media (Chou et al. 2009). Such information can quickly spread within a community and motivate more patients with similar issues to share (Price 2016). Indeed, independent ADR signals from others can significantly strengthen an individual's

conviction about his/her suspected ADR, resulting in more sharing and more ADR signals. By interacting with and receiving prompt feedback from others, consumers are naturally more willing to share more personal experiences in the future. Empirically, Freifeld et al. (2014) found that there are nearly three times as many ADR discussions on Twitter than ADR reports to the FDA by consumers. According to Vilar et al. (2017), social media is particularly useful for inspecting drug-drug interactions, which are poorly explored using traditional sources.

The above arguments for the information effect conceptually demonstrate the value of social media pharmacovigilance. Both academic researchers and corporate managers seem to be convinced. As noted in the literature review, many academic researchers are developing algorithms to extract ADR signals from social media data (Leaman and Wojtulewicz 2010, Sarker and Gonzalez 2015, Nikfarjam et al. 2015, Bian et al. 2012). Oracle reported that social media is becoming an important data source through which pharmaceutical companies detect safety signals.[7] As another example, Advera Health Analytics, a global leader in pharmacovigilance software, added Booz Allen Hamilton's social media data and analysis to its existing drug safety analytics platform.[8]

3.2. Publicity Effect

In contrast to FAERS, which is akin to a repository of ADR reports from consumers and professionals, social media ADR posts are not only easy for the public to access but also straightforward for individuals to *disseminate*. This unique "social" aspect can be concerning for the involved pharmaceutical company because widely spread social media discussions of product defects can cause damaging negative publicity (Price 2016).

Previous research has established that negative publicity is particularly harmful to product ratings and sales. Moreover, consumers often consider negative information to be more informative than positive information when forming their overall evaluations of a firm/brand (Fiske 1980, Ahluwalia et al. 2000). For instance, Chevalier and Mayzlin (2006) showed that an incremental negative review

---

[7] See http://www.oracle.com/us/industries/health-sciences/address-data-challenges-pharma-wp-5018953.pdf

[8] See http://www.prweb.com/releases/2018/01/prweb15079580.htm

is more potent in reducing book sales on Amazon.com than an additional positive review is in increasing sales. We believe the same effect exists for prescription drugs. For physicians, who play a vital role in prescription drug demands, negative publicity of a drug updates their belief regarding drug quality (Kalra et al. 2011). For consumers, negative publicity of a drug lowers their trust of and satisfaction with the drug, which, in turn, affects their preference for the drug (Chintagunta et al. 2009). Along with lowering short-term product sales, negative publicity damages long-term brand equity, a valuable yet fragile asset defined as a composite measure of brand attitudes, perceptions of quality, brand trust, and brand purchase likelihood (Dawar and Pillutla 2000, Erdem and Swait 1998, Liu and Shankar 2015).

An effective defensive strategy against negative publicity is crucial for a firm in the midst of a product-harm crisis. A passive, no-response strategy will almost certainly fail in the social media age because it would solidify public suspicion of an irresponsible institution. The combination of a service failure and a failed intervention, which is referred to in the literature as double deviation, is much more explosive and riskier for firms (Bitner et al. 1990, Grégoire et al. 2015). Consumers may seek revenge against the firm by sharing their experiences or withdrawing from future interactions with the firm (Grégoire et al. 2009). Dutta and Pullig (2011) compared three commonly observed response strategies: *denial*, defined as the refusal to explain; *reduction-of-offensiveness*, defined as only offering an explanation; and *corrective action*, defined as the explanation and assurance of prevention. They found that corrective action was the most effective defensive strategy when negative publicity was mainly about product defects.

However, the effectiveness of a corrective action strategy crucially depends on the response speed (Henthorne and Henthorne 1994). By taking swift action, a firm can not only quickly stop the harmful impact of its defective product on consumers and, consequently, on its liability but also prevent further escalation of the simmering negative publicity, especially on social media. In summary, by amplifying the negative publicity of ADR events, we expect social media ADR discussions to put more pressure on firms to accelerate investigations and take more prompt action.

Based on our hypothesized information effect and publicity effect, we propose the following hypothesis for empirical testing in the context of drug recall:

Acceleration Hypothesis: *Social media ADR discussions accelerate the product recall process.*

4. Data, Variables, and the Empirical Model
4.1. Data Description

We use two FDA data sets:[9] *recall enforcement reports* and *adverse event reports.*[10] Each enforcement report consists of two major sections. The first section lists general information about the recall, such as when the drug was recalled and whether the recall was voluntary. The second section, which includes detailed drug information, especially the National Drug Code (NDC),[11] is critical for creating a recalled drugs list.

We use three exclusion criteria for the enforcement reports. First, enforcement reports without an NDC are excluded. This is an important criterion because we can leverage NDCs to find information in the FDA database for reports with missing details. Second, reports concerning OTC products are excluded for two reasons: 1) the definition of pharmacovigilance indicates that it is implemented to detect safety signals for prescription drugs only (World Health Organization 2002); and 2) consumers can purchase OTC drugs from pharmacy stores and supermarkets, making it hard to collect market size information, which is important for alleviating potential endogeneity in our model. Third, reports on products that were recalled multiple times are excluded.[12] The main reason for excluding products with multiple recalls is the lack of information on the recalls, such as the improvements made by firms

[9] We collected data from openFDA, an environment maintained by the FDA to make it easy for researchers and the public to access and use health data sets. openFDA provides open application programming interfaces (APIs) for a collection of FDA public data sets.

[10] According to the FDA, the agency started recording recall event information through its Recall Enterprise System since 2002, but openFDA only provides public reports dated after 2012. Therefore, we request the remaining enforcement reports from the FDA to ensure that our data set was complete.

[11] NDC is a unique product identifier used in the US for drugs intended for human use.

[12] A drug recall temporarily removes the drug from the market. The drug can return to the market with FDA approval. Therefore, a drug may be recalled multiple times.

to regain FDA approval. Since the effectiveness of those improvements might influence the number of ADR discussions on social media and the likelihood of the drug being recalled in the future, the inclusion of such products would introduce endogeneity issues into our analysis. After preprocessing the reports, our final sample consists of 629 drug recalls from 2002 to 2016. Figure 2 reports the distribution of drug recalls by year.

The adverse events database stores data collected from FAERS. For each of the 629 drugs, we collect corresponding adverse event reports from 2004 to 2016.[13] According to the FDA, reporting adverse events is voluntary in the US, and the reporting sources are mainly consumers and healthcare professionals. Figure 3 plots the distribution of adverse event reports from different sources by year and shows dramatic increases in the number of reports from healthcare professionals and consumers. It is worth noting that healthcare professionals were the dominant source of ADRs at first, however, after 2012, more reports were from consumers. This trend shows that consumers are playing an increasingly more important role in reporting ADRs and suggests the necessity of investigating the influence of consumer voices on product recalls.

As noted above, the drug-related social media data come from six online forums: *MedHelp*, *DailyStrength*, *SteadyHealth*, *MedsChat*, *Drugs.com*, and *DrugBuyersGuide*. These forums were launched in different years and each focuses on a different area. For instance, *Drugs.com*, an online pharmaceutical encyclopedia, is a drug-related information sharing platform, while *DailyStrength* is a social network where users support each other by sharing their experiences. Since firms and researchers also consider Twitter as an essential data source for social media pharmacovigilance, we use Twitter Advanced Search[14] and Selenium[15] to collect drug-related tweets from 2006 to 2016.

To identify relevant posts and tweets, we first compile a list of unique drug names as keywords. The FDA provides two names for each drug: the nonproprietary name (also known as the generic

---

[13] The adverse events database includes reports submitted to the FDA through FAERS since 2004.

[14] See https://twitter.com/search-advanced?lang=en

[15] Selenium is a Python package that can automatically execute the actions performed in a web browser, such as navigating to a website.

name) and the proprietary name (also known as the brand name). A nonproprietary name is assigned by the US Adopted Names Council and is usually a shorthand version of the drug's chemical name, structure, or formula, which are complicated and hard to remember. In contrast, a proprietary name is developed by a pharmaceutical company for marketing purposes and is relatively easy for healthcare professionals and consumers to remember. In this study, we use the proprietary name as the keyword for two reasons. First, since proprietary names are easier to remember than nonproprietary names, healthcare professionals tend to use the proprietary name for prescriptions and consumers also prefer the proprietary name for convenience.[16] Second, since multiple proprietary names can relate to the same nonproprietary name, using the nonproprietary name as the keyword would lead to inaccurate measurement of social media ADR discussions.

4.2. Variable

Table 3 provides a summary of the key variables and their definitions. The correlation matrix is included in Appendix A (Table A.1.).

The key construct in the acceleration hypothesis is social media ADR discussions. To quantify this construct, we use two variables, $adrPredictionsSVM_{i,t}$ and $adrPredictionsBERT_{i,t}$, that are based on different supervised machine learning algorithms to reduce the potential impact of measurement error. The first algorithm, SVM, performs significantly better than naïve Bayes and maximum entropy classifiers for ADR detection (Sarker and Gonzalez 2015). The second algorithm, bidirectional encoder representations from transformers (BERT), is a deep learning method that can learn contextual relations between words in a text (Devlin et al. 2018). It has achieved state-of-the-art performance on multiple natural language processing (NLP) tasks.

To implement these methods, we create data sets that include tweets/posts, which are manually labeled as ADR-related. We randomly and proportionally select tweets for each of the drugs in our sample to build a comprehensive data set, TW. There are 2,467 ADR-related instances out of 30,209

---

[16] See https://www.merckmanuals.com/home/drugs/brand-name-and-generic-drugs/overview-of-generic-drugs-and-drug-naming

tweets, which is close to the 10% ratio reported by Ginn et al. (2014). Similarly, Forums consists of

5,100 posts randomly selected for each drug. We describe the implementation of these two methods

in detail in Appendix B. Table 4 reports the predictive performances of SVM and BERT on each

data set. As shown, the BERT classifier performs much better than the SVM classifier with ADR

F-scores of 0.675 and 0.724, for the Forums and TW data sets, respectively.

We obtain three variables from FAERS: $FAERS_{i,t}$ assesses the total number of adverse events

reported to the FDA, $professionals_{i,t}$ captures the number of events reported by health professionals,

and $consumers_{i,t}$ reflects the number of events reported by consumers or non-health professionals.

We control for a set of time-varying covariates. $substitute_{i,t}$ is a binary variable indicating the exis-

tence of a prescription drug that is bioequivalent to drug i until month t. $clinicalTrials_{i,t}$ measures

the total number of clinical trials that have been completed until month t, collected from the US

National Library of Medicine.[17] News mentions are controlled because they may influence both the

time to recall and social media mentions. We use drug names as keywords to search historical news-

papers from NewsPaperArchive[18] and construct $news_{i,t}$ to measure the number of offline newspapers

that mention drug i at month t.[19] Another variable that can introduce bias is the market size of a

drug. We collect market size information from the MEPS summary tables,[20] which provide annual

statistics on total purchases of prescription medicines or therapeutic class groups from 1996 to 2016.

$marketSize_{c,t}$, measures the total purchases in thousands by drug i's therapeutic class c because the

MEPS summary tables do not cover most of the drugs in our sample.

We create a list of covariates to control for firm characteristics. $numRecalls_{j,t}$ measures the total

number of prescription drugs that have been recalled by firm j until month t, which captures the

---

[17] See https://www.clinicaltrials.gov/

[18] See https://newspaperarchive.com/

[19] To alleviate the concern about $news_{i,t}$ being an intermediate outcome, we conduct a robustness check by re-estimating our main model without $news_{i,t}$ as a control. The results are consistent with our main findings and are available upon request. We thank an anonymous reviewer for this suggestion.

[20] See https://meps.ahrq.gov/mepstrends/hc_pmed/

firm's prior experience with recalling drugs. We construct $numDrugs_{j,t}$ to assess the total number of OTC and prescription drugs marketed by firm j at month t, which indirectly reflects the firm size and its financial situation. Following Ball et al. (2018b), we construct $competition_{j,t}$ as the proxy for the product competition faced by firm j, which may affect its recall speed. We control for product complexity with $complexity_{j,t}$, which measures the ratio of prescription drugs to the total number of drugs.

Several drug features are also controlled. $abusive_i$ indicates whether a drug is addictive based on the Drug Enforcement Administration's (DEA) assessment. Drug labeling information is also collected. $adverseReactions_i$ is a binary variable indicating whether a label lists possible ADRs. $warnings_i$ equals 1 if drug i lists warnings on the label and 0 otherwise. The FDA asks manufacturers to list important and serious warnings in a black box section to draw customers' attention. Therefore, we use $boxedWarnings_i$ to indicate whether a drug's label has a boxed section for warnings. $priorityReview_i$ is equal to 1 if drug i received priority review status from the FDA, which means that the FDA completed the initial regulatory review within six months instead of the usual ten months.[21] $orphanDrug_i$ is a binary variable indicating whether drug i is designated as an orphan drug, which confers a longer period of market exclusivity to offset the small size of the patient population.[22] $wideDistribution_i$ is a binary variable indicating whether drug i is distributed to the whole nation. $approvalType_i$ is a binary variable indicating a drug's approval type and equals 1 if drug i went through a New Drug Application (NDA) process.

4.3. Empirical Model

The data set in this study is a time-to-event data set, with the recall of a drug as the event under study. Survival analysis is applied to study the time until the occurrence of an event. Here, we use discrete-time survival analysis since the observations take values over a discrete grid (i.e., months).

---

[21] See https://www.fda.gov/patients/fast-track-breakthrough-therapy-accelerated-approval-priority-review/priority-review

[22] See https://www.fda.gov/industry/developing-products-rare-diseases-conditions/designating-orphan-product-drugs-and-biological-products

Previous literature has suggested that discrete-time survival analysis is more appropriate when data is interval-censored (Singer and Willett 1993, Willett and Singer 1993, Jenkins 2005). This method has also been applied in information systems (IS) studies (Peng and Dey 2013, Greenwood et al. 2019).

Two discrete-time survival analysis specifications have been widely used in the literature: logit hazard and complementary log-log hazard (cloglog). Kalbfleisch and Prentice (2011) suggested that the cloglog hazard model is more appropriate when time is continuous but only observed in grouped form and is reported to be more robust. To alleviate reverse causality concerns, we follow Peng and Dey (2013) by using lagged independent variables in the survival analysis and obtain the following model on the *cloglog* of the hazard or conditional probability of recall at time t given survival up to that time:

$$cloglog\lambda_{i,t} = \delta_t + u_i + \beta_0 + \beta_1 ADR_{i,t-1} + \beta_2 professionals_{i,t-1} + \beta_3 consumers_{i,t-1} + \beta_4 X_{i,t-1}$$

$$+ \beta_5 Z_{j,t-1} + \beta_6 D_i$$

where $\delta_t = cloglog\lambda_{0,t}$ is the complementary log-log transformation of the baseline hazard. There are several options for the baseline hazard function. In the main analysis, we specify the baseline hazard function as the logarithm of time, although we later use alternative baselines for robustness checks. We are interested in β_1, the coefficient of $ADR_{i,t-1}$, which indicates the volume of ADR discussions on social media and is measured either by $adrPredictionsSVM_{i,t-1}$ or $adrPredictionsBERT_{i,t-1}$. We include a comprehensive set of control variables, including drugs' time-varying controls (i.e., $X_{i,t-1}$), company characteristics (i.e., $Z_{j,t-1}$), and drugs' time-invariant characteristics (i.e., D_i).

Because drugs were recalled only once in our data set, it is difficult to include fixed effect (Wooldridge 2010). To alleviate the concern that unobserved drug heterogeneity affects both social media mentions and the speed of drug recalls, we incorporate drug random effects, u_i, because using random effects to account for unobserved heterogeneity is appropriate if the sampled units are "drawn from a large population" (Greene 2003) and drugs in our sample are those recalled from the US over a long period.

5. Empirical Results
5.1. Main Results

Table 5 reports the regression results of the discrete-time survival analysis. Columns 1 - 5 report the results when the volume of ADR discussions on social media is measured by $adrPredictionsSVM_{i,t-1}$, and Columns 6 - 10 report the results when the volume is measured by $adrPredictionsBERT_{i,t-1}$. We gradually add different variables from FAERS. As shown in Columns 2 and 7, we include $FAERS_{i,t-1}$ to investigate the influence of FAERS on the speed of drug recalls. To explore the effects of adverse events reported by different sources, we add $professionals_{i,t-1}$ and $consumers_{i,t-1}$ separately (i.e., Columns 3, 4, 8, and 9) and then add them simultaneously (i.e., Columns 5 and 10). All key results are qualitatively the same across these specifications. The coefficients of $adrPredictionsSVM_{i,t-1}$, and $adrPredictionsBERT_{i,t-1}$ are consistently positive and significant ($p < 0.01$), suggesting that more ADR discussions on social media indeed increase the hazard rate of being recalled and thus decrease the time to recall. Hence, we find support for the acceleration hypothesis. In terms of the magnitude, an approximate quantification based on Column 10 suggests that 10 social media ADR posts per month can accelerate the recall process by reducing the duration from 105 months to 99 months, a 6 month or 6.5% reduction.[23]

As shown in Table 5, the coefficient of $FAERS_{i,t-1}$ is significantly positive ($p < 0.01$), indicating that FAERS, as an official platform, provides credible information for drug recalls. Moreover, the coefficients of $professionals_{i,t-1}$ remain positive and significant ($p < 0.01$) while the coefficients of $consumers_{i,t-1}$ become insignificant. This suggests that, compared to consumers, health professionals can draw more attention from pharmaceutical companies or the FDA and thus have a stronger impact when reporting adverse events in FAERS.

---

[23] The approximate quantification is implemented in the following manner. We use the econometric model to predict the duration of a drug in the following scenarios: 1) no social media ADR posts per month; and 2) 10 social media ADR posts per month. All the other covariates are fixed at their mean or median. The predicted duration for a drug in the first scenario is 105 months, while the predicted duration for the second scenario is 99 months.

5.2. ADR Heterogeneity

In the main analysis, we focused on how the volume of ADR discussions on social media affects drug recall speed. Here we conduct two sets of analyses to alleviate the concern of ADR heterogeneity.

First, we consider the severity of the ADRs as a confounding factor because severe ADRs are likely to trigger more social media mentions and drive faster recalls. To alleviate this concern, we follow Gottlieb et al. (2015) by defining severe ADRs as those ranked above the 95^{th} percentile in a ranking list of 2929 ADRs, which is evaluated by more than 2000 Amazon Mechanical Turk workers.[24] Next, we classify all ADR terms in the lexicon as severe or non-severe[25] and construct a variable, $ratioSevereADRSVM_{i,t-1}/ratioSevereADRBERT_{i,t-1}$, to measure the ratio of severe ADR predictions to total ADR predictions for drug i at month t. Columns 1 - 4 of Table 6 show that the coefficients of $adrPredictionsSVM_{i,t-1}$ and $adrPredictionsBERT_{i,t-1}$ remain significantly positive after controlling for ADR severity.

Second, we develop the idea of ADR embedding to control for "horizontal" differences between various ADRs. Using over 13,000 ADR terms provided by Sarker and Gonzalez (2015), we generate a drug-ADR matrix, in which each row corresponds to a drug recall and each column corresponds to an ADR term. Each cell of the matrix contains the number of social media posts that mention the corresponding ADR term regarding the corresponding drug recall over the entire sample period. Next, we use the singular value decomposition (SVD) to reduce the matrix to a certain number of dimensions that capture the most relevant ADR characteristics for each drug. Columns 5 - 12 in Table 6 report the results when 20 and 30 dimensions of ADR characteristics are included in the econometric model.[26] As shown, the coefficients of ADR discussions are significantly positive.

[24] We also try 90^{th} percentile and the results are qualitatively the same.

[25] The ideal way to measure the severity is to follow a set of criteria for the standardized classification of adverse effects, such as the Common Terminology Criteria for Adverse Events (CTCAE). However, this method requires domain knowledge and detailed reports provided by patients, which are absent in our study. Even though the ranking list provided by Gottlieb et al. (2015) is limited because Amazon Mechanical Turk workers are not experts with deep understandings of ADR severity, it provides a useful way for the construction of ADR severity variable.

[26] Our results are robust to alternative numbers of dimensions.

5.3. Alternative Survival Models

Accelerated failure time (AFT) models are a class of survival models that are considered as an alternative to hazard models. We estimate an often-used AFT model where the error term follows the Weibull distribution. The results, reported in Table 7, are consistent with our main finding. Using alternative error distributions (e.g., exponential, log-normal, and log-logistic distributions) results in qualitatively the same findings.

A polynomial in time and a non-parametric baseline are alternative specifications of the baseline hazard function for discrete-time survival analysis. For the polynomial specification, we use the square of time as the baseline hazard function. For the non-parametric baseline, we create a duration-interval-specific dummy variable for each month at risk. Columns 1 - 8 in Table 8 show that our results are robust to quadratic and non-parametric baseline hazard functions.

An alternative link function for discrete-time survival analysis is the logit function. Columns 9 - 12 of Table 8 report the estimation results where we replace *cloglog* with *logit*. Again, we obtain qualitatively the same findings.

5.4. Falsification Test

The causal mechanism of the acceleration effect is through the following chain of events: social media ADR discussions trigger a comprehensive investigation to understand the cause of the ADR thanks to social media pharmacovigilance; the investigation either leads to no evidence or to the uncovering of the exact cause of the ADR (e.g., biochemical factors); in the latter case, a drug recall may be initiated. This underlying causal chain suggests an idea of falsification test: if certain recalled drugs are unlikely to cause any adverse health reaction, then the aforementioned chain of causal relation cannot initiate and we should not be able to detect the acceleration effect. Otherwise, our results are likely driven by unobserved confounding factors. To construct such a falsification test, we take advantage of the FDA's recall classification. The FDA classifies a recall into one of the following three classes after the drug recall is initiated.[27]

---

[27] https://www.fda.gov/drugs/drug-recalls/fdas-role-drug-recalls.

- **Class I**: dangerous or defective products that predictably could cause serious health problems or even death.

- **Class II**: products that might cause a temporary health problem or pose only a slight threat of a serious nature.

- **Class III**: products that are unlikely to cause any adverse health reaction but violate FDA labeling or manufacturing laws.

By these definitions, Class III recalls are unlikely triggered by adverse events while Class I and II recalls can be triggered by adverse events. Hence, the causal mechanism underlying our acceleration hypothesis should be absent for Class III drug recalls. Therefore, if our econometric analysis can nevertheless detect a significant effect of social media ADR discussions on the speed of Class III drug recalls, our main results are probably driven by unobserved confounding factors.

To implement this falsification test, we split the data into two subsamples where the first subsample consists of Class I & II drug recalls, and the second subsample consists of Class III recalls. The summary statistics of each subsample are reported in Table C.1. of Appendix C. We estimate our main model on these two subsamples and report the estimation results in Table 9. We find that the coefficients of $adrPredictionsSVM_{i,t-1}$ and $adrPredictionsBERT_{i,t-1}$ for Class I and II recalls in Columns 1 and 2 are positively significant ($p < 0.05$), while the coefficients for Class III recalls, as shown in Columns 3 and 4, are statistically insignificant. Therefore, the falsification test results support our main finding.

5.5. Mechanism Test: Information Effect

To empirically test the validity of the proposed information effect, we first check whether an ADR mentioned in a social media post is described in the ADR section of the drug's label. If the mentioned ADR is not described in the label, we consider it to be a *new* ADR that could provide novel information for companies and government agencies.[28] The logic of this mechanism test is as follows. If the acceleration effect only exists due to public pressure, then new and known ADR discussions

---

[28] We thank an anonymous reviewer for this valuable suggestion.

24   Article submitted to *Information Systems Research*; manuscript no. (Please, provide the manuscript number!)

Authors' names blinded for peer review

should play the same role in the recall process. In other words, when both the volume of new and known ADR discussions are included in the econometric model, we should not find any difference in the coefficients of these two measures.

To implement this idea, we construct two variables, $adrPredictionsNew_{i,t}$ and $adrPredictionsKnown_{i,t}$, to measure the number of ADR-related social media posts that contain new and known ADRs respectively. Table 10 reports the estimation results of this mechanism test, in which we find significantly positive coefficients of $adrPredictionsSVMNew_{i,t-1}$ and $adrPredictionsBERTNew_{i,t-1}$ and statistically insignificant coefficients of $adrPredictionsSVMKnown_{i,t-1}$ and $adrPredictionsBERTKnown_{i,t-1}$. The lack of evidence indicating the effect of social media posts discussing known ADRs on recall speed, and the strong evidence of such an effect from new ADRs, suggest that the acceleration effect is at least partially driven by the underlying information effect. Indeed, if the publicity effect is the only driver of the acceleration effect, we would not expect to find the consistent pattern shown in Table 10.

5.6. Mechanism Test: Publicity Effect

To test the publicity effect, we focus on Twitter data because only Twitter provides the number of followers for each user and the number of likes and retweets for each tweet in order to measure the exposure of a tweet. We construct three variables as alternative measures of publicity: $avgLogFollowers_{i,t}$ is the average of the log-transformed number of followers, $avgLogLikes_{i,t}$ is the average of the log-transformed number of likes, and $avgLogRetweets_{i,t}$ is the average of the log-transformed number of retweets. Table C.2. in Appendix C reports the summary statistics of the key variables used in this analysis.

As shown in Table 11, the coefficient estimates of $adrPredictionsSVMTwitter_{i,t-1}$ and $adrPredictionsBERTTwitter_{i,t-1}$ remain significantly positive ($p < 0.01$), which is consistent with the previous findings. More interestingly, we find that the coefficients of the publicity measures (i.e., $avgLogFollowers_{i,t-1}$, $avgLogLikes_{i,t-1}$, and $avgLogRetweets_{i,t-1}$) are positive and statistically significant ($p < 0.01$), suggesting that online influence also matters. Since social media pharmacovigilance is

based on the content of social media posts, not the influence of the patients or the post, our hypothesized information effect should not depend on online influence levels.[29] Therefore, this evidence supports the existence of the publicity effect.

5.7. Heterogeneous Effect by Drug Distribution Type

The acceleration effect of social media on drug recalls may be heterogeneous over the drug distribution type (i.e., state-level and nationwide). Specifically, we expect the acceleration effect to be weaker for nationwide drugs for two reasons. First, the economic cost of recalling a widely distributed drug may be significantly higher than that of a state-level drug, which could hinder pharmaceutical companies from making a quick decision. Second, due to the larger volume of social media voices for widely distributed drugs, it can cost much more time for companies and government agencies to detect a safety signal and complete a comprehensive investigation.

To empirically examine the heterogeneous effect, we include an interaction term between social media ADR discussions and $wideDistribution_i$ in the econometric model. Table 12 shows that the coefficient estimates of ADR discussions remain positive and significant ($p < 0.01$), which is consistent with the main findings. Interestingly, the coefficients of the interaction term between ADR discussions and $wideDistribution_i$ are consistently negative and significant ($p < 0.10$), implying that the acceleration effect of social media is indeed weaker for drugs distributed nationwide.

5.8. Fake Social Media ADR Posts

As misinformation on social media is prevalent nowadays, it is important to discuss the potential implications of fake social media ADR reports for our study. In particular, we need to understand whether the potential existence of fake ADR reports may affect our empirical findings about the existence of the acceleration effect. To this end, we address the following two questions.

1. Will the existence of fake ADR posts lead to an overestimation of the acceleration effect through the channel of information effect?

---

[29] One caveat is that patients with different online influence levels may have different ADR signal precision. While we cannot empirically test this, we believe that patients with different online influence levels are unlikely to have significant differences in their medical knowledge.

Authors' names blinded for peer review

26   Article submitted to *Information Systems Research*; manuscript no. (Please, provide the manuscript number!)

2. Will the existence of fake ADR posts lead to an overestimation of the acceleration effect through the channel of publicity effect?

For the first question, we believe fake ADR reports, if existing in large amount, will inflate our main independent variables (i.e., *adrPredictionsSVM* and *adrPredictionBERT)*, but should not affect the dependent variable through the channel of information effect. The reason is that any ADR report has to be investigated by pharmaceutical companies or the FDA during the investigation stage (see Figure 1). Fake ADR posts will be filtered out, hence would not affect the recall speed through the information effect. In fact, investigating these fake posts may distract firms or the FDA and possibly even slow down the investigation of true ADR posts, if any. As such, the presence of fake ADR posts would only result in an underestimation, rather than an overestimation, of the true acceleration effect through the channel of information effect.

For the second question, we find that Google search returns no media report of fake social media ADR posts as of June 2021, which is several years after our sample period. This suggests that fake ADR posts, even if they exist, seem to have failed to generate much attention. Hence, their publicity effect is likely minimal too. We believe the rarity of fake ADR posts is not surprising because there is little incentive for anyone to publish fake ADR posts. Pharmaceutical companies are not interested in directly competing with each other. For brand drugs, companies are protected by the patent law from any competition. For generic drugs, companies indeed face intense competition, but they have no incentive to spread negative and false information about each other's product, because their products are therapeutically the same which makes them vulnerable to the negative spillover of fake ADR posts.

Therefore, our qualitative conclusion regarding the existence of the acceleration effect is unlikely affected by the possibility of fake ADR reports.

6. Conclusions

This study investigates whether and why social media pharmacovigilance influences the speed of drug recalls for prescription drugs that are eventually recalled. Conceptually, we propose two channels

through which social media may accelerate the drug recall process: the information effect and the publicity effect. Empirically, we find evidence that more social media ADR discussions lead to more prompt drug recalls. Furthermore, estimation results from two mechanism tests support the existence of both channels. Our results also suggest that the acceleration effect is stronger for drugs that are not nationally distributed. Overall, these findings offer new and important implications for the impact of social media on product recalls.

6.1. Contributions to Literature

The current paper makes three contributions to the literature. First, this study contributes to the product-harm crisis literature by investigating the important and timely question of whether consumer voices on social media can accelerate prescription drug recalls. To the best of our knowledge, the impact of social media on the speed of product recalls has been largely unexplored (see Table 1). The literature on product recalls has primarily focused on the effects of firm and product characteristics on product recall decisions (Thirumalai and Sinha 2011, Shah et al. 2017) and thus far has neglected the role played by consumer voices on social media. Given the increasingly important role of social media in society, this study fills this research gap and reveals for the first time the *acceleration* effect of social media on the speed of product recalls. The context of our study (i.e., the pharmaceutical industry) also enriches the product-harm crisis literature, which has mainly focused on other industries.

Second, this study contributes theoretically to the literature by proposing the information effect and the publicity effect as two channels through which social media influences the product recall process. There are several works in the IS literature investigating the impact of information from social media on business decisions other than product recalls (Xu and Zhang 2013, Goh et al. 2013, Wu 2013, Han et al. 2020). For example, Goh et al. (2013) focused on how information richness influences consumer purchase expenditure and Wu (2013) quantified the effects of information diversity on productivity and job security. However, few studies have investigated the impact of product-related information extracted from social media on product recalls. Similarly, although the link between publicity on

social media and certain business outcomes or decisions, such as product sales, has been explored (Chevalier and Mayzlin 2006, Berger et al. 2010, Ching et al. 2016, Borah and Tellis 2016, Gunarathne et al. 2018), no study has connected publicity on social media with product recall decisions. Our theoretical contribution of the information effect and the publicity effect is further strengthened by supportive empirical evidence in the context of drug recalls.

Third, as the first examination of how social media ADR discussions influence the speed of drug recalls, this study contributes to the social media pharmacovigilance literature in a uniquely important way. Previous literature on social media pharmacovigilance has almost exclusively focused on the development of data mining techniques for the identification of ADR discussions on social media. While Harpaz et al. (2016) and Price (2016) discussed the strengths and weaknesses of social media pharmacovigilance in a qualitative way, no study has attempted to empirically test the practical impact of social media pharmacovigilance on the speed of drug recalls, which is clearly important to justify the significant investments by firms, government agencies, and academic researchers who are devoting time and resources to improve social media pharmacovigilance techniques.

6.2. Contributions to Practice

Our paper offers important insights to practitioners. First, the empirical setting of our study suggests that ADR discussions on social media can speed up the recall process for drugs that are presumably harmful to patients and therefore to pharmaceutical companies. Although the adoption rate of social media pharmacovigilance is on the rise, it has not yet been widely adopted. According to a survey by Oracle, 22% of respondents have implemented social media pharmacovigilance by 2018.[30] Hence, our findings offer encouraging signs that managerially actionable insights are buried in ADR discussions on social media, and pharmaceutical companies should take advantage of this rich information source by actively monitoring and mining social media data.

For policymakers, the validation of the acceleration hypothesis is an important call for attention, regardless of the relative importance of the information effect and the publicity effect. The current

[30] For details, see http://www.oracle.com/us/industries/health-sciences/address-data-challenges-pharma-wp-5018953.pdf

system of monitoring and recalling drugs is slow and inefficient. For instance, the FDA did not begin the recall of a blood pressure drug that posed a cancer risk to users until July 13, 2018, four years after the first safety signal was detected.[31] The recalled drug was widely dispensed in US pharmacies, and more than 60 million patients suffered the risk of developing cancer from the slow drug recall process. Our findings suggest that there are certain complementary features of social media, compared to FAERS, that can speed up the recall process. Therefore, we recommend policymakers consider two improvements to FAERS. First, because safety signals may emerge on social media earlier than in FAERS and some may never reach FAERS, policymakers should consider the incorporation of social media data into the FAERS platform to more efficiently utilize this information source and to encourage more patients to share safety signals. Second, our analysis of the publicity effect suggests that consumer voices on social media present public pressure on firms to act more promptly. Policymakers can learn from this design feature by building "social functionalities" into FAERS. For example, they could add features so that patients can easily share their adverse event reports on social media, and also connect and discuss with other patients who report on FAERS.

Finally, despite our focus on the pharmaceutical industry, we believe that our findings are generalizable to industries with at least one of the following two characteristics. First, there need to be credible product safety signals on social media for the information effect to work. Taking the food industry as an example, Soylent was able to promptly recall its food bars thanks to consumer discussions on social media sites regarding their experiences of nausea, vomiting, and diarrhea. However, there are probably very few safety signals on social media related to industrial equipment or enterprise software, and firms in those industries likely would not gain much insight from the current study. Second, for the publicity effect to have a direct impact, firms in the industry must be susceptible to public pressure on social media. However, even without this susceptibility, the publicity effect might work indirectly through its impact on traditional media, which, in turn, can place public pressure

---

[31] During an interview with CNN, a spokeswoman for the FDA revealed that the first possible appearance of impurity was in 2014. For details, see https://www.cnn.com/2019/02/22/health/fda-recall-valsartan-arbs/index.html

on firms. Since many industries are investing heavily in social media in order to listen to customer voices and gather intelligence, and traditional media often pick up stories from social media, the mechanisms needed for the acceleration effect to work should exist in many industries. Therefore, we recommend companies in such industries consider their own "social media product vigilance" system.

6.3. Limitations and Future Research

This study has limitations that bear noting. First, like many recent papers that take advantage of unstructured data and machine learning algorithms, our detection of ADR posts on social media is imperfect. Although we have conducted multiple robustness tests and always use two different ADR measures to alleviate this concern, measurement error is almost inevitable and may bias our estimation if it is systematically different for drugs that are recalled at different speeds. Second, like many survival analyses of non-repeatable events, our econometric analysis is subject to bias due to unobserved confounding factors. Even though we have included a large number of control variables and conducted a series of robustness checks, the concern of omitted variable bias cannot be fully eliminated.

Given the established acceleration effect of social media on drug recalls, we believe the following research topics are worth future investigation. First, the acceleration effect might be heterogeneous over different types of social media users. For instance, since healthcare professionals are usually deemed a more reliable information source than consumers, their online discussions could draw more attention and thus play a stronger role in accelerating the recalls. Therefore, if richer data sets with more personal information become available, a future research direction is to examine the heterogeneous effect of ADR discussions from different types of social media users on the speed of product recalls. Second, because consumer voices on Facebook and other health-related forums, such as PatientsLikeMe, are not included in our sample due to a lack of data authorization, future studies using those sources would be particularly valuable because the motivation and behavior of consumers on Facebook and PatientsLikeMe could be different from consumers on the social media platforms considered in the current paper. Third, social media users may also discuss non-ADR topics about

drugs, such as the packaging, labeling, and efficacy of a drug. It will be interesting to investigate whether these discussions have an impact on pharmaceutical companies including but not limited to their recall decisions.

Of course, a much broader future research direction is studying the role of social media in the acceleration of product recalls across different industries or even in the acceleration of regulatory or legislative reform in public institutions. The loss of 346 lives due to the defect of the Boeing 737 MAX is a tragic reminder that an efficient and speedy process to identify product defects is of utmost importance, and we hope social media can be more widely utilized and studied to facilitate such a process.

Authors' names blinded for peer review

32   Article submitted to *Information Systems Research*; manuscript no. (Please, provide the manuscript number!)

References

Abbasi A, Li J, Adjeroh D, Abate M, Zheng W (2019) Don't mention it? Analyzing user-generated content signals for early adverse event warnings. *Information Systems Research* 30(3):1007–1028.

Abbasi A, Zhou Y, Deng S, Zhang P (2018) Text analytics to support sense-making in social media: A language-action perspective. *MIS Quarterly* 42(2):427–A38.

Ahluwalia R, Burnkrant RE, Unnava HR (2000) Consumer response to negative publicity: The moderating role of commitment. *Journal of Marketing Research* 37(2):203–214.

Ahmad SR (2003) Adverse drug event monitoring at the Food and Drug Administration: your report can make a difference. *Journal of General Internal Medicine* 18(1):57–60.

Appleby J, Krantz M (2004) Merck estimates $2.5 B impact from pulling Vioxx plug. *USA Today, September* 30.

Ball GP, Shah R, Donohue K (2018a) The decision to recall: A behavioral investigation in the medical device industry. *Journal of Operations Management* 62:1–15.

Ball GP, Shah R, Wowak KD (2018b) Product competition, managerial discretion, and manufacturing recalls in the US pharmaceutical industry. *Journal of Operations Management* 58:59–72.

Benton A, Ungar L, Hill S, Hennessy S, Mao J, Chung A, Leonard CE, Holmes JH (2011) Identifying potential adverse effects using the web: A new approach to medical hypothesis generation. *Journal of Biomedical Informatics* 44(6):989–996.

Berger J, Sorensen AT, Rasmussen SJ (2010) Positive effects of negative publicity: When negative reviews increase sales. *Marketing Science* 29(5):815–827.

Bhattacharya D, Ram S (2012) Sharing news articles using 140 characters: A diffusion analysis on Twitter. *2012 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, 966–971.

Bian J, Topaloglu U, Yu F (2012) Towards large-scale twitter mining for drug-related adverse events. *Proceedings of the 2012 international workshop on smart health and wellbeing*, 25–32.

Bitner MJ, Booms BH, Tetreault MS (1990) The service encounter: diagnosing favorable and unfavorable incidents. *Journal of Marketing* 54(1):71–84.

Authors' names blinded for peer review
Article submitted to *Information Systems Research*; manuscript no. (Please, provide the manuscript number!)

33

Borah A, Tellis GJ (2016) Halo (spillover) effects in social media: do product recalls of one brand hurt or help rival brands? *Journal of Marketing Research* 53(2):143–160.

Chen Y, Ganesan S, Liu Y (2009) Does a firm's product-recall strategy affect its financial value? An examination of strategic alternatives during product-harm crises. *Journal of Marketing* 73(6):214–226.

Chevalier JA, Mayzlin D (2006) The effect of word of mouth on sales: online book reviews. *Journal of Marketing Research* 43(3):345–354.

Ching AT, Clark R, Horstmann I, Lim H (2016) The effects of publicity on demand: The case of anti-cholesterol drugs. *Marketing Science* 35(1):158–181.

Chintagunta PK, Jiang R, Jin GZ (2009) Information, learning, and drug diffusion: The case of Cox-2 inhibitors. *Quantitative Marketing and Economics* 7(4):399–443.

Chou WYS, Hunt YM, Beckjord EB, Moser RP, Hesse BW (2009) Social media use in the United States: Implications for health communication. *Journal of Medical Internet Research* 11(4):48.

Cleeren K, Dekimpe MG, Helsen K (2008) Weathering product-harm crises. *Journal of the Academy of Marketing Science* 36(2):262–270.

Cleeren K, van Heerde HJ, Dekimpe MG (2013) Rising from the ashes: How brands and categories can overcome product-harm crises. *Journal of Marketing* 77(2):58–77.

Culotta A, Cutler J (2016) Mining brand perceptions from Twitter social networks. *Marketing Science* 35(3):343–362.

Dawar N, Pillutla MM (2000) Impact of product-harm crises on brand equity: The moderating role of consumer expectations. *Journal of Marketing Research* 37(2):215–226.

Devlin J, Chang MW, Lee K, Toutanova K (2018) Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* .

Dow PA, Adamic L, Friggeri A (2013) The anatomy of large facebook cascades. *Proceedings of the International AAAI Conference on Web and Social Media* 7(1).

Downing NS, Shah ND, Aminawung JA, Pease AM, Zeitoun JD, Krumholz HM, Ross JS (2017) Postmarket safety events among novel therapeutics approved by the US Food and Drug Administration between 2001 and 2010. *JAMA* 317(18):1854.

Dutta S, Pullig C (2011) Effectiveness of corporate responses to brand crises: The role of crisis type and response strategies. *Journal of Business Research* 64(12):1281–1287.

Edwards IR, Aronson JK (2000) Adverse drug reactions: Definitions, diagnosis, and management. *Lancet* 356(9237):1255–1259.

Eilert M, Jayachandran S, Kalaignanam K, Swartz TA (2017) Does it pay to recall your product early? An empirical investigation in the automobile industry. *Journal of Marketing* 81(3):111–129.

Erdem T, Swait J (1998) Brand equity as a signaling phenomenon. *Journal of Consumer Psychology* 7(2):131–157.

Fiske ST (1980) Attention and weight in person perception: The impact of negative and extreme behavior. *Journal of Personality and Social Psychology* 38(6):889.

Forster AJ, Jennings A, Chow C, Leeder C, van Walraven C (2012) A systematic review to evaluate the accuracy of electronic adverse drug event detection. *Journal of the American Medical Informatics Association* 19(1):31–38.

Freifeld CC, Brownstein JS, Menone CM, Bao W, Filice R, Kass-Hout T, Dasgupta N (2014) Digital drug safety surveillance: Monitoring pharmaceutical products in Twitter. *Drug Safety* 37(5):343–350.

Gao H, Xie J, Wang Q, Wilbur KC (2015) Should ad spending increase or decrease before a recall announcement? The marketing–finance interface in product-harm crisis management. *Journal of Marketing* 79(5):80–99.

Ginn R, Pimpalkhute P, Nikfarjam A, Patki A, O'connor K, Sarker A, Smith K, Gonzalez G (2014) Mining Twitter for adverse drug reaction mentions: A corpus and classification benchmark. *In Proceedings of the fourth workshop on building and evaluating resources for health and biomedical text processing.*, 1–8.

Goh KY, Heng CS, Lin Z (2013) Social media brand community and consumer behavior: Quantifying the relative impact of user-and marketer-generated content. *Information Systems Research* 24(1):88–107.

Gong S, Zhang J, Zhao P, Jiang X (2017) Tweeting as a marketing tool: A field experiment in the TV industry. *Journal of Marketing Research* 54(6):833–850.

Gottlieb A, Hoehndorf R, Dumontier M, Altman RB (2015) Ranking adverse drug reactions with crowd-sourcing. *Journal of Medical Internet Research* 17(3):e80.

Greene WH (2003) *Econometric analysis* (Pearson Education India).

Greenwood BN, Ganju KK, Angst CM (2019) How does the implementation of enterprise information systems affect a professional's mobility? An empirical study. *Information Systems Research* 30(2):563–594.

Grégoire Y, Salle A, Tripp TM (2015) Managing social media crises with your customers: The good, the bad, and the ugly. *Business Horizons* 58(2):173–182.

Grégoire Y, Tripp TM, Legoux R (2009) When customer love turns into lasting hate: the effects of relationship strength and time on customer revenge and avoidance. *Journal of Marketing* 73(6):18–32.

Guerini M, Gatti L, Turchi M (2013) Sentiment analysis: How to derive prior polarities from SentiWordNet. *arXiv preprint arXiv:1309.5843* .

Gunarathne P, Rui H, Seidmann A (2017) Whose and what social media complaints have happier resolutions? Evidence from Twitter. *Journal of Management Information Systems* 34(2):314–340.

Gunarathne P, Rui H, Seidmann A (2018) When social media delivers customer service: Differential customer treatment in the airline industry. *MIS Quarterly* 42(2):489–520.

Han Y, Lappas T, Sabnis G (2020) The Importance of Interactions Between Content Characteristics and Creator Characteristics for Studying Virality in Social Media. *Information Systems Research* 31(2):576–588.

Harpaz R, DuMochel W, Shah NH (2016) Big data and adverse drug reaction detection. *Clinical Pharmacology & Therapeutics* 99(3):268–270.

Harpaz R, DuMouchel W, Shah NH, Madigan D, Ryan P, Friedman C (2012) Novel data-mining methodologies for adverse drug event discovery and analysis. *Clinical Pharmacology & Therapeutics* 91(6):1010–1021.

Hauben M, Aronson JK (2009) Defining 'signal'and its subtypes in pharmacovigilance based on a systematic review of previous definitions. *Drug Safety* 32(2):99–110.

Haunschild PR, Rhee M (2004) The role of volition in organizational learning: The case of automotive product recalls. *Management Science* 50(11):1545–1560.

He S, Rui H, Whinston AB (2017) Social media strategies in product-Harm Crises. *Information Systems Research* 29(2):362–380.

Henthorne BH, Henthorne TL (1994) The tarnished image: Anticipating and minimizing the impact of negative publicity in health services organizations. *Journal of Consumer Marketing* 11(3):44–54.

Hora M, Bapuji H, Roth AV (2011) Safety hazard and time to recall: The role of recall strategy, product defect type, and supply chain player in the U.S. toy industry. *Journal of Operations Management* 29(7-8):766–777.

Hsu L, Lawrence B (2016) The role of social media and brand equity during a product recall crisis: A shareholder value perspective. *International Journal of Research in Marketing* 33(1):59–77.

Hu Y, Tafti A, Gal D (2019a) Read this, please? The role of politeness in customer service engagement on social media. *Proceedings of the 52nd Hawaii International Conference on System Sciences.*

Hu Y, Xu A, Hong Y, Gal D, Sinha V, Akkiraju R (2019b) Generating business intelligence through social media analytics: Measuring brand personality with consumer-, employee-, and firm-generated content. *Journal of Management Information Systems* 36(3):893–930.

Jarrell G, Peltzman S (1985) The impact of product recalls on the wealth of sellers. *Journal of Political Economy* 93(3):512–536.

Jenkins SP (2005) Survival analysis. *Unpublished manuscript, Institute for Social and Economic Research, University of Essex, Colchester, UK* 42:54–56.

Kalaignanam K, Kushwaha T, Eilert M (2013) The impact of product recalls on future product reliability and future accidents: Evidence from the automobile industry. *Journal of Marketing* 77(2):41–57.

Kalbfleisch JD, Prentice RL (2011) *The statistical analysis of failure time data* (John Wiley & Sons).

Kalra A, Li S, Zhang W (2011) Understanding responses to contradictory information about products. *Marketing Science* 30(6):1098–1114.

Kanuri VK, Chen Y, Sridhar SH (2018) Scheduling content on social media: Theory, evidence, and application. *Journal of Marketing* 82(6):89–108.

Lazarou J, Pomeranz BH, Corey PN (1998) Incidence of adverse drug reactions in hospitalized patients: a meta-analysis of prospective studies. *JAMA* 279(15):1200–1205.

Leaman R, Wojtulewicz L (2010) Towards internet-age pharmacovigilance: extracting adverse drug reactions from user posts to health-related social networks. *Proceedings of the 2010 Workshop on Biomedical Natural Language Processing* 117–125.

Lee D, Hosanagar K, Nair HS (2018) Advertising content and consumer engagement on social media: evidence from Facebook. *Management Science* 64(11):5105–5131.

Lee LF, Hutton AP, Shu S (2015) The Role of Social Media in the Capital Market: Evidence from Consumer Product Recalls. *Journal of Accounting Research* 53(2):367–404.

Liu X, Chen H (2013) AZDrugMiner: an information extraction system for mining patient-reported adverse drug events in online patient forums. *International Conference on Smart Health*, 134–150.

Liu X, Yong Liu, Ting Luo (2016) What drives a firm's choice of product recall remedy? The impact remedy cost, product hazard, and the CEO. *Journal of Marketing* 80(3):79–95.

Liu Y, Shankar V (2015) The dynamic impact of product-harm crises on brand preference and advertising effectiveness: An empirical analysis of the automobile industry. *Management Science* 61(10):2514–2535.

Liu Y, Shankar V, Yun W (2017) Crisis management strategies and the long-term effects of product recalls on firm value. *Journal of Marketing* 81(5):30–48.

Ma L, Sun B, Kekre S (2015) The squeaky wheel gets the grease—An empirical analysis of customer voice and firm intervention on Twitter. *Marketing Science* 34(5):627–645.

Marcus A (1989) The deterrent to dubious corporate behavior: Profitability, probability and safety recalls. *Strategic Management Journal* 10(3):233–250.

Miller AR, Tucker C (2013) Active social media management: The case of health care. *Information Systems Research* 24(1):52–70.

Mukherjee A, Chauhan SS (2019) The impact of product recall on advertising decisions and firm profit while envisioning crisis or being hazard myopic. *European Journal of Operational Research* 288(3):953–970.

Mukherjee UK, Sinha KK (2018) Product Recall Decisions in Medical Device Supply Chains: A Big Data Analytic Approach to Evaluating Judgment Bias. *Production & Operations Management* 27(10):1816–1833.

Nikfarjam A, Gonzalez GH (2011) Pattern mining for extraction of mentions of Adverse Drug Reactions from user comments. *AMIA Annual Symposium Proceedings* 2011:1019.

Nikfarjam A, Sarker A, O'Connor K, Ginn R, Gonzalez G (2015) Pharmacovigilance from social media: Mining adverse drug reaction mentions using sequence labeling with word embedding cluster features. *Journal of the American Medical Informatics Association* 22(3):671–681.

Peng G, Dey D (2013) Research Note—A dynamic view of the impact of network structure on technology adoption: The case of OSS development. *Information Systems Research* 24(4):1087–1099.

Price J (2016) What can big data offer the pharmacovigilance of orphan drugs? *Clinical Therapeutics* 38(12):2533–2545.

Ransohoff JD, Nikfarjam A, Jones E, Loew B, Kwong BY, Sarin KY, Shah NH (2018) Detecting chemotherapeutic skin adverse reactions in social health networks using deep learning. *JAMA Oncology* 4(4):581–583.

Rose BJ, Fritsch BF (2013) FDA's MedWatch program turns 20: What's new? *Pharmacy Today* 19(10):68–69.

Rubel O, Naik PA, Srinivasan S (2011) Optimal advertising when envisioning a product-harm crisis. *Marketing Science* 30(6):1048–1065.

Sarker A, Gonzalez G (2015) Portable automatic text classification for adverse drug reaction detection via multi-corpus training. *Journal of Biomedical Informatics* 53:196–207.

Sarker A, Molla D, Paris C (2013) Automatic prediction of evidence-based recommendations via sentence-level polarity classification. *Proceedings of the Sixth International Joint Conference on Natural Language Processing*, 712–718.

Shah R, Ball GP, Netessine S (2017) Plant operations and product recalls in the automotive industry: An empirical investigation. *Management Science* 63(8):2439–2459.

Singer JD, Willett JB (1993) It's about time: Using discrete-time survival analysis to study duration and the timing of events. *Journal of Educational and Behavioral Statistics* 18(2):155–195.

Steven AB, Britto RA (2016) Emerging market presence, inventory, and product recall linkages. *Journal of Operations Management* 46:55–68.

Steven AB, Dong Y, Corsi T (2014) Global sourcing and quality recalls: An empirical study of outsourcing-supplier concentration-product recalls linkages. *Journal of Operations Management* 32(5):241–253.

Stoll JD (2019) Did Twitter help ground the Boeing 737 MAX? *Wall Street Journal* .

Thirumalai S, Sinha KK (2011) Product recalls in the dedical device industry: An empirical exploration of the sources and financial consequences. *Management Science* 57(2):376–392.

Toki T, Ono S (2018) Spontaneous reporting on adverse events by consumers in the United States: An analysis of the Food and Drug Administration Adverse Event Reporting System database. *Drugs - Real World Outcomes* 5(2):117–128.

Van Heerde H, Helsen K, Dekimpe MG (2007) The Impact of a product-harm crisis on marketing effectiveness. *Marketing Science* 26(2):230–245.

Velichety S, Ram S (2020) Finding a needle in the haystack-Recommending online communities on social media platforms using network and design science. *Forthcoming at JAIS* .

Vilar S, Friedman C, Hripcsak G (2017) Detection of drug–drug interactions through data mining studies using clinical sources, scientific literature and social media. *Briefings in Bioinformatics* 19(5):863–877.

Wang YY, Wang T, Calantone R (2020) The effect of competitive actions and social media perceptions on offline car sales after automobile recalls. *International Journal of Information Management* 102257.

Willett JB, Singer JD (1993) Investigating onset, cessation, relapse, and recovery: why you should, and how you can, use discrete-time survival analysis to examine event occurrence. *Journal of Consulting and Clinical Psychology* 61(6):952.

Wooldridge JM (2010) *Econometric analysis of cross section and panel data* (MIT press).

World Health Organization (2002) The importance of pharmacovigilance. *World Health Organization* .

Wowak KD, Ball GP, Post C, Ketchen DJ (2020) The Influence of Female Directors on Product Recall Decisions. *Manufacturing & Service Operations Management* .

Wu L (2013) Social Network Effects on Productivity and Job Security: Evidence from the Adoption of a Social Networking Tool. *Information Systems Research* 24(1):30–51.

Xie K, Lee YJ (2015) Social media and brand purchase: Quantifying the effects of exposures to earned and owned social media activities in a two-Stage decision making model. *Journal of Management Information Systems* 32(2):204–238.

Xu R, Wang Q (2014) Large-scale combining signals from both biomedical literature and the FDA Adverse Event Reporting System (FAERS) to improve post-marketing drug safety signal detection. *BMC bioinformatics* 15(1):17.

Xu SX, Zhang XM (2013) Impact of Wikipedia on Market Information Environment: Evidence on Management Disclosure and Investor Reaction. *MIS Quarterly* 37(4):1043–1068.

Yang CC, Yang H, Jiang L (2014) Postmarketing drug safety surveillance using publicly available health-consumer-contributed content in social media. *ACM Transactions on Management Information Systems* 51(1):1–21.

Yang M, Wang X, Kiang MY (2013) Identification of consumer adverse drug reaction messages on social media. *PACIS*, 193.

Yeleswarapu S, Rao A, Joseph T, Saipradeep VG, Srinivasan R (2014) A pipeline to extract drug-adverse event pairs from multiple data sources. *BMC Medical Informatics and Decision Making* 14(1):13.

Zhao Y, Zhao Y, Helsen K (2011) Consumer learning in a turbulent market environment: Modeling consumer choice dynamics after a product-harm crisis. *Journal of Marketing Research* 48(2):255–267.

Figure 1   **Conceptual Illustration**

Figure 2   **Drug Recalls By Year**


Figure 3   **Adverse Event Reports By Year**

Table 1   Positioning and Contribution of This Paper in Product-Harm Crisis (Firm Level) Literature

Study	Social Media	Pre-crisis/Post-crisis	Industry	Dependent Variable(s)	Findings
Haunschild and Rhee (2004)		Post-crisis	Automotive	• Number of recalls	Voluntary recalls result in more learning than mandated recalls when learning is measured as a reduction in subsequent involuntary recalls.
Van Heerde et al. (2007)		Post-crisis	Food	• Brand sales	A brand suffers a loss in sales and the effectiveness of marketing instruments from a product-harm crisis. It also becomes more vulnerable to competitive effects.
Chen et al. (2009)		Post-crisis	Various industries	• Abnormal returns	Proactive strategies have a more negative effect on firm value than more passive strategies.
Rubel et al. (2011)		Post-crisis	Automotive	• Brand sales	Managers should reduce advertising spending before the crisis but advertise more after it occurs.
Thirumalai and Sinha (2011)		Both	Medical devices	• Number of recalls • Abnormal returns	The magnitude of financial consequences of device recalls is affected by the product scope, sales, growth prospects, and the capital structure of a firm. The likelihood of recalls decreases with prior recall experience.
Hora et al. (2011)		Pre-crisis	Toys	• Time to recall	The time to recall is associated with the recall strategy, the type of product defect, and the supply chain entity.
Cleeren et al. (2013)		Post-crisis	Consumer goods	• Brand share • Category purchases	The effectiveness of advertising and price depends on blame and publicity.
Kalaignanam et al. (2013)		Post-crisis	Automotive	• Product reliability • Number of injuries • Number of recalls	Increases in recall magnitude lead to decreases in future number of injuries and recalls. The effect is moderated by several factors.
Liu and Shankar (2015)		Post-crisis	Automotive	• Market share	The negative impact of recalls on brand sales and advertising effectiveness is stronger when the recall is severe, faces greater attention, and involves higher-quality brands.
Steven et al. (2014)		Post-crisis	Consumer goods	• Number of recalls	Offshore outsourcing has a greater impact on recalls than offshoring without outsourcing; outsourcing domestically has the least influence. Outsourcing to a smaller supplier base may lead to fewer recalls at low levels of outsourcing.
Lee et al. (2015)	✓	Post-crisis	Consumer goods	• Abnormal return	Corporate social media attenuates the negative price reaction to product recall announcements. As social media channel become more active, the attenuation benefits lessened but are not fully lost. Increased frequency of tweets by other users exacerbates the negative market reaction while increased frequency of tweets by the firm attenuates the negative market reaction.
Gao et al. (2015)		Pre-crisis	Automotive	• Abnormal return	Adjustments to a firm's pre-recall advertising expenditure can either mitigate or amplify the negative effect of the recall on stock market value, depending on the direction of advertising adjustment and the recall characteristic.
Borah and Tellis (2016)	✓	Post-crisis	Automotive	• Brand sales	Negative publicity caused by product recalls have spillover effects and amplifies the negative impact of recalls on sales.
Liu et al. (2016)		Post-crisis	Consumer goods	• Remedy choice	Companies prefer to avoid full remedy when remedy cost is high, yet they are more likely to provide full remedy for more severe product hazards. CEO's personal interests interfere with remedy decisions.
Hsu and Lawrence (2016)	✓	Post-crisis	Various industries	• Abnormal returns	Online WOM exacerbate the negative effect of a product recall on firm value. The negative effect is lower for high-equity brands.
Steven and Britto (2016)		Pre-crisis	Consumer goods Medical products	• Number of recalls	Outsourcing to emerging markets increase recalls and that sales penetration into emerging markets reduces recalls. In-house offshoring to emerging markets appears to mitigate the positive relationship between outsourcing to these markets and recalls.

Notes. This review of the relevant literature demonstrates the uniqueness of our study in the context of product-harm crisis. The analyses in those papers are from firms' perspectives. The other stream of literature aims to analyze the impact of crises on consumers' perceptions of a brand or their purchase intentions (Dawar and Pillutla 2000, Cleeren et al. 2008, Zhao et al. 2011). Those studies are not included in the table since they are not directly related to our work.

Table 1   Continued - Positioning and Contribution of This Paper in Product-Harm Crisis (Firm Level) Literature

Study	Social Media	Pre-crisis/Post-crisis	Industry	Dependent Variable(s)	Findings
Shah et al. (2017)		Post-crisis	Automotive	• Number of recalls	Manufacturing-related recalls are positively associated with product variety and plant utilization, but not with plant variety.
Liu et al. (2017)		Post-crisis	Automotive	• Abnormal returns	The negative impact of product recall volume lingers over time. Moderating effects of voluntary recall initiation and post-recall remedial efforts differ between short- and long-run evaluation windows.
He et al. (2017)	✓	Post-crisis	Airlines	• Marketing tweets	Nonfocal airlines increase defensive marketing efforts but decrease offensive marketing efforts after the product-harm crisis. The strategic adjustment is attenuated by the competition between nonfocal airlines and the focal one.
Eilert et al. (2017)		Pre-crisis	Automotive	• Time to recall • Abnormal returns	Problem severity increases time to recall. The relationship is moderated by the brand's reliability and intensity of recalls in the past. The time to recall is negatively related to stock market performance.
Mukherjee and Sinha (2018)		Pre-crisis	Medical devices	• Judgement bias in recall decisions	Noisy signals in user feedback (i.e., high noise-to-signal ratio) are associated with under-reaction likelihood. User feedback related to adverse events characterized by high severity is associated with high over-reaction likelihood.
Ball et al. (2018a)		Pre-crisis	Medical devices	• Likelihood to recall	Physicians' ability to detect a defect prior to product use decrease the likelihood to recall, perceived patient customer harm and perceived cost of the recall both mediate the relationship; managers' understanding of the root cause of the defect increases the likelihood to recall, perceived cost of the recall mediates the relationship.
Ball et al. (2018b)		Post-crisis	Pharmaceutical	• Number of recalls	Firms that face more product competition have significantly more manufacturing-related recalls. The product competition-recalls relationship is contingent on managerial discretion.
Mukherjee and Chauhan (2019)		Post-crisis	-	• Advertising decisions • Firm profit	Advertising levels of firms should differ in pre-recall and post-recall regimes depending on the impact and likelihood of the recall. A product-recall with a minor impact can increase the focal firm's long-term expected profit.
Wang et al. (2020)	✓	Post-crisis	Automotive	• Brand sales	Fewer competitive actions by the recalling firm are more effective in influencing car sales. Although social media discussions can increase awareness and sales, the change in perceptions of the brand on social media can negatively affect sales. The positive effect of social media is amplified when the recalling firm conducts fewer competitive actions.
Wowak et al. (2020)		Both	Medical products	• Number of recalls • Time to recall	Adding female directors to a firm's board of directors leads to more medical product recalls that are low in severity and accelerates recall decisions for the most serious defects that are high in severity and dangerous for customers.
The current study	✓	Pre-crisis	Pharmaceutical	• Time to recall	Consumer voices on social media can reduce the time to recall through the information effect and the publicity effect.

Notes. This review of the relevant literature demonstrates the uniqueness of our study in the context of product-harm crisis. The analyses in those papers are from firms' perspectives. The other stream of literature aims to analyze the impact of crises on consumers' perceptions of a brand or their purchase intentions (Dawar and Pillutla 2000, Cleeren et al. 2008, Zhao et al. 2011). Those studies are not included in the table since they are not directly related to our work.

Authors' names blinded for peer review

Article submitted to *Information Systems Research*; manuscript no. (Please, provide the manuscript number!)

Table 2   Sample Posts Related to ADRs on Social Media

Drug	Content	Time	Source
Brilinta	Anyone else have internal bleeding with #Brilinta? Is there a class action suit for Brilinta making people have internal bleeding yet?	2017-03-01	Twitter
Brilinta	I had 3 stents put in 3-1/2 months ago doc put me on brilinta. I had uncontrollable bladder bleeding felt shaky and weak and dizzy. Doc changed me to plavix side affects gone.	2017-04-25	SteadyHealth
Warfarin Sodium	Since being on warfarin sodium I have experienced all over joint pain and stiffness which I did not have before	2013-03-16	Twitter
Benicar	Benicar may be responsible for rapid weight loss which may require medical intervention 1-877-292-1500	2014-06-28	Twitter

Table 3   **Summary Statistics**

Variable	Mean	S. D.	Definition
Social Media Discussions			
$adrPredictionsSVM_{i,t}$	6.36	27.72	Number of tweets and posts that mention drug i's proprietary name and are predicted by the SVM classifier to be ADR-related at month t
$adrPredictionsBERT_{i,t}$	5.26	19.47	Number of tweets and posts that mention drug i's proprietary name and are predicted by the BERT classifier to be ADR-related at month t
FAERS			
$FAERS_{i,t}$	41.34	147.01	Number of adverse events that are related to drug i and reported to FAERS at month t
$professionals_{i,t}$	18.87	43.63	Number of adverse events that are related to drug i and reported by professionals to FAERS at month t
$consumers_{i,t}$	19.29	124.35	Number of adverse events that are related to drug i and reported by consumers to FAERS at month t
Time-varying Controls			
$substitute_{i,t}$	0.85	0.36	Binary variable indicating whether drug i has a bioequivalent product
$clinicalTrials_{i,t}$	2.93	6.05	Total number of clinical trials that have done for drug i till month t
$news_{i,t}$	3.00	9.51	Number of offline newspapers that mentioned drug i at month t
$marketSize_{c,t}$	350.97	241.72	Total purchases in millions at month t by the therapeutic class c of drug i
Firm Characteristics			
$numRecalls_{j,t}$	3.16	5.93	Total number of prescription drugs that have been recalled by firm j till month t
$numDrugs_{j,t}$	83.85	95.18	Total number of drugs marketed by firm j at month t
$competition_{j,t}$	0.67	0.37	Ratio of Abbreviated New Drug Application (ANDA) drugs to the total number of drugs marketed by firm j at month t
$complexity_{j,t}$	0.97	0.07	Ratio of prescription drugs to the total number of drugs marketed by firm j at month t
Drug Characteristics			
$abusive_i$	0.10	0.29	Binary variable indicating whether drug i has the potential for abuse based on DEA's assessment
$adverseReactions_i$	0.75	0.43	Binary variable indicating whether the label of drug i lists possible adverse reactions
$warnings_i$	0.74	0.44	Binary variable indicating whether the label of drug i lists warnings
$boxedWarnings_i$	0.30	0.46	Binary variable indicating whether the label of drug i lists warnings in a boxed section
$priorityReview_i$	0.10	0.30	Binary variable indicating whether drug i is approved through the priority review
$orphanDrug_i$	0.02	0.15	Binary variable indicating whether drug i receives designation as an orphan drug
$wideDistribution_i$	0.85	0.36	Binary variable indicating whether drug i is distributed to the entire nation
$approvalType_i$	0.27	0.44	Binary variable indicating whether drug i is a New Drug Application (NDA) drug

Note. Number of observations is 31,048.

Table 4   Classification Performances

Data set	Method	Accuracy	AUC	ADR			NADR		
				Precision	Recall	F_1 score	Precision	Recall	F_1 score
Forums	SVM	0.855	0.777	0.662	0.647	0.654	0.906	0.912	0.909
		(0.005)	(0.008)	(0.015)	(0.017)	(0.012)	(0.005)	(0.005)	(0.003)
TW	SVM	0.945	0.831	0.667	0.694	0.680	0.972	0.968	0.970
		(0.001)	(0.006)	(0.012)	(0.012)	(0.007)	(0.001)	(0.002)	(0.001)
Forums	BERT	0.870	0.786	0.711	0.642	**0.675**	0.907	0.930	**0.919**
		(0.009)	(0.016)	(0.047)	(0.039)	(0.025)	(0.011)	(0.016)	(0.007)
TW	BERT	0.957	0.833	0.768	0.685	**0.724**	0.972	0.981	**0.976**
		(0.003)	(0.012)	(0.021)	(0.025)	(0.015)	(0.003)	(0.002)	(0.001)

Note. Standard deviations in the parentheses. SVM stands for Support Vector Machine. BERT stands for Bidirectional Encoder Representation from Transformers

Table 5   Main Analyses

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
$adrPredictionsSVM_{i,t-1}$	0.0030***	0.0030***	0.0029***	0.0030***	0.0029***					
	(0.0009)	(0.0009)	(0.0009)	(0.0009)	(0.0009)					
$adrPredictionsBERT_{i,t-1}$						0.0054***	0.0053***	0.0052***	0.0053***	0.0052***
						(0.0013)	(0.0014)	(0.0014)	(0.0013)	(0.0014)
$FAERS_{i,t-1}$		0.0003***					0.0003***			
		(0.0001)					(0.0001)			
$professionals_{i,t-1}$			0.0010***		0.0010***			0.0008***		0.0008***
			(0.0003)		(0.0003)			(0.0003)		(0.0003)
$consumers_{i,t-1}$				0.0003**	0.0002				0.0003**	0.0002
				(0.0001)	(0.0001)				(0.0001)	(0.0001)
Baseline Hazard	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Time-varying Controls	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Firm Characteristics	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Drug Characteristics	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Random Effects	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
AIC	6006.44	6003.19	6002.90	6005.32	6003.38	6002.92	5999.94	5999.64	6001.96	6000.20
BIC	6172.90	6177.97	6177.68	6180.10	6186.49	6169.38	6174.72	6174.42	6176.74	6183.30
Observations	30419	30419	30419	30419	30419	30419	30419	30419	30419	30419

Notes. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Robust standard errors in parentheses. The baseline hazard function is specified as the logarithm of time, *lnt*.

Table 6   ADR Heterogeneity

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
$adrPredictionsSVM_{i,t-1}$	0.0030*** (0.0009)	0.0029*** (0.0009)			0.0043*** (0.0010)	0.0043*** (0.0010)			0.0043*** (0.0010)	0.0043*** (0.0012)		
$ratioSevereADRSVM_{i,t-1}$	0.2942 (0.5301)	0.2902 (0.5297)										
$adrPredictionsBERT_{i,t-1}$			0.0053*** (0.0014)	0.0052*** (0.0014)			0.0080*** (0.0016)	0.0080*** (0.0016)			0.0078*** (0.0017)	0.0078*** (0.0016)
$ratioSevereADRBERT_{i,t-1}$			−0.1529 (0.7746)	−0.1565 (0.7752)								
$FAERS_{i,t-1}$	0.0003*** (0.0001)		0.0003*** (0.0001)		0.0003*** (0.0001)		0.0003*** (0.0001)		0.0003* (0.0002)		0.0004** (0.0002)	
$professionals_{i,t-1}$		0.0008*** (0.0003)		0.0008*** (0.0003)		0.0009*** (0.0003)		0.0009*** (0.0003)		0.0009*** (0.0003)		0.0008*** (0.0003)
$consumers_{i,t-1}$		0.0002 (0.0001)		0.0002 (0.0001)		0.0002 (0.0001)		0.0002* (0.0001)		0.0002 (0.0003)		0.0003 (0.0002)
ADR Embedding			Y	Y	20	20	20	20	30	30	30	30
Baseline Hazard	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Time-varying Controls	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Firm Characteristics	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Drug Characteristics	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Random Effects	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
AIC	6004.96	6005.16	6001.89	6002.15	6017.74	6017.37	6008.38	6008.66	6026.57	6026.02	6014.97	6015.90
BIC	6188.06	6196.58	6185.00	6193.57	6358.98	6366.93	6341.30	6349.90	6451.04	6458.80	6422.79	6432.04
Observations	30419	30419	30419	30419	30419	30419	30419	30419	30419	30419	30419	30419

Notes. ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$. Robust standard errors in parentheses. The baseline hazard function is specified as the logarithm of time, *lnt*. Columns 1 - 4 report regression results after controlling for the severity of ADRs. Columns 5 - 8 report regression results after controlling for 20 dimensions of ADR characteristics. Columns 9 - 12 report regression results after controlling for 30 dimensions of ADR characteristics.

Table 7	Accelerated Failure Time Models			
	(1)	(2)	(3)	(4)
$adrPredictionsSVM_{i,t-1}$	-0.0023^{***}	-0.0023^{***}		
	(0.0007)	(0.0007)		
$adrPredictionsBERT_{i,t-1}$			-0.0039^{***}	-0.0039^{***}
			(0.0010)	(0.0010)
$FAERS_{i,t-1}$	-0.0002^{***}		-0.0002^{***}	
	(0.0001)		(0.0001)	
$professionals_{i,t-1}$		-0.0007^{***}		-0.0007^{***}
		(0.0002)		(0.0002)
$consumers_{i,t-1}$		-0.0002^{**}		-0.0002^{**}
		(0.0001)		(0.0001)
Time-varying Controls	Y	Y	Y	Y
Firm Characteristics	Y	Y	Y	Y
Drug Characteristics	Y	Y	Y	Y
Weibull Distribution	Y	Y	Y	Y
AIC	1532.16	1532.41	1529.39	1529.68
BIC	1698.61	1707.19	1695.84	1704.46
Observations	30419	30419	30419	30419

Notes. $^{***}p < 0.01$, $^{**}\ p < 0.05$, $^{*}\ p < 0.1$. Robust standard errors in parentheses.

Table 8   Alternative Specifications for the Discrete-time Survival Analysis

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
$adrPredictionsSVM_{i,t-1}$	0.0029*** (0.0009)	0.0029*** (0.0009)			0.0030*** (0.0008)	0.0029*** (0.0008)			0.0031*** (0.0010)	0.0031*** (0.0010)		
$adrPredictionsBERT_{i,t-1}$			0.0052*** (0.0014)	0.0052*** (0.0014)			0.0053*** (0.0013)	0.0052*** (0.0013)			0.0053*** (0.0014)	0.0052*** (0.0014)
$FAERS_{i,t-1}$	0.0003*** (0.0001)		0.0003*** (0.0001)		0.0003*** (0.0001)		0.0003*** (0.0001)		0.0003** (0.0001)		0.0003** (0.0001)	
$professionals_{i,t-1}$		0.0008*** (0.0003)		0.0008*** (0.0003)		0.0009*** (0.0003)		0.0009*** (0.0003)		0.0011* (0.0006)		0.0011** (0.0005)
$consumers_{i,t-1}$		0.0002** (0.0001)		0.0002** (0.0001)		0.0002* (0.0001)		0.0002* (0.0001)		0.0002 (0.0001)		0.0002 (0.0001)
Baseline Hazard	t^2	t^2	t^2	t^2	NP	NP	NP	NP	lnt	lnt	lnt	lnt
Link Function	$cloglog$	$cloglog$	$cloglog$	$cloglog$	$cloglog$	$cloglog$	$cloglog$	$cloglog$	$logit$	$logit$	$logit$	$logit$
Time-varying Controls	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Firm Characteristics	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Drug Characteristics	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Random Effects	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
AIC	6004.41	6004.78	6001.14	6001.57	5997.67	5997.72	5994.23	5994.36	6002.13	6001.96	5999.18	5999.10
BIC	6179.19	6187.88	6175.91	6184.67	6272.32	6280.69	6268.89	6277.33	6176.91	6185.06	6173.96	6182.20
Observations	30419	30419	30419	30419	30419	30419	30419	30419	30419	30419	30419	30419

Notes. $***p < 0.01$, $** p < 0.05$, $* p < 0.1$. Robust standard errors in parentheses. NP stands for non-parametric. Columns 1 - 4 report the results when the baseline hazard function is specified as the square of time, t^2. Columns 5 - 8 report the results when the baseline hazard function is specified as a set of time dummies. Columns 9 - 12 reports the results when the baseline hazard function is the logarithm of time. Columns 1 - 8 report the results where the link function is logit and the baseline hazard function is the logarithm of time.

Table 9   Falsification Tests on Different Classes of Drug Recalls

	Class I & II		Class III	
	(1)	(2)	(3)	(4)
$adrPredictionsSVM_{i,t-1}$	0.0030**		0.0022	
	(0.0012)		(0.0018)	
$adrPredictionsBERT_{i,t-1}$		0.0054***		0.0034
		(0.0016)		(0.0034)
$professionals_{i,t-1}$	0.0008**	0.0008**	0.0012**	0.0012**
	(0.0003)	(0.0003)	(0.0005)	(0.0005)
$consumers_{i,t-1}$	0.0001	0.0001	0.0003	0.0003
	(0.0001)	(0.0001)	(0.0002)	(0.0002)
Baseline Hazard	Y	Y	Y	Y
Time-varying Controls	Y	Y	Y	Y
Firm Characteristics	Y	Y	Y	Y
Drug Characteristics	Y	Y	Y	Y
Random Effects	Y	Y	Y	Y
AIC	3478.72	3475.73	2534.12	2534.11
BIC	3650.63	3647.64	2697.00	2696.99
Observations	18288	18288	12131	12131

Notes. ***$p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Robust standard errors in parentheses.
The baseline hazard function is specified as the logarithm of time, *lnt*.

Table 10   Mechanism Test on Information Effect

	(1)	(2)	(3)	(4)
$adrPredictionsSVMNew_{i,t-1}$	0.0060**	0.0059**		
	(0.0026)	(0.0026)		
$adrPredictionsSVMKnown_{i,t-1}$	0.0019	0.0019		
	(0.0023)	(0.0023)		
$adrPredictionsBERTNew_{i,t-1}$			0.0067***	0.0068***
			(0.0022)	(0.0022)
$adrPredictionsBERTKnown_{i,t-1}$			0.0043	0.0041
			(0.0046)	(0.0045)
$FAERS_{i,t-1}$	0.0003***		0.0003***	
	(0.0001)		(0.0001)	
$professionals_{i,t-1}$		0.0008**		0.0008**
		(0.0003)		(0.0003)
$consumers_{i,t-1}$		0.0003**		0.0003**
		(0.0001)		(0.0001)
Baseline Hazard	Y	Y	Y	Y
Time-varying Controls	Y	Y	Y	Y
Firm Characteristics	Y	Y	Y	Y
Drug Characteristics	Y	Y	Y	Y
Random Effects	Y	Y	Y	Y
AIC	3695.91	3697.27	3692.86	3694.20
BIC	3861.06	3870.28	3858.00	3867.21
Observations	19226	19226	19226	19226

Notes. ***$p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Robust standard errors in parentheses. The
baseline hazard function is specified as the logarithm of time, *lnt*.

Table 11   Mechanism Tests on Publicity Effect

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$adrPredictionsSVMTwitter_{i,t-1}$	0.0032***	0.0025***	0.0029***	0.0031***				
	(0.0009)	(0.0009)	(0.0009)	(0.0008)				
$adrPredictionsBERTTwitter_{i,t-1}$					0.0063***	0.0049***	0.0060***	0.0060***
					(0.0012)	(0.0017)	(0.0013)	(0.0013)
$avgLogFollowers_{i,t-1}$			0.0547***			0.0493***		
			(0.0175)			(0.0166)		
$avgLogLikes_{i,t-1}$				0.9899***			0.8927***	
				(0.1810)			(0.1569)	
$avgLogRetweets_{i,t-1}$					0.7834***			0.8927***
					(0.2046)			(0.1686)
$professionals_{i,t-1}$	0.008***	0.0007**	0.0008***	0.0008***	0.0008***	0.0007**	0.0008***	0.0008***
	(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0003)
$consumers_{i,t-1}$	0.0002*	0.0002*	0.0002*	0.0002*	0.0002*	0.0002*	0.0002*	0.0002*
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Baseline Hazard	Y	Y	Y	Y	Y	Y	Y	Y
Time-varying Controls	Y	Y	Y	Y	Y	Y	Y	Y
Firm Characteristics	Y	Y	Y	Y	Y	Y	Y	Y
Drug Characteristics	Y	Y	Y	Y	Y	Y	Y	Y
Random Effects	Y	Y	Y	Y	Y	Y	Y	Y
AIC	6003.09	5993.17	5985.99	5995.13	5998.04	5988.53	5972.26	5978.63
BIC	6186.20	6176.27	6169.10	6178.23	6181.14	6171.63	6155.36	6161.73
Observations	30419	30419	30419	30419	30419	30419	30419	30419

Notes. ***$p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Robust standard errors in parentheses. The baseline hazard function is specified as the logarithm of time, lnt.

52

Authors' names blinded for peer review
Article submitted to *Information Systems Research*; manuscript no. (Please, provide the manuscript number!)

Table 12   Interaction between ADR Discussions and Distribution

	(1)	(2)	(3)	(4)
$adrPredictionsSVM_{i,t-1}$	0.0130***	0.0129***		
	(0.0040)	(0.0040)		
$adrPredictionsSVM_{i,t-1} \times wideDistribution_i$	−0.0103**	−0.0103**		
	(0.0040)	(0.0040)		
$adrPredictionsBERT_{i,t-1}$			0.0147***	0.0146***
			(0.0055)	(0.0055)
$adrPredictionsBERT_{i,t-1} \times wideDistribution_i$			−0.0099*	−0.0098*
			(0.0055)	(0.0055)
$wideDistribution_i$	−0.0626	−0.0650	−0.0612	−0.0636
	(0.2480)	(0.2412)	(0.2628)	(0.2675)
$FAERS_{i,t-1}$	0.0003**		0.0003**	
	(0.0001)		(0.0001)	
$professionals_{i,t-1}$		0.0008***		0.0008***
		(0.0003)		(0.0003)
$consumers_{i,t-1}$		0.0002		0.0002
		(0.0002)		(0.0002)
Baseline Hazard	Y	Y	Y	Y
Time-varying Controls	Y	Y	Y	Y
Firm Characteristics	Y	Y	Y	Y
Drug Characteristics	Y	Y	Y	Y
Random Effects	Y	Y	Y	Y
AIC	6000.54	6000.82	5999.15	5999.45
BIC	6183.65	6192.25	6182.25	6190.87
Observations	30419	30419	30419	30419

Notes. ***$p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Robust standard errors in parentheses. The baseline hazard function is specified as the logarithm of time, *lnt*.

Appendix A: Correlation Matrix

Table A.1.   Correlation Matrix

Variable	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21
$adrPredictionsSVM_{i,t}$	-																				
$adrPredictionsBERT_{i,t}$	0.90*	-																			
$FAERS_{i,t}$	0.09*	0.12*	-																		
$professionals_{i,t}$	0.14*	0.17*	0.59*	-																	
$consumers_{i,t}$	0.05*	0.07*	0.95*	0.33*	-																
$substitute_{i,t}$	0.02*	0.05*	0.01	0.02*	0.00	-															
$clinicalTrials_{i,t}$	0.02*	0.03*	0.06*	0.09*	0.03*	-0.03*	-														
$news_{i,t}$	0.29*	0.31*	0.08*	0.09*	0.05*	0.02*	-0.02*	-													
$marketSize_{c,t}$	0.05*	0.04*	0.01	0.01	0.01	0.16*	0.07*	0.03*	-												
$numRecalls_{j,t}$	0.01*	0.00	0.02*	0.03*	0.02*	0.11*	0.11*	-0.05*	0.04*	-											
$numDrugs_{j,t}$	0.04*	0.04*	-0.01	-0.01	-0.00	0.12*	0.03*	0.01	0.07*	0.50*	-										
$competition_{j,t}$	0.04*	0.08*	0.04*	0.06*	0.03*	0.43*	0.02*	0.01	0.20*	0.13*	0.21*	-									
$complexity_{j,t}$	-0.01	-0.02*	0.02*	0.03*	0.02*	0.00	0.06*	-0.02*	-0.01	0.09*	0.11*	-0.05*	-								
$abuse_i$	0.06*	0.03*	-0.02*	-0.03*	-0.01*	0.02*	0.02*	0.05*	0.20*	0.04*	0.10*	-0.00	-0.03*	-							
$adverseReactions_i$	0.00	0.01	0.02*	0.03*	0.01*	-0.07*	0.22*	0.03*	0.01	0.05*	0.06*	-0.04*	0.09*	-0.02*	-						
$warnings_i$	0.02*	0.03*	0.00	0.01	-0.00	-0.07*	0.21*	0.03*	0.01	0.06*	0.07*	-0.06*	0.08*	-0.02*	0.91*	-					
$boxedWarnings_i$	0.04*	0.04*	0.10*	0.14*	0.06*	-0.05*	0.11*	0.01	0.06*	-0.01	0.00	0.03*	0.06*	0.19*	0.34*	0.34*	-				
$priorityReview_i$	0.02*	-0.02*	0.02*	0.04*	0.02*	-0.42*	0.04*	0.01*	-0.12*	-0.11*	-0.21*	-0.53*	0.04*	0.02*	0.09*	0.08*	0.09*	-			
$orphanDrug_i$	-0.03*	-0.04*	0.01	0.01	0.00	-0.16*	0.03*	-0.02*	-0.05*	-0.08*	-0.12*	-0.26*	0.04*	-0.05*	0.01*	0.02*	0.02*	0.47*	-		
$wideDistribution_i$	0.03*	0.02*	-0.00	-0.00	0.00	0.03*	0.02*	0.00	0.03*	0.10*	-0.05*	0.13*	0.01	0.05*	-0.00	-0.03*	-0.05*	-0.10*	-0.13*	-	
NDA_i	-0.03*	-0.06*	-0.07*	-0.09*	-0.05*	-0.40*	-0.01	-0.03*	-0.17*	-0.01*	-0.05*	-0.69*	0.11*	-0.04*	0.14*	0.15*	-0.04*	0.46*	0.14*	0.00	-

Note. * $p < 0.05$

Appendix B: Text Classification for ADR Detection

B.1.   Support Vector Machine (SVM)

We follow the procedure proposed by Sarker and Gonzalez (2015) to build an SVM classifier for ADR detection. First, we preprocess social media posts by lowercasing, stemming, and parsing all the tokens. Next, we extract the following features from the preprocessed tokens.

- *Unified Medical Language System (UMLS) semantic types and concept IDs (CUIs).* We use MetaMap toolbox[32] to identify terms that are related to the UMLS semantic types and the CUIs, which represent categories of medical concepts. To reflect the importance of a term, we compute the term frequency-inverse document frequency (TF-IDF) values for this set of features.

- *Syn-set expansion.* For each adjective, noun, or verb in a sentence, we use WordNet[33] to identify the synonyms of that term and add the synonymous terms as features. Similar to the previous feature set, we calculate the TF-IDF value for each derived synonym.

- *Change phrases.* To capture the polarity of the text, we use the lexicon proposed by Sarker et al. (2013) to derive the following four binary features: MORE-GOOD, MORE-BAD, LESS-GOOD, and LESS-BAD. For example, if a MORE-word and a GOOD-word both occur in a text, we denote the feature MORE-GOOD as one. The other three features are constructed in a similar manner.

[32] See http://metamap.nlm.nih.gov/

[33] See http://wordnet.princeton.edu/

- *ADR lexicon matches.* To incorporate domain-specific knowledge, we use the lexicon proposed by Leaman and Wojtulewicz (2010) to derive two features measuring ADR mentions in the text. The first feature is a binary feature indicating the presence of ADR mentions and the second feature is a numeric feature indicating the ratio of ADR mentions to the total number of words in the text segment.

- *Sentiword scores.* We use the lexicon proposed by Guerini et al. (2013) to assign each term with a sentiment score between -1 to 1 and then calculate the overall sentiment of a text by summing the individual scores divided by the length of the sentence in words.

- *Topic-based feature.* We use a topic modeling technique called Mallet[34] to generate two features: the topic terms and the sums of all the relevance scores of the terms.

- *Other features.* We include three simple features, including the length of the text segments in words, the presence of comparatives and superlatives, and the presence of modals.

After creating such a comprehensive set of features, we split a data set into a training set (80%) and a test set (20%) and then conduct a grid search with 10-fold cross-validation over the training set to find the best hyperparameters, such as the kernel function, the regularization parameter, and the weight assigned to each class. Next, we build an SVM classifier supervised by the training set and evaluate it on the test set. We report a standard set of performance measures for evaluation, such as accuracy, area under the ROC curve (AUC), precision, recall, and F_1 score.

B.2. Bidirectional Encoder Representations from Transformers (BERT)

BERT is a language representation model that pre-trains deep bidirectional representations from the unlabeled text by jointly conditioning on both left and right context in all layers (Devlin et al. 2018). Due to the ability to learn contextual relations between words in a text, the BERT model has achieved state-of-the-art performance on multiple NLP tasks, such as the General Language Understanding Evaluation (GLEU) task sets, Stanford Question Answering Dataset (SQuAD), and Situations With Adversarial Generations (SWAG).

There are two steps to implement a BERT model: *pre-training* and *fine-tuning.* In the pre-training stage, the BERT model is trained on unlabeled data over two tasks: masked language model, the objective of which is to predict the original token of a masked token based only on its context; and next sentence prediction,

---

[34] See http://mallet.cs.umass.edu/index.php

the objective of which is to understand the relationship between two sentences. Devlin et al. (2018) provides two pre-trained models in their seminal work: $BERT_{BASE}$ (12 layers, 768 hidden vector size, 12 heads, and 110 million parameters) and $BERT_{LARGE}$ (24 layers, 1024 hidden vector size, 16 heads, and 340 million parameters). In the fine-tuning stage, the BERT model is first initialized with the pre-trained parameters, and then all parameters are fine-tuned using labeled data from the specific NLP task.

In this study, we consider the $BERT_{BASE}$ pre-trained by Devlin et al. (2018) as a starting point and utilize transfer learning to fine-tune the neural network with the annotated social media posts. More specifically, we split the data set into a training set (80%) and a test set (20%) and fine-tune the BERT model with the training set. Next, we evaluate the fine-tuned BERT model on the test set and report a standard set of performance measures.

Appendix C: Summary Statistics

Table C.1.   **Summary Statistics on Different Classes of Drug Recalls**

Variable	No. of obs.	Mean	S. D.	Min.	Max.
Class I & II Recall					
$adrPredictionsSVM_{i,t}$	18,652	5.72	24.27	0	879
$adrPredictionsBERT_{i,t}$	18,652	5.07	18.36	0	562
$professionals_{i,t}$	18,652	21.42	49.00	0	3275
$consumers_{i,t}$	18,652	20.62	109.29	0	10525
Class III Recall					
$adrPredictionsSVM_{i,t}$	12,396	7.33	32.19	0	572
$adrPredictionsBERT_{i,t}$	12,396	5.54	21.02	0	392
$professionals_{i,t}$	12,396	15.04	33.64	0	1985
$consumers_{i,t}$	12,396	17.30	144.06	0	6762

Table C.2.   **Summary Statistics on ADR-related Tweets**

Variable	Mean	S. D.	Definition
adrPredictionsSVM			
$adrPredictionsSVMTwitter_{i,t}$	4.63	24.49	Number of tweets that mention drug i's proprietary name and are predicted by SVM to be ADR-related at month t
$avgLogFollowers_{i,t}$	1.14	2.25	Average of logarithm-transformed number of followers for drug i at month t
$avgLogLikes_{i,t}$	0.01	0.10	Average of logarithm-transformed number of likes for drug i at month t
$avgLogRetweets_{i,t}$	0.01	0.09	Average of logarithm-transformed number of retweets for drug i at month t
adrPredictionsBERT			
$adrPredictionsBERTTwitter_{i,t}$	3.38	15.53	Number of tweets that mention drug i's proprietary name and are predicted by BERT to be ADR-related at month t
$avgLogFollowers_{i,t}$	1.30	2.52	Average of logarithm-transformed number of followers for drug i at month t
$avgLogLikes_{i,t}$	0.01	0.10	Average of logarithm-transformed number of likes for drug i at month t
$avgLogRetweets_{i,t}$	0.01	0.10	Average of logarithm-transformed number of retweets for drug i at month t

Note. Number of observations is 31,048.