Article

Journal of Marketing
2021, Vol. 85(2) 70-88

DOI: 10.1177/0022242920937703
journals.sagepub.com/home/jmx

Do Spoilers Really Spoil? Using Topic Modeling to Measure the Effect of Spoiler Reviews on Box Office Revenue

Jun Hyun (Joseph) Ryoo, Xin (Shane) Wang, and Shijie Lu

Abstract

A sizable portion of online movie reviews contain spoilers, defined as information that prematurely resolves plot uncertainty. In this research, the authors study the consequences of spoiler reviews using data on box office revenue and online word of mouth for movies released in the United States. To capture the degree of information in spoiler review text that reduces plot uncertainty, the authors propose a spoiler intensity metric and measure it using a correlated topic model. Using a dynamic panel model with movie fixed effects and instrumental variables, the authors find a significant and positive relationship between spoiler intensity and box office revenue with an elasticity of .06. The positive effect of spoiler intensity is greater for movies with a limited release, smaller advertising spending, and moderate user ratings, and is stronger in the earlier days after the movie's release. Using an event study and online experiments, the authors provide further evidence that spoiler reviews can help consumers reduce their uncertainty about the quality of movies, consequently encouraging theater visits. Thus, movie studios may benefit from consumers' access to plot-intense reviews and should actively monitor the content of spoiler reviews to better forecast box office performance.

Keywords

machine learning, motion pictures, online word of mouth, spoilers, topic modeling

Online supplement: https://doi.org/10.1177/0022242920937703

In April 2019, the directors of *Avengers: Endgame* issued a stern warning to fans about the much anticipated blockbuster film: "When you see *Endgame* in the coming weeks, please don't spoil it for others, the same way you wouldn't want it spoiled for you" (Kooser 2019). As a marketing tactic, this ploy was successful, generating significant buzz on social media. However, the directors' true intention behind their statement remains ambiguous. Did they truly want to silence viewers? What is the relationship between spoilers and box office revenue? Should movie studios be concerned about the exchange of spoilers among consumers? Extant marketing research is unequivocal that online word of mouth (WOM) is vital for the financial success of new products such as movies (e.g., Babić Rosario et al. 2016; Kerrigan 2017). However, the understanding of spoilers and how they influence consumer purchase decisions is still limited.

In the context of movies, a "spoiler review" refers to a movie review that contains spoilers, and a "nonspoiler review" refers to a movie review without any spoilers, where a "spoiler" is defined as information that prematurely resolves plot uncertainty for those who have yet to see the movie. According to Internet Movie Database (IMDb) data, approximately 93% of movies released between January 2013 and December 2017 in the United States garnered at least one spoiler review throughout their screenings, and approximately 31% of total movie reviews contained spoilers, suggesting the prevalence of spoiler reviews in the movie industry. With the growth of social media, spoiler reviews can spread rapidly throughout the internet to reach a broad audience. Conventional wisdom suggests a negative relationship between spoiler reviews and consumer demand, as exemplified by the concern raised by the directors of *Avengers: Endgame*. However, previous research has shown either mixed or null effects of spoilers on consumer behavior (Johnson and Rosenbaum 2015; Leavitt and Christenfeld 2011). Thus, the prevalence of spoilers in the movie industry and its unclear ramifications call for a deeper understanding of whether and

Jun Hyun (Joseph) Ryoo is a PhD Candidate, Department of Marketing, Western University, Canada (email: jryoo.phd@ivey.ca). Xin (Shane) Wang is Associate Professor of Marketing and MBA '80 Faculty Fellow, Western University, Canada (email: xwang@ivey.ca). Shijie Lu is Assistant Professor of Marketing, University of Houston, USA (email: slu@bauer.uh.edu).

how spoiler reviews affect consumers' moviegoing decisions—questions we attempt to address in this research.

We provide a conceptual discussion of spoilers that guides the development of *spoiler intensity*, which we define as the degree of information in spoiler reviews that reduces plot uncertainty. Although previous marketing research has examined the relationship between consumer demand and various aspects of online WOM such as volume (Godes and Mayzlin 2004; Liu 2006), valence (Chevalier and Mayzlin 2006; Chintagunta, Gopinath, and Venkataraman 2010; Moon, Bergey, and Iacobucci 2010), and variance (Sun 2012), most studies have not considered the information within review content beyond the sentiment. Unlike spoiler volume, spoiler intensity is a latent construct that needs to be inferred from review text. In this study, we use a correlated topic model (CTM; Blei and Lafferty 2005) to identify key topics in movie reviews and propose a spoiler intensity metric as a function of these topics.

We assembled a data set of 140,869 reviews for 993 movies released in the United States between January 2013 and December 2017. We collected both spoiler and nonspoiler reviews from IMDb and exploited the review platform's spoiler labels for movie reviews as a training sample to identify topics that are more likely to appear in spoiler than nonspoiler reviews, which we then used in the construction of the spoiler intensity metric. Using a dynamic panel model with movie fixed effects, we quantified the association between spoiler reviews and box office revenue. We alleviated the potential endogeneity concern arising from the inclusion of WOM-related variables and marketing mix variables using instrumental variables. We find that the spoiler intensity of a movie's reviews is positively associated with subsequent box office revenue, whereas the association between spoiler volume and subsequent box office revenue is not evident. We also provide evidence that these findings are robust to alternative specifications of spoiler intensity.

We further investigate the behavioral mechanism that may drive the positive relationship between spoiler intensity and demand. Moviegoers often visit online review platforms to seek diagnostic information from their peers and resolve uncertainty about movie quality (Dellarocas 2003; Goh, Heng, and Lin 2013). Unlike nonspoiler reviews, spoiler reviews can reveal plot-related information as justifications when critiquing a movie and therefore tend to be more diagnostic for potential moviegoers. As such, we expected that the diagnostic value of spoiler reviews would help consumers reduce uncertainty about movie quality, which in turn would encourage theater visits. To indirectly test the uncertainty-reduction mechanism of spoiler reviews, we considered four potential moderators of the effect of spoiler intensity: (1) release type (limited release vs. wide release), (2) movie age, (3) advertising, and (4) average user rating. We find that the positive effect of spoiler intensity is larger for movies characterized by greater uncertainty for moviegoers, such as limited release movies and movies with smaller advertising spending. In addition, the effect of spoiler intensity decays over time, which is consistent with the higher uncertainty in the earlier (rather than later) stages of a movie's life cycle. We also find an inverted U relationship between average user ratings and the effect of spoiler intensity, which suggests that the positive spoiling effect is stronger for movies that receive moderate or mixed ratings compared to movies that receive extreme ratings (i.e., either very high or low). This finding is likely driven by the fact that user ratings in the middle range tend to convey more ambiguous signals about movie quality than extreme ratings (Tang, Fang, and Wang 2014). Thus, potential consumers of movies with moderate user ratings have greater incentive to seek diagnostic information to reduce their uncertainty about future consumption.

Moreover, we present additional evidence in support of the uncertainty-reduction mechanism of spoiler reviews from an event study. In particular, we examine the change in the effect of spoiler intensity on box office revenue after an exogenous update on the IMDb website that increased both consumers' cost for reading spoiler reviews and the diagnosticity of nonspoiler reviews. If the uncertainty-reduction mechanism is indeed important, we would expect the positive effect of spoiler intensity on demand to be weakened after the website update because of the decrease in the relative diagnostic value and the increase in the cost of reading spoiler reviews. Our results from the event study are consistent with this expectation and therefore provide additional support for the proposed mechanism.

To complement the findings from the field study, we ran online experiments for two different movies to test causal links between spoiler intensity and moviegoing decisions. Specifically, we employed a 2 (high spoiler intensity vs. low spoiler intensity) $\times$ 2 (high uncertainty vs. low uncertainty) between-subjects design, in which we asked each subject to read a spoiler review and a nonspoiler review before expressing their willingness to watch the movie. We manipulated movie uncertainty by showing subjects video clips with different levels of plot-related information. We then presented artificially created spoiler reviews that were the same except for one sentence to manipulate spoiler intensity between the two conditions. The results show that for subjects in the high-uncertainty condition, reading a more "spoiled" review increases the willingness to watch the movie. However, for those in the low-uncertainty condition, the effect is statistically nonsignificant. These findings are consistent with the proposed uncertainty-reduction mechanism of spoiler reviews. When consumers have low uncertainty about movie quality, they do not benefit much from the additional reduction in uncertainty from spoiler reviews, which explains the null effect we find in the low-uncertainty condition.

With this research, we aim to make three contributions. First, we provide a conceptual background of spoilers by formally defining what constitutes spoiling information in a movie review and discussing several key properties that a spoiler intensity metric needs to capture. Second, we make substantive contributions by showing a positive association between spoiler reviews and consumer demand driven by spoiler intensity rather than spoiler volume. Furthermore, we show that the effect of spoiler intensity is more prominent for movies with a limited release, smaller advertising spending, and moderate user ratings. The positive effect of spoiler intensity is also stronger in earlier periods of a movie's life cycle. Finally, using both field and experimental data, we find data patterns that support the behavioral mechanism that uncertainty reduction drives the positive effect of spoiler intensity.

Related Literature

Given our focus on spoiler reviews, this research builds on the literature on online WOM. Extant marketing research conceptualizes the influence of online WOM on demand through two distinct channels (e.g., Babić Rosario et al. 2016; Seiler, Yao, and Wang 2017): (1) the informative effect of online WOM involves increasing consumer awareness about the existence of a product and providing information about the product that consumers seek and value, and (2) the persuasive effect of online WOM involves increasing consumers' appreciation for a product without delivering specific product information. The informative role of online WOM is supported by the positive relationship found between the number of reviews and box office sales (Duan, Gu, and Whinston 2008; Liu 2006) and between the amount of online conversation and television ratings (Godes and Mayzlin 2004). The persuasive effect of online WOM is supported by the positive relationship found between valence (e.g., review ratings, sentiment) and demand (Chevalier and Mayzlin 2006; Chintagunta, Gopinath, and Venkataraman 2010; Dellarocas, Zhang, and Awad 2007). Regarding the variance of online WOM, measured by the statistical dispersion of ratings, previous findings are less consistent. This is in part because of the complex ways in which variance may affect sales (Clemons, Gao, and Hitt 2006; Sun 2012).

In addition to the summary statistics of online WOM (e.g., ratings and volume), marketing scholars have explored specific types and patterns of online WOM observed in the movie industry. For example, Hennig-Thurau, Wiertz, and Feldhaus (2015) examine Twitter to study the diagnostic value of microblogging WOM and find that negative tweets are potentially harmful to a movie's early box office revenue. Gelper, Peres, and Eliashberg (2018) note that sporadic volume bursts, or spikes of online WOM prior to a movie's release, are positively associated with opening weekend box office revenue. Recently, academics have paid increasing attention to online WOM content beyond its overall valence. Gopinath, Thomas, and Krishnamurthi (2014) use human coders to examine the attribute-, emotion-, and recommendation-oriented dimensions of online WOM and find that only the valence of the recommendation-oriented dimensions has an impact on sales. Liu, Singh, and Srinivasan (2016) use the principal components of words in tweets to show that the content of online WOM can significantly increase the accuracy of predictions about television show ratings.

It is particularly important to account for the WOM content when examining the impact of online WOM in the entertainment industry for at least two reasons. First, summary statistics alone cannot provide a full picture. For instance, previous research has shown that review ratings are subject to inflation (Chevalier and Mayzlin 2006) and selection bias (Dellarocas 2003; Godes and Silva 2012; Li and Hitt 2008), suggesting that ratings can sometimes be misleading in signaling a movie's true quality. Second, to minimize the risk of watching movies of poor quality, potential moviegoers have incentives to read detailed content (Mudambi and Schuff 2010), especially content related to plots, to seek diagnostic information. We contribute to the online WOM literature by presenting the first empirical study of the relationship between plot-related WOM, which often appears in spoiler reviews, and consumer demand in the movie industry.

Conceptual Discussion

What Are Spoilers?

Previous research in the field of literature finds that consumption of stories involves a prospective orientation in the minds of consumers, related to forming predictions and looking ahead to what will happen next in the plot (Olson, Mack, and Duffy 1981). As a result, plot uncertainty, which stimulates tension and suspense, serves as an important source of utility in story consumption (Ely, Frankel, and Kamenica 2015). For example, consumers often become emotionally invested in the protagonist, who might encounter danger in a story, and the protagonist's uncertain fate creates suspense that causes consumers to yearn for the story's resolution (Zillmann 1995). In the context of movies, plot uncertainty can be resolved either by watching the movie or by reading reviews that include plot-related information prior to watching it. We therefore define spoilers as information that prematurely resolves plot uncertainty for those who have yet to see the movie.

Effect of Spoilers

Extant research in psychology and communication has revealed mixed findings regarding the impact of spoilers on story enjoyment. By manipulating the types of short stories read by subjects in laboratory conditions, Leavitt and Christenfeld (2011) find that spoilers can have a positive effect on media enjoyment. The authors later explain that this effect is caused by consumers' increased ease of understanding the media experience due to spoilers, which frees cognitive resources and allows them to enjoy media at a deeper level (Leavitt and Christenfeld 2013). In contrast, Johnson and Rosenbaum (2015) find that spoiled stories are less fun and suspenseful when using a multidimensional approach to measure enjoyment. They explain their findings using excitation transfer theory (Zillmann, Hay, and Bryant 1975), positing that spoilers have a negative effect on media enjoyment because they displace the physiological arousal generated by suspense that should be resolved by media consumption.

The relationship between spoilers and consumer demand is arguably more relevant to marketers. In contrast to the conventional knowledge that spoilers harm demand, Johnson and Rosenbaum (2015) fail to find a significant effect of spoilers on media selection. When subjects were presented with a choice between spoiled and unspoiled short stories, they were just as likely to choose the spoiled stories as the unspoiled stories.

However, researchers have not yet examined the relationship between spoiler reviews and movie demand.

On the one hand, spoiler reviews might discourage theater visits. By prematurely revealing plot-related information, spoiler reviews can ruin the element of surprise in a movie experience and consequently decrease consumption utility. Such a "surprise burst" effect can be triggered by different types of plot-related information of movies from different genres. For example, the death of a character could be a surprising event for a dramatic movie, while the proposal and marriage between characters could be the ultimate surprise for a romantic movie.

On the other hand, spoiler reviews might help consumers reduce uncertainty about product fit. Due to their subjective nature, the quality of experiential products such as movies is difficult for consumers to evaluate prior to consumption (Alba and Williams 2013). By revealing important plot details and increasing the informative value of WOM, spoiler reviews could have a positive effect on movie demand. It is unclear whether this positive uncertainty-reduction effect outweighs the negative surprise burst effect of spoiler reviews in the movie industry. Using both field and experimental data, we seek to extend the literature on spoilers by investigating the net effect of spoiler reviews on movie demand (i.e., the sum of the positive effect from uncertainty reduction and the negative effect from the burst of surprise).

Definition and Properties of Spoiler Intensity

Studying the consequences of spoiler reviews requires measuring both the volume and intensity of spoilers, where spoiler intensity is defined as the degree to which the information in spoiler reviews reduces plot uncertainty. Consider a movie that receives multiple spoiler reviews. Measuring only the number of spoiler reviews is inadequate at capturing the spoiling effect because these reviews may provide similar plot-related information and therefore do not accumulate in resolving plot uncertainty. As such, spoiler intensity is an important construct that differs from spoiler volume. Next, we present and explain several key properties that an adequate measure of spoiler intensity should capture.

Property 1: Spoiler intensity should be a continuous rather than dichotomous variable because the extent to which a spoiler review resolves plot uncertainty depends on the level of detail in the review. For example, a spoiler review for the movie *Avengers: Endgame* can reveal not only the names of characters who died at the end (e.g., "Iron Man dies") but also the causes and consequences of the deaths (e.g., "Iron Man sacrifices himself to defeat Thanos"), which further resolve plot uncertainty for consumers. Therefore, a dichotomous variable is insufficient to capture the level of plot uncertainty prematurely resolved in a spoiler review.

Property 2: Spoiler intensity should capture a multitude of plot-related topics that are involved in the structure of a story. Previous research suggests that stories in general share similar patterns and plot structures, and stories in movies are no exception (Deighton, Romer, and McQueen 1989). In particular, movie plots typically unfold in a three-act structure: exposition, rising action, and climax (Trottier 1998), where exposition is used to introduce the major characters, rising action occurs when the protagonist encounters some sort of crisis that creates tension, and the climax features the resolution of the main tensions of the story. For each act, the screenwriter can craft the story using various elements, which we call plot-related topics (e.g., topics related to "fight" often appear in the climax of action movies, whereas topics related to "emotion" often appear in the climax of romantic movies). Because of the similar patterns and structures of stories in the movie industry, we assume that movie reviews convey a discrete number of plot-related topics.

Property 3: Spoiler intensity should allow the degree to which a spoiler review resolves uncertainty regarding a specific topic to vary across movies. For example, although both *Avengers: Endgame* and *The Lego Movie* might include the topic of "survival," the level of suspense resolved by reading plot details related to the topic of "survival" in a spoiler review is likely to be greater for *Avengers: Endgame* than for *The Lego Movie* due to the overall storyline and plot structure. Thus, an adequate measure of spoiler intensity should account for the potential heterogeneity in each topic's contribution to resolving plot uncertainty across movies.

Property 4: Spoiler intensity should discount the degree to which a spoiler review resolves plot uncertainty regarding a certain topic when the topic has appeared in previous reviews. This property captures the potential dynamics in the spoiling process when a consumer reads multiple reviews. For instance, suppose a consumer has already read several spoiler reviews. Three scenarios might occur when this consumer reads a subsequent spoiler review. First, the new spoiler review could include information about new plot-related topics that have not appeared in previous reviews. Given that a new facet of plot uncertainty can be resolved by reading this new spoiler review, the degree to which this additional spoiler review resolves plot uncertainty should not be discounted when assessing the overall spoiler intensity of multiple reviews. Second, the new spoiler review could include information on plot-related topics that have already appeared in previous reviews but provide additional details for these existing topics. In this case, the new spoiler review would further resolve plot uncertainty pertaining to existing topics because of the additional information it provides. Third, the new spoiler review could include information on plot-related topics that has appeared in previous reviews but not provide any new information for these topics. This new spoiler review's contribution to reducing a consumer's plot uncertainty needs to be discounted because previous reviews would still be driving the consumer's feeling of suspense.

Setting and Data

We obtained a list of movies released in the United States between January 2013 and December 2017 from WildAboutMovies.com. From this list, we sampled 993 movies that have their daily box office revenue data available on BoxOfficeMojo.com.

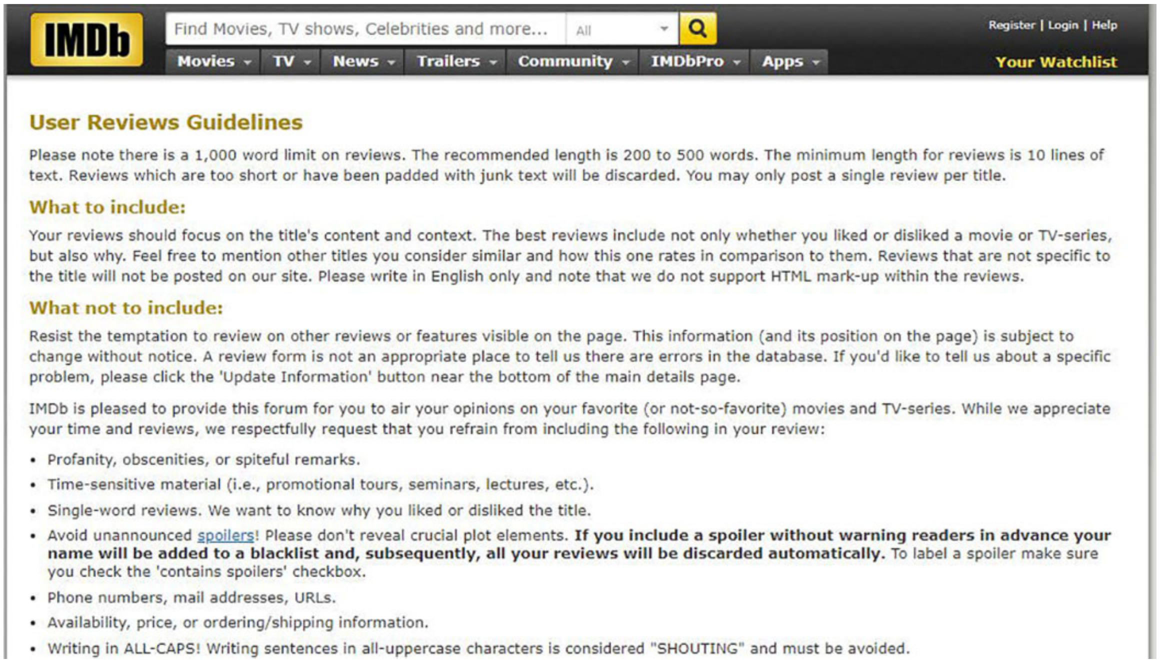

IMDb

Find Movies, TV shows, Celebrities and more... All

Register | Login | Help

Movies TV News Trailers Community IMDbPro Apps

Your Watchlist

User Reviews Guidelines

Please note there is a 1,000 word limit on reviews. The recommended length is 200 to 500 words. The minimum length for reviews is 10 lines of text. Reviews which are too short or have been padded with junk text will be discarded. You may only post a single review per title.

What to include:

Your reviews should focus on the title's content and context. The best reviews include not only whether you liked or disliked a movie or TV-series, but also why. Feel free to mention other titles you consider similar and how this one rates in comparison to them. Reviews that are not specific to the title will not be posted on our site. Please write in English only and note that we do not support HTML mark-up within the reviews.

What not to include:

Resist the temptation to review on other reviews or features visible on the page. This information (and its position on the page) is subject to change without notice. A review form is not an appropriate place to tell us there are errors in the database. If you'd like to tell us about a specific problem, please click the 'Update Information' button near the bottom of the main details page.

IMDb is pleased to provide this forum for you to air your opinions on your favorite (or not-so-favorite) movies and TV-series. While we appreciate your time and reviews, we respectfully request that you refrain from including the following in your review:

- Profanity, obscenities, or spiteful remarks.
- Time-sensitive material (i.e., promotional tours, seminars, lectures, etc.).
- Single-word reviews. We want to know why you liked or disliked the title.
- Avoid unannounced spoilers! Please don't reveal crucial plot elements. **If you include a spoiler without warning readers in advance your name will be added to a blacklist and, subsequently, all your reviews will be discarded automatically.** To label a spoiler make sure you check the 'contains spoilers' checkbox.
- Phone numbers, mail addresses, URLs.
- Availability, price, or ordering/shipping information.
- Writing in ALL-CAPS! Writing sentences in all-uppercase characters is considered "SHOUTING" and must be avoided.

Figure 1. User review guidelines on IMDb.

Table 1. Variable Definitions.

Variable Name	Description
DAILYREV	Box office revenue on day t for movie i.
INTENSITY	Spoiler intensity of spoiler reviews within the ten days prior to day t for movie i.
PROP	Moving average of proportion of spoiler reviews within the ten days prior to day t for movie i.
CUMRATING	Mean ratings of cumulative movie reviews on day t for movie i.
CUMVOL	Number of cumulative movie reviews on day t for movie i.
ADVERT	Average daily advertising expenditure on day t for movie i.
THEATERS	Number of theaters that screen movie i on day t.
AGE (t)	Number of days since the release of movie i in theaters.
HOLIDAY	Dummy variable for the ten federal holidays in the United States.
DAYOFWEEK	Indicator variables for each day of the week.

We focused on the first eight weeks of daily box office revenue because 97% of total box office revenue is accrued within the first eight weeks of a movie's release (Liu 2006). We collected daily box office revenue and daily number of theaters in which a movie was playing, as well as other movie characteristics (e.g., Motion Picture Association of America rating, genre, and release type) from both BoxOfficeMojo.com and IMDb. We matched our movie sample with advertising spending data provided by Kantar Media.

We used IMDb to collect online WOM data for two reasons. First, IMDb is by far the most popular online movie review platform in the United States.[1] Second, IMDb requires users to label their reviews with spoiler warnings if a user believes that their review discloses any critical plot elements of a movie. As Figure 1 shows, IMDb penalizes users who do not label spoiler reviews by blacklisting their accounts and deleting their reviews automatically. This institutional feature gives us a data set with a clear classification between spoiler and nonspoiler reviews.

Table 1 lists key time-varying variables in this study, along with their descriptions. Table 2 presents summary statistics of time-varying variables and time-invariant movie characteristics. On average, each movie's daily box office revenue was \$1.04 million. Each movie received approximately one spoiler review and two nonspoiler reviews per day.[2] As shown in Figure 2, Panel A, both the volume of spoiler reviews and the

[1] IMDb was ranked 25th, Rotten Tomatoes 322nd, and Metacritic 841st for websites in the United States on Alexa.com, accessed July 2019.

[2] See Figure A1 in the Web Appendix for a Pareto chart of the distribution of spoiler reviews across movies.

Table 2. Descriptive Statistics.

Variable	Mean	Standard Deviation	Minimum	Maximum
Daily level				
DAILYREV (in $)	1,039,985	3,265,580	5	119,119,282
INTENSITY	2.48	2.69	0	45.17
PROP	.18	.15	0	1
CUMRATING	6.27	1.48	0	10
CUMVOL	12.87	247.73	0	4,276
ADVERT (in $1,000)	126.3	621.7	0	6,807
THEATERS	1,240	1,309	1	4,535
Movie level				
MPAA ratings				
G & PG	.15	.36	0	1
PG-13	.40	.49	0	1
R	.40	.49	0	1
Unrated	.05	.21	0	1
Genres				
Action	.09	.28	0	1
Adventure/Sci-Fi	.10	.30	0	1
Comedy	.20	.40	0	1
Drama	.32	.47	0	1
Family	.10	.30	0	1
Foreign	.02	.14	0	1
Horror	.06	.24	0	1
Musical	.02	.12	0	1
Romance	.02	.14	0	1
Thriller	.08	.27	0	1
Release type				
Limited release	.40	.49	0	1

volume of total reviews grow over time, though with greater momentum in the earlier rather than later days after movie release. We also plot the dynamics in the proportion of spoiler reviews in Figure 2, Panel B. The average proportion of spoiler reviews across movies is 26% on day one, which gradually increases to 31% by the end of the eighth week.

Measuring Spoiler Intensity

Uncovering Topics from Review Text

The construction of spoiler intensity requires revealing a multitude of plot-related topics from review texts (Property 2). We use machine learning—in particular, CTM (Blei and Lafferty 2005)—to uncover the set of topics included in movie reviews. CTM is an extension of latent Dirichlet allocation (LDA; Blei, Ng, and Jordan 2003), which marketing researchers have used to study emerging topics in scholarly articles (Wang et al. 2015), dimensions of customer product reviews (Tirunillai and Tellis 2014), and the predictive power of text in peer-to-peer loan applications (Netzer, Lemaire, and Herzenstein 2019). CTM replaces the Dirichlet distribution in LDA with a multinomial distribution in its data generation process. This modification allows flexible correlations between topics and therefore leads to an improved fit with the data (Blei and Lafferty 2005, 2007). Indeed, we find that CTM consistently outperforms LDA in terms of model fit in our empirical context, which provides support for the use of CTM in this study.[3] To apply the CTM, we prepared the textual data by removing stop words, tokenizing each word using a standard stemming algorithm, and removing sparse words that appear in less than 1% of movie reviews. This procedure yielded a preprocessed document-term matrix of 140,869 reviews (including both spoiler and nonspoiler reviews) represented by 1,624 unique words.

We refer to a movie review as a "document" and the collection of movie reviews as a "corpus." The CTM of each document from the corpus can be described as follows:

1. Draw $\eta | \{\mu, \Sigma\} \sim N(\mu, \Sigma)$.
2. For each word w contained in the document:
   a. Draw topic assignment variable $z | \eta$ from multinomial $(f(\eta))$,
   b. Draw a word $w | z, \beta$ from multinomial (β)

where $f(\eta)$ in step 2a maps a natural parameterization of $\eta = (\eta_1, \ldots, \eta_K)$ to the vector of topic probabilities $\theta = (\theta_1, \ldots, \theta_K)$ expressed as:

$$\theta_k = f(\eta_k) = \frac{\exp(\eta_k)}{\sum_{k=1}^{K} \exp(\eta_k)} \tag{1}$$

The data generation process of CTM can be interpreted as follows. When a user starts writing a movie review, they first decide on the weight of each topic (θ_k) that will appear in the movie review from a fixed number of topics (K). When choosing which word to write, the user selects a topic (z) according to its probabilistic distribution (multinomial (θ)). Conditional on the topic (z), the user's word choice (w) is then drawn from the associated distribution (multinomial (β)). The mapping of η to θ in Equation 1 allows the $K \times 1$ vector of topic probabilities for each document to carry a correlational relationship from Σ. We estimated the posterior distribution of the latent variables using a variational expectation-maximization algorithm (Roberts, Stewart, and Airoldi 2016). We refer interested readers to Blei and Lafferty (2007) for the derivation of the posterior distribution for CTM.

The CTM assumes a fixed number of topics K, which is a hyperparameter that researchers must predetermine (Blundell et al. 2009). We used the algorithm proposed by Lee and Mimno (2017), which estimates the vertices of the convex hull of word co-occurrences using a method of t-distributed stochastic neighbor embedding. Compared to cross-validation, an advantage of this algorithm is the computational efficiency for large data sets like the one in this study. We find that $K = 61$ is the optimal number of topics for movie reviews (including spoiler and nonspoiler reviews). We named each topic using its representative words and present all topics in Table A1 in the Web Appendix.

[3] See Figure A2 in the Web Appendix for details regarding the model fit comparison between CTM and LDA.

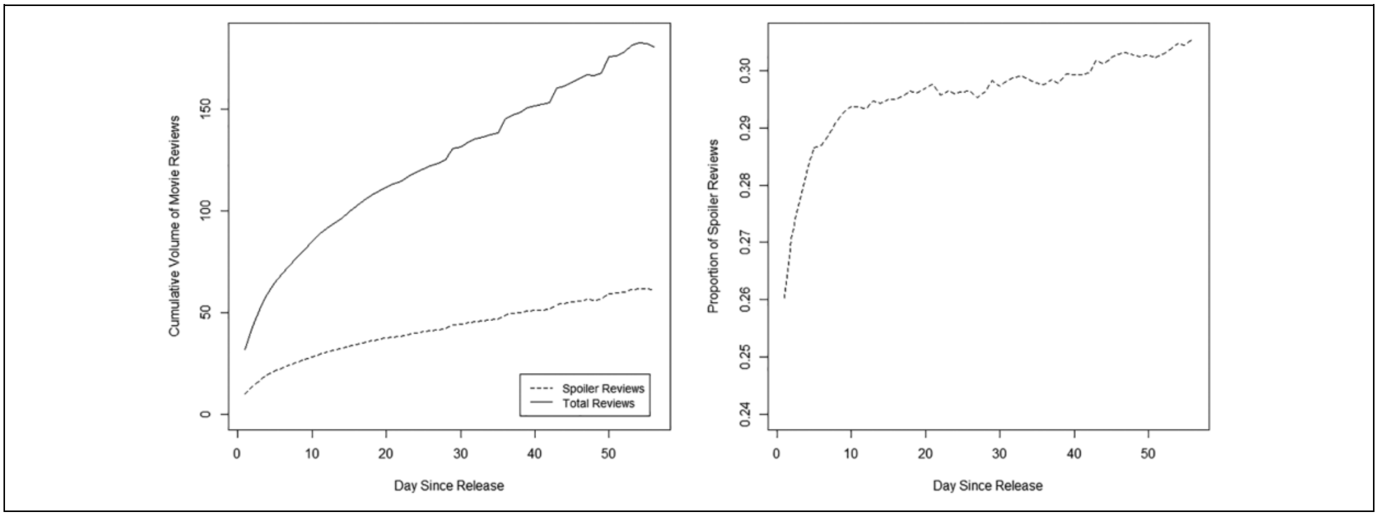

Figure 2. Cumulative volume and proportion of spoiler reviews over time.

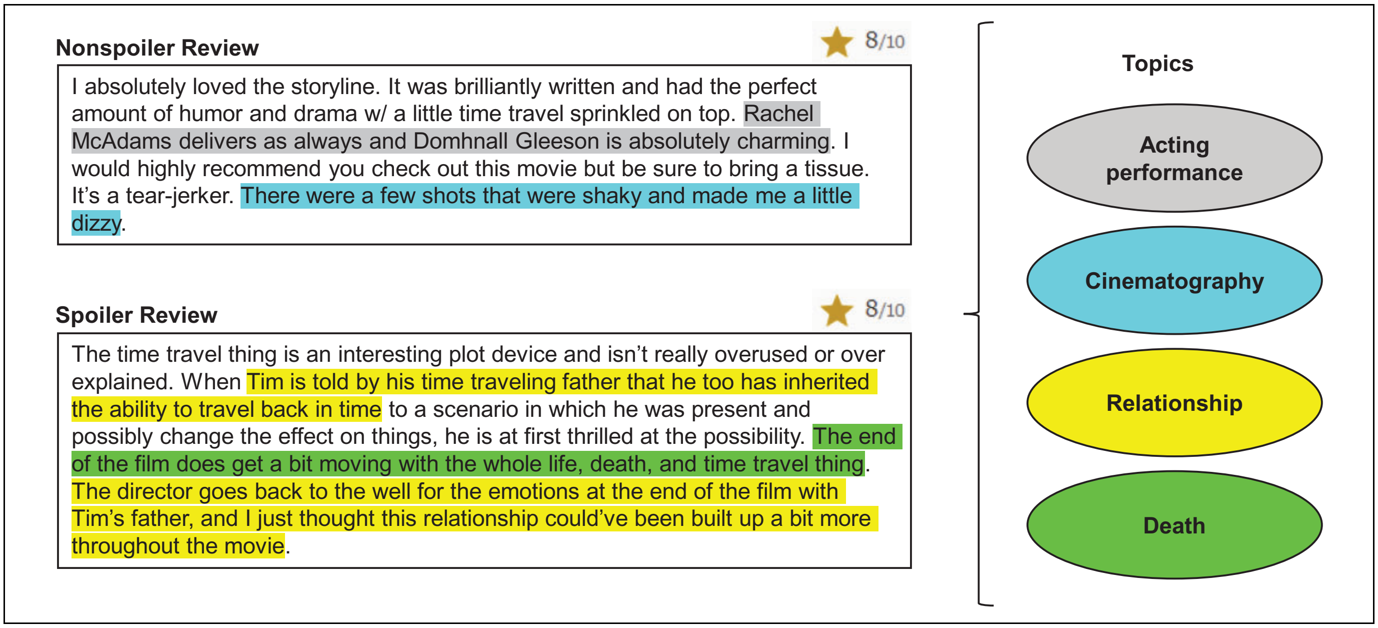

Figure 3. Examples of spoiler and nonspoiler reviews for the movie *About Time*.

Identifying Spoiling Topics

As not all topics resolve plot uncertainty, we further relied on the difference between text in spoiler and nonspoiler reviews to identify the set of topics deemed important in resolving plot uncertainty. To better explain the intuition behind the identification strategy, we provide examples of a spoiler review and a nonspoiler review in Figure 3, both of which are real reviews for the movie *About Time*. Notably, the text of each review can be well summarized by its underlying topics. For example, the nonspoiler review includes topics related to "cinematography" and "acting performance," and the spoiler review includes topics related to "relationship" and "death," as evidenced by sentences in their associated colors.

The topics revealed in the nonspoiler review and those revealed in the spoiler review are different in terms of the amount of plot-related information. Two plot-related topics that we clearly observe in the spoiler review are "death" (which occurred at the end of the movie) and "relationship" (between the protagonist Tim and his father). Although the nonspoiler review describes the movie as a "tear-jerker," the plot details as to why the movie is a "tear-jerker" are not provided. Another observation is that both reviews mentioned "time travel," suggesting that the topic of time travel is not regarded as spoiling

Table 3. Topics Associated with Spoiler Reviews.

Topic Name	Coefficient	Standard Error	Significance
America	1.582e−02	5.973e−03	**
Book	2.796e−02	3.489e−03	***
Character development	1.021e−02	2.610e−03	***
Death	5.816e−02	5.282e−03	***
Disappointment	4.457e−01	6.030e−02	***
Emotion	3.205e−02	3.718e−03	***
Fight	1.472e−02	2.448e−03	***
Ghost	9.521e−03	3.969e−03	*
Historical	7.063e−03	2.646e−03	**
Humans and robots	1.657e−02	5.245e−03	**
Kill	1.889e−01	4.407e−03	***
Length of movie	1.349e−02	2.329e−03	***
Lesson	7.624e−03	2.867e−03	**
Office	2.797e−02	3.144e−03	***
Overall evaluation	5.410e−02	2.892e−03	***
Relationship	3.160e−02	2.848e−03	***
Romance	2.939e−02	4.327e−03	***
Science fiction	8.301e−03	3.566e−03	*
Soundtrack	2.186e−02	3.471e−03	***
Space travel	2.850e−02	2.932e−03	***
Star Wars characters	3.177e−02	2.074e−03	***
Survival	1.425e−02	3.644e−03	***
Western	9.939e−03	4.860e−03	*

*$p < .05$.
**$p < .01$.
***$p < .001$.

for this movie. Therefore, not all topics that appear in movie reviews are regarded as spoiling: a spoiling topic is more likely to occur in spoiler than nonspoiler reviews, whereas a nonspoiling topic has either equal or higher likelihood to appear in nonspoiler reviews.

To identify spoiling topics, we ran a logistic regression in which the outcome variable was the review type (i.e., 1 = spoiler, and 0 = nonspoiler) and predictors were the number of words in a review associated with each topic. We operationalized the number of words from topic j in review l as $w_{jl} = \theta_{jl} \times n_l$, where θ_{jl} is the weight of topic j in review l from the estimation of CTM and n_l is the number of words in review l.[4] We also included movie dummies in the regression to account for movie heterogeneity.

We report in Table 3 the 23 topics that have significantly larger weights ($p < .05$) in spoiler reviews than in nonspoiler reviews. The top three spoiler-related topics (i.e., topics that weigh the most in spoiler reviews) are "disappointment," "kill," and "death." Not surprisingly, "kill" and "death" are often involved in critical plot points of movies (e.g., death of the main character). The topic "disappointment" is associated with words "worst," "ruin," and "disappoint." These are common words one might use when expressing one's unsatisfactory movie experience followed by the reveal of plot information as a justification. Although not presented in Table 3, the top three topics related to nonspoiler reviews (i.e., topics that weigh the most in nonspoiler reviews) are: "cinematography," "expectation," and "acting performance." The topic "cinematography" is associated with the words "beautiful," "visual," and "set," which are related to the visual appeal of the movie and therefore unrelated to movie plot. Similarly, "expectation" (associated with the words "time," "expect," and "watch") and "acting performance" (associated with the words "actor," "perform," and "role") are not directly associated with the plot of the movie. This comparison between the top topics related to spoiler and nonspoiler reviews provides some face validity to our identification of spoiling topics.

Constructing the Spoiler Intensity Metric

With the uncovered set of topics that constitute spoiling information, we further constructed the spoiler intensity metric following the guidance of the remaining three properties previously discussed (i.e., Property 1, 3, and 4). To better illustrate these properties, we provide two spoiler reviews for the movie *About Time* (one of which is the same review as in Figure 3), and one spoiler review for *The Lego Batman Movie* in Figure 4.

Comparing the two spoiler reviews for *About Time*, we noticed that spoiling topics may receive different degrees of elaboration. Spoiler review A provides more details on the topic of "relationship" than spoiler review B. In particular, spoiler review A reveals that the protagonist's father can time travel, as the ability is heritable, and that the father is involved in the movie's emotional ending. Spoiler review B provides a relatively limited description that the topic of relationship is not saccharine, and that it is taken seriously by the movie. As the degree of elaboration is often associated with the length of description, we use the number of words related to each topic (w_{jl}) as a proxy for the amount of plot-related information revealed in a spoiler review. This specification renders the spoiler intensity variable continuous and therefore satisfies Property 1.

Property 3 suggests that the degree of spoiling per topic might vary across movies. For example, in addition to spoiler reviews A and B, spoiler review C for *The Lego Batman Movie* in Figure 4 also discusses the topic of relationship (between Batman and the Joker). However, since *The Lego Batman Movie* is a comedy, the degree of spoiling from reading the topic of relationship is potentially less than for a romantic movie like *About Time.* As such, for each J = 23 plot-related topic, we quantify the degree of spoiling of topic j for movie i. Recall that the probability of review l associated with movie i being a spoiler review is predicted by the logistic model as follows:

$$y_{il} = \frac{\exp(\gamma w_l + \omega_i)}{1 + \exp(\gamma w_l + \omega_i)} \quad (2)$$

where $w_l = (w_{1l}, \ldots, w_{Kl})$, and ω_i is the fixed effect of movie i.

[4] We report in Web Appendix B.1 more details about the predictive power of topics from CTM.

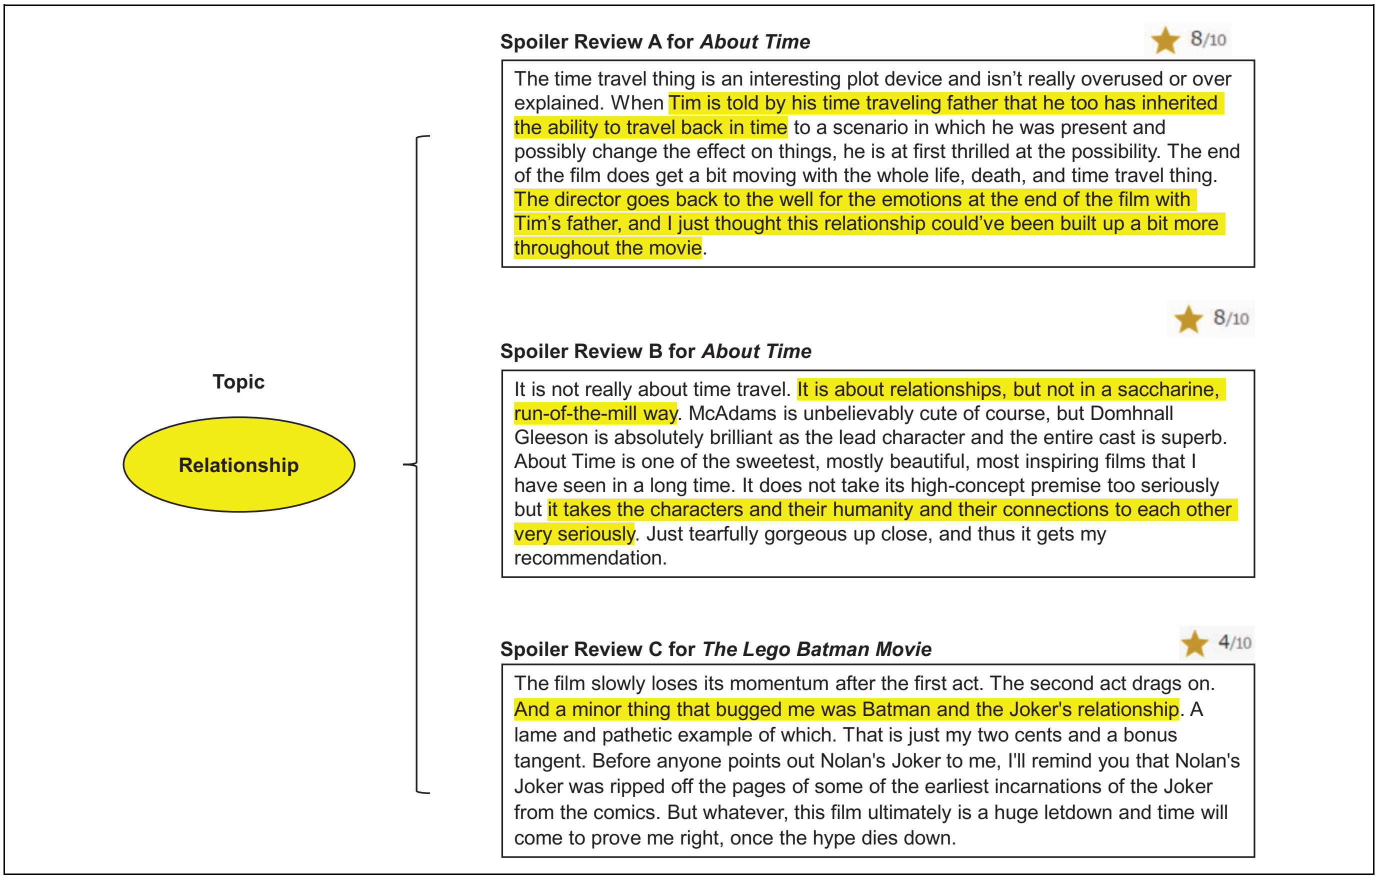

Figure 4. Example of topic distributed in spoiler reviews across movies.

We calculated the contribution of spoiling information from topic j in review l as follows:

$$c_{ijl} = \frac{\exp\left(\gamma(w_{jl}+1)+\delta w_{-jl}+\omega_i\right)}{1+\exp\left(\gamma(w_{jl}+1)+\delta w_{-jl}+\omega_i\right)} - \frac{\exp\left(\gamma w_{jl}+\delta w_{-jl}+\omega_i\right)}{1+\exp\left(\gamma w_{jl}+\delta w_{-jl}+\omega_i\right)} \quad (3)$$

where w_{-jl} is a vector of the number of words from other topics. The difference of the two terms on the right-hand side of Equation 3 measures the change in the likelihood (y_{il}) in response to a topic share increase in θ_{jl} (i.e., $(\theta_{jl}+\Delta\theta)\times n_l = w_{jl}+1$, where $\Delta\theta = \frac{1}{n_l}$). A greater c_{ijl} suggests a higher degree of spoiling from topic j in review l for movie i.

We aggregated the degree of spoiling of topic j for movie i, denoted by α_{ij}, using the normalized sum of c_{ijl} across all reviews of movie i as follows:

$$\alpha_{ij} = \frac{\sum_{l=1}^{L_i} c_{ijl}}{\sum_{j=1}^{J}\left(\sum_{l=1}^{L_i} c_{ijl}\right)} \quad (4)$$

where L_i represents the set of all reviews associated with movie i and $\sum_{j=1}^{J}\alpha_{ij} = 1$. The parameter α_{ij} in Equation 4 measures the spoiling effect of topic j for movie i, suggesting that the inclusion of α_{ij} in the spoiler intensity metric will satisfy Property 3.

Let S_{it} denote the set of spoiler reviews for movie i generated within a lagged time window of day t. We operationalized spoiler intensity of movie i on day t using all spoiler reviews from S_{it} as follows:

$$\text{INTENSITY}_{it} = \sum_{j=1}^{J} \alpha_{ij} \times \max_{l\in S_{it}} w_{jl} \quad (5)$$

Property 4 suggests that once the spoiling information related to a certain topic has been revealed, information from the same topic does not further reduce plot uncertainty when it reappears in subsequent reviews, unless additional information is provided. Consider again the two spoiler reviews for *About Time* in Figure 4. If an individual reads spoiler review A after spoiler review B, this individual can further reduce plot uncertainty because spoiler review A contains more specific plot details regarding the topic of relationship (e.g., with the protagonist's father, his involvement in the emotional ending) than spoiler review B, which only indicates that the movie treats relationships between characters seriously. However, if the order is reversed (i.e., reading spoiler review B after spoiler review A), it is unlikely for the individual to reduce plot uncertainty with spoiler review B because much of the plot-related information has been covered by spoiler review A. As such, we used the maximum function in Equation 5 to capture Property 4 of the spoiler intensity metric. The maximum function ensures that once a piece of information has been spoiled, it cannot be spoiled again. We provide evidence for the validity of the proposed spoiler intensity metric in capturing the level of spoiling information perceived by real people in Web Appendix B.2.

We chose the lag window used to construct spoiler-related variables (i.e., spoiler intensity and spoiler volume) to be ten days on the basis of a separate panel data set of movie reviews that we collected for 45 movies released in the United States in April 2019. For each movie, we tracked first-page spoiler reviews on IMDb daily in April 2019. The recency of spoiler reviews on the first page had a mean of 9.53 days, where we calculated the recency of each spoiler review by the difference between the date of observation and the date of creation. Therefore, we assumed that consumers typically read spoiler reviews generated within the last ten days.

Empirical Analysis

Model of Box Office Revenue

Let i denote movies and t denote the days after release. The dependent variable is $\ln(\text{DAILYREV})_{it}$, which represents the log-transformed daily box office revenue for movie i on day t. To examine the relationship between spoiler reviews and box office revenue, we considered the following model specification:

$$\begin{aligned}\ln(\text{DAILYREV})_{it} &= \beta_1 \ln(\text{DAILYREV})_{i,t-1} \\ &+ \beta_2 \ln(\text{INTENSITY})_{i,t-1} + \beta_3 \text{PROP}_{i,t-1} \\ &+ \beta_4 \ln(\text{CUMRATING})_{i,t-1} + \beta_5 \ln(\text{CUMVOL})_{i,t-1} \\ &+ \beta_6 \ln(\text{ADVERT})_{i,t-1} + \beta_7 \ln(\text{THEATERS})_{it} \\ &+ \beta_8 t + \beta_9 \text{HOLIDAY}_{it} \\ &+ \sum_{d=1}^{6} \gamma_j I\{\text{DAYOFWEEK}_{it} = d\} + \omega_i + \epsilon_{it}\end{aligned} \tag{6}$$

We included the lagged dependent variable, $\ln(\text{DAILYREV})_{i,t-1}$, on the right-hand side of Equation 6 to better capture the dynamics and indirectly control for past realizations of independent variables (e.g., WOM-related variables), which can persist to influence contemporaneous box office revenue (Keele and Kelly 2006). INTENSITY denotes the spoiler intensity described in Equation 5, and PROP measures the proportion of spoiler volume, defined as the moving average of the proportion of spoiler reviews to total movie reviews within the last ten days (i.e., from $t-10$ to $t-1$).[5]

For controls, we included the mean rating (CUMRATING) and volume (CUMVOL) of cumulative movie reviews because IMDb presents these summary statistics on the main page of each movie. We also included marketing mix variables, which comprise log-transformed advertising expenditure (ADVERT) and theater release count (THEATERS), as well as time-related variables, which comprise days after movie release (t), a dummy variable for federal holidays in the United States (HOLIDAY), and indicator variables ($I\{n\}$) for each day of the week (DAYOFWEEK).

In line with previous research (e.g., Duan, Gu, and Whinston 2008; Liu 2006), we lagged WOM-related and marketing mix variables except for the number of theaters to alleviate simultaneity concerns. We included ω, the movie fixed effect, to control for time-invariant heterogeneity of movies that include observable factors (e.g., budget, genre, star power) and unobservable factors (e.g., quality of the script, plot). Finally, ϵ is the idiosyncratic error term with a mean of zero.

Endogeneity Issues

It is well known that the inclusion of the lagged dependent variable as a predictor leads to a specific endogeneity issue known as the dynamic panel bias (Nickell 1981). As such, we estimated Equation 6 using the generalized method of moments (GMM) proposed by Blundell and Bond (1998). This estimation approach involves instrumenting the lagged dependent variable using both of its lagged levels and lagged differences. Our panel data allows the use of multiple lags (i.e., lags 2 and up) as GMM-type instruments to increase the efficiency of our estimation (Blundell and Bond 1998).

Unobserved time-variant characteristics of movies can induce a correlation between the regressors and the error term. The first potential source of endogeneity stems from the WOM-related variables. For example, unobserved offline WOM may increase both the demand for movies and the number of movie reviews. In addition, a user's interest in writing a spoiler review may also be associated with unobserved demand factors. Our solution follows Anderson and Hsiao (1981, 1982) to instrument the endogenous variable (CUMVOL, PROP, and INTENSITY) using its lagged level. The lagged levels of the endogenous variables are valid instruments under zero second-order autocorrelation (Anderson and Hsiao 1981, 1982), an assumption we empirically checked and confirmed.

Moreover, strategic information held by movie studios may also be a potential source of endogeneity. After a movie's release, studio managers may obtain private market information, allowing for adjustments of THEATERS and ADVERT. We followed previous research (e.g., Chintagunta, Gopinath, and Venkataraman 2010; Lu, Wang, and Bendle 2020) to use the means of

[5] We log-transform INTENSITY using $\ln(x+1)$. PROP is not in log because it is bounded between 0 and 1.

Table 4. Estimation Results of the Model of Box Office Revenue.

	OLS (1)	FE (2)	GMM with IVs for Lagged DV (3)	GMM with IVs for Lagged DV, WOM, and Marketing Mix (4)
Intercept	5.958***	-	-	-
	(.025)			
$\ln(DAILYREV)_{i,\ t-1}$	-	.474***	.606***	.638***
		(.011)	(.013)	(.017)
$\ln(INTENSITY)_{i,\ t-1}$	.180***	.045***	.077***	.060***
	(.010)	(.008)	(.014)	(.014)
$PROP_{i,\ t-1}$	.564***	−.016	.170**	.075
	(.041)	(.031)	(.055)	(.062)
$\ln(CUMRATING)_{i,\ t-1}$	.457***	−.004	.202***	.171***
	(.012)	(.020)	(.029)	(.027)
$\ln(CUMVOL)_{i,\ t-1}$	.140***	−.192***	.048***	.037***
	(.004)	(.014)	(.008)	(.010)
$\ln(ADVERT)_{i,\ t-1}$	.123***	.018***	.051***	.097***
	(.002)	(.002)	(.003)	(.007)
$\ln(THEATERS)_{it}$	.894***	.431***	.356***	.337***
	(.002)	(.010)	(.012)	(.020)
AGE(t)	−.033***	−.019***	−.012***	−.008***
	(3.33e−4)	(.001)	(.001)	(.001)
$HOLIDAY_{it}$	.759***	.572***	.558***	.533***
	(.023)	(.016)	(.018)	(.018)
DAYOFWEEK dummies	Yes	Yes	Yes	Yes
Movie fixed effects	No	Yes	Yes	Yes
Adjusted R-squared	.899	.952	-	-
Cluster-robust standard error	No	Yes	Yes	Yes
Number of observations			49,057	

Notes: OLS = ordinary least squares; FE = fixed effect; GMM = generalized method of moments; IV = independent variable; DV = dependent variable; WOM = word of mouth.
*$p < .05$.
**$p < .01$.
***$p < .001$.

THEATERS and ADVERT of other movies from the same genre as movie i and the same number of days t from the release as instruments for $THEATERS_{it}$ and $ADVERT_{it}$. The rationale for the relevance of these instruments is similar to that provided by Chintagunta, Gopinath, and Venkataraman (2010): movies of the same genre are likely to share similar release patterns and promotional strategies. The exclusion restrictions of these instruments stem from the fact that the means of marketing mixes set by other movies at different times are unlikely to be correlated with the current demand shock of the focal movie.

Empirical Findings

We begin with a standard ordinary least squares (OLS) regression of the model of box office revenue without the lagged dependent variable.[6] We report the results in Table 4, Column 1, which provides preliminary evidence that the association between spoiler intensity and box office revenue is positive and significant (.180, $p < .001$). We find that the association between spoiler volume and box office revenue is also positive and significant (.564, $p < .001$). Estimates for the control variables are of expected signs. For example, both CUMRATING and CUMVOL have positive associations with box office revenue. In addition, box office revenue is greater for movies that played in a larger number of theaters and spent more on advertising.

In Table 4, Column 2, we present model estimates using standard fixed effects regression and report robust standard errors clustered at the movie level. After controlling for time-invariant heterogeneity of movies, the association between spoiler intensity and box office revenue remains positive and significant (.045, $p < .001$), whereas the association between spoiler volume and box office revenue becomes nonsignificant (−.016, $p > .05$).

We report estimates using the GMM method (Blundell and Bond 1998) in Table 4, Columns 3 and 4, where robust standard errors clustered at the movie level are reported. We show the results with endogeneity correction for only the lagged dependent variable in Column 3 and endogeneity corrections for the lagged dependent variable, WOM-related variables, and marketing mix variables in Column 4. We conduct Hansen's J-test and the Arellano-Bond test for AR(2) to check the validity of over-identifying restrictions and second-order

[6] Including the lagged dependent variable in OLS leads to an almost perfect linear relationship (adjusted R-square of 1); therefore, the results are uninformative.

Table 5. Estimation Results from Robustness Checks.

	Simpler Measures of Spoiler Reviews		Alternative Specifications of Spoiler Intensity				Heterogeneous Trend
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$\ln(\text{DAILYREV})_{i,t-1}$	.637***	.625***	.647***	.649***	.661***	.652***	.622***
	(.017)	(.018)	(.017)	(.017)	(.016)	(.017)	(.018)
$\ln(\text{INTENSITY})_{i,t-1}$	-	-	-	-	-	-	.067***
							(.015)
$\ln(\text{INTENSITY}^{A})_{i,t-1}$	-	-	.072***	.102**	.064***	.060*	-
			(.012)	(.038)	(.009)	(.026)	
$\text{SPOILER}_{i,t-1}$	.082***	-	-	-	-	-	-
	(.017)						
$\ln(\text{Nwords})_{i,t-1}$	-	.036***	-	-	-	-	-
		(.005)					
$\text{PROP}_{i,t-1}$	-	−.072	4.68e−4	.100	−.021	.052	.097
		(.066)	(.056)	(.061)	(.055)	(.087)	(.069)
$\ln(\text{CUMRATING})_{i,t-1}$	.161***	.168***	.166***	.171***	.171***	.164***	.115**
	(.028)	(.027)	(.027)	(.027)	(.025)	(.027)	(.037)
$\ln(\text{CUMVOL})_{i,t-1}$	.050***	.033***	.028**	.037***	.023*	.040***	.043***
	(.010)	(.010)	(.009)	(.010)	(.010)	(.011)	(.011)
$\ln(\text{ADVERT})_{i,t-1}$	.101***	.101***	.089***	.085***	.080***	.094***	.103***
	(.007)	(.007)	(.006)	(.006)	(.006)	(.006)	(.007)
$\ln(\text{THEATERS})_{it}$	.338***	.340***	.335***	.332***	.309***	.317***	.344***
	(.020)	(.020)	(.019)	(.019)	(.017)	(.018)	(.020)
$\text{AGE}(t)$	-.008***	-.007***	-.007***	-.008***	-.007***	-.008***	−.013***
	(.001)	(.001)	(.001)	(.001)	(.001)	(.001)	(.003)
HOLIDAY_{it}	.534***	.535***	.531***	.532***	.528***	.528***	.536***
	(.018)	(.018)	(.018)	(.018)	(.018)	(.018)	(.018)
$\ln(\text{CUMRATING})_{i,t-1} \times t$	-	-	-	-	-	-	.003*
							(.001)
DAYOFWEEK dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Movie fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Endogeneity corrections	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Cluster-robust standard error	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Number of observations				49,057			

*$p < .05$.
**$p < .01$.
***$p < .001$.

autocorrelation, respectively. The p-values of the J-test and the test for AR(2) are .362 and .147 for Column 3 and .488 and .134 for Column 4, supporting the validity of proposed instruments and providing no evidence of second-order autocorrelation. Results from both GMM specifications show that the coefficient of INTENSITY is positive and significant. Although the estimate of PROP is positive and significant in Column 3, it becomes nonsignificant in Column 4 after the endogeneity corrections for WOM-related and marketing mix variables. We focus on the results in Column 4 in the rest of the article because of the more careful endogeneity corrections. The log-log model indicates that one percentage increase in spoiler intensity for movie i on day t is associated with a .06 percentage increase in box office revenue on the following day.

Robustness Checks

We checked the robustness of our findings against alternative measures of spoiler reviews and report estimation results from GMM with endogeneity corrections in Table 5.[7] We first re-estimated the model in Equation 6 using simpler measures of spoiler reviews. In Column 1, we consider a benchmark model to include SPOILER, a dummy variable that equals 1 if there is at least one spoiler review within the last ten days, to capture the relationship between the availability of spoiler reviews and box office revenue. In Column 2, we replace INTENSITY in Equation 6 with Nwords, the total count of words associated with spoiling topics in spoiler reviews within the last ten days. Consistent with the main findings, both SPOILER and $\ln(\text{Nwords})$ have positive and significant associations with box office revenue.

We further checked the sensitivity of our results against various aspects in the spoiler intensity specification. In particular, we considered an alternative spoiler intensity metric denoted by INTENSITY^{A}, which (1) assumed equal weight

[7] We conduct additional robustness checks to spoiler intensity from nonspoiler reviews in Web Appendix B.3.

(i.e., $\alpha_{ij} = \frac{1}{J}$ in Equation 5) among spoiling topics (Column 3) or (2) used an average function for aggregation (Column 4) or (3) used a sum function for aggregation (Column 5) or (4) used a longer lag window of three weeks (Column 6),[8] which covered 92.1% of first-page spoiler reviews according to the data collected in April 2019. Across Column 3 to Column 6, the coefficient of $\ln(\text{INTENSITY}^A)$ is positive and significant, whereas the coefficient of PROP is nonsignificant, supporting the robustness of our findings.

Lastly, we considered the possibility that high-quality movies can attract more intense spoiler reviews over time, creating the risk that the cross-sectional differences in box office dynamics can load onto the spoiler intensity variable. To test this possibility, we allowed for a heterogeneous time trend across movies by including an interaction term between CUMRATING and AGE, in which CUMRATING served as a proxy for movie quality. The results in Column 7 confirm that the positive effect of spoiler intensity still holds.

Underlying Mechanism

We further investigate the behavioral mechanism that may drive the positive effect of spoiler reviews on demand. For experiential products like movies, potential consumers often visit online review platforms to seek diagnostic information to resolve uncertainty (Dellarocas 2003; Goh, Heng, and Lin 2013). Compared to nonspoiler reviews, spoiler reviews are more diagnostic in reducing uncertainty because spoiler reviews can reveal important plot-related information as justification when critiquing a movie whereas nonspoiler reviews cannot. The reduction in potential moviegoers' uncertainty about movie quality due to spoiler reviews might lead to higher demand.

Moderator Analysis

To indirectly test the uncertainty-reduction mechanism of spoiler reviews, we considered four potential moderators of the effect of spoiler intensity: (1) release type (limited vs. wide release), (2) movie age, (3) advertising, and (4) average user rating. If uncertainty reduction is important, we expect the positive effect of spoiler intensity to be stronger under greater movie uncertainty.

We first considered whether the positive effect of spoiler intensity varies by the release type of a movie. Intuitively, it is in the movie studio's financial interest to play the movie in as many theaters as possible. However, a wide release strategy typically requires significant marketing investment and substantial negotiating power on behalf of the distributor (Kerrigan 2017). As a result, studios often reserve this strategy for mainstream and potential blockbuster movies, whereas they often employ a limited release strategy for independent movies. Compared to mainstream movies (which are typically developed to appeal to the masses and are thus more predictable or formulaic), independent movies are generally avant-garde and associated with higher uncertainty in terms of artistic quality (Holbrook 1999). As such, we anticipated that the positive effect of spoiler intensity would be stronger for independent movies, which, as noted, often use limited release.[9] We followed the literature to define a limited release movie, denoted by a dummy variable LIMITED, as a movie that plays in less than 700 theaters on its opening day, and a wide release movie (i.e., LIMITED = 0) as a movie that plays in more than 700 theaters (Fellman 2006; Kerrigan 2017).

We also examined the moderating role of movie age. We expected that consumers have higher movie uncertainty in the earlier (vs. later) period of a movie's life cycle because more quality signals (e.g., online WOM) become available as time progresses. For instance, past box office revenue can serve as a quality signal for potential moviegoers because high-quality movies tend to accrue greater ticket sales over time than low-quality movies (Moon, Bergey, and Iacobucci 2010). Following this rationale, we anticipated that the positive effect of spoiler intensity would be greater in the earlier period after the movie release due to the higher movie uncertainty.

It is well known that the informative function of advertising can reduce product uncertainty for potential buyers (Bagwell 2007; Hoch and Ha 1986). For example, Kim and Krishnan (2015) find that product descriptions and video commercials provided by online market platforms have a significant effect in reducing product uncertainty for intangible products. Moreover, Basuroy, Desai, and Talukdar (2006) suggest that advertising can serve as a credible signal of quality in the movie industry because any upward deviation of true quality in advertising content (i.e., overselling) can result in negative WOM and long-term harms, and therefore movie studios will not adopt this strategy. Considering these findings, we expected that the positive effect of spoiler intensity would be more salient for movies that spend less on advertising.

The last moderator we considered was average user rating. Compared to extreme ratings (either very high or low), ratings in the middle range tend to convey more ambiguous signals about movie quality (Tang, Fang, and Wang 2014). Thus, we expected that for movies with moderate or mixed ratings, consumers would be more likely to seek additional information to reduce movie uncertainty. Following this line of thought, we hypothesized an inverted-U relationship between the effect of spoiler intensity and average user ratings. To test this relationship, we classified our movie sample into quartiles on the basis of the average user ratings. We then included the first and fourth quartile dummies—denoted as QUART1 and QUART4, respectively—as moderators for spoiler intensity.[10]

We examined the moderators by re-estimating Equation 6 with additional interaction terms with spoiler intensity using GMM. We report estimation results in Table 6, in which

[8] We update the variable PROP using the three-week window accordingly.

[9] Release type is more objective and well-defined in the industry than movie type, which is subjective in nature.

[10] We do not include the interaction terms between spoiler intensity and average user rating and its squared term directly because of multicollinearity: the variance inflation factor has a mean of 13.44 and a maximum of 69.10.

Table 6. Estimation Results Examining Spoiler Intensity Interactions.

	Excluding Interaction with Quartile Dummies of Average User Ratings (1)	Including Interaction with Quartile Dummies of Average User Ratings (2)
$\ln(\text{INTENSITY})^{MC}_{i,\ t-1} \times \text{LIMITED}_i$	.351***	.439***
	(.067)	(.071)
$\ln(\text{INTENSITY})^{MC}_{i,\ t-1} \times t$	−.005***	−.002**
	(.001)	(.001)
$\ln(\text{INTENSITY})^{MC}_{i,\ t-1} \times \ln(\text{ADVERT})^{MC}_{i,\ t-1}$	−.186***	−.163***
	(.026)	(.022)
$\ln(\text{INTENSITY})^{MC}_{i,\ t-1} \times \text{QUART1}_i$	-	−.409***
		(.079)
$\ln(\text{INTENSITY})^{MC}_{i,\ t-1} \times \text{QUART4}_i$	-	−.250***
		(.052)
$\ln(\text{DAILYREV})_{i,\ t-1}$	.529***	.499***
	(.019)	(.018)
$\ln(\text{INTENSITY})^{MC}_{i,\ t-1}$	.085*	.154**
	(.041)	(.050)
$\text{PROP}_{i,\ t-1}$	−.119	−.136
	(.101)	(.109)
$\ln(\text{CUMRATING})_{i,\ t-1}$	.214***	.194***
	(.052)	(.054)
$\ln(\text{CUMVOL})_{i,\ t-1}$	.058***	.036*
	(.016)	(.017)
$\ln(\text{ADVERT})^{MC}_{i,\ t-1}$	.193***	.215***
	(.018)	(.018)
$\ln(\text{THEATERS})_{it}$	.430***	.484***
	(.021)	(.020)
AGE(t)	−.012***	−.010***
	(.001)	(.001)
HOLIDAY_{it}	.552***	.560***
	(.018)	(.018)
DAYOFWEEK dummies	Yes	Yes
Movie fixed effects	Yes	Yes
Endogeneity corrections	Yes	Yes
Cluster-robust standard error	Yes	Yes
Number of observations	49,057	

Notes: "MC" denotes mean-centered.
*$p < .05$.
**$p < .01$.
***$p < .001$.

Column 1 presents the results without the interactions between spoiler intensity and quartile dummies of average user ratings and Column 2 presents results using the complete set of moderators.[11] Given the consistency of estimates, we summarize findings by focusing on the results in Column 2. In line with our hypotheses, the positive and significant coefficient of the interaction between INTENSITY and LIMITED suggests that the positive effect of spoiler intensity is greater for limited release movies than for wide release movies. The negative and significant coefficient of the interaction between INTENSITY and t indicates a decay in the effect of spoiler reviews over time. Furthermore, the positive effect of spoiler intensity is negatively associated with advertising spending and is stronger for movies with moderate user ratings. The negative and significant coefficients of the interaction terms between INTENSITY and the two quartile dummies (i.e., QUART1 and QUART4) reveal an inverted-U relationship between average user rating and the effect of spoiler intensity on box office revenue. In summary, the results from the moderator analysis are consistent with the uncertainty-reduction mechanism of spoiler reviews.

[11] P-values of Hansen's J-test and the test for AR(2) are .713 and .306 for Column 1, and .645 and .498 for Column 2.

Event Study

We provide additional support for the uncertainty-reduction mechanism using an event study, which focuses on an exogenous IMDb website update made on December 11, 2017 (IMDb 2017). This update made two major changes to the way movie reviews are displayed on IMDb: (1) reviews are displayed *only* in the order of "helpfulness" (from most helpful to least helpful), whereas prior to the update reviews could be sorted in a variety of ways (e.g., by date, most positive/negative, etc.), and (2) the content of spoiler reviews is hidden by default and IMDb requires users to manually click on the spoiler review to see the content. The ability to sort reviews by methods other than helpfulness was restored after a subsequent update on February 10, 2018.

Theoretically, displaying all reviews according to their helpfulness should increase the diagnostic value of nonspoiler reviews and therefore decrease the relative usefulness of spoiler reviews in reducing movie uncertainty. In addition, hiding spoiler reviews by default increases the consumers' cost of reading spoiler reviews. Because of the decrease in relative benefit and the increase in cost of reading, we expected the positive effect of spoiler reviews on demand to be smaller after the IMDb update. We tested this hypothesis using data of 47 movies that were in play both before and after the IMDb update. Again, the effect of spoiler intensity is positive and statistically significant before the update. Furthermore, this positive effect is attenuated after the update, as predicted by the uncertainty-reduction mechanism. We provide more details on the data and estimations results in Web Appendix C.

Experiments

To directly test for the uncertainty-reduction mechanism of spoiler intensity, we designed surveys and conduct experiments online using Amazon Mechanical Turk (MTurk). The randomized experiment also allowed us to examine the causal effect of spoiler intensity on moviegoing decisions under varying levels of movie uncertainty.

Design and procedure. For the first experiment, we focused on the movie *Before I Go to Sleep,* an R-rated thriller. We recruited 545 participants from MTurk (59.2% female, mean age = 38.1 years) who resided in the United States and had completed at least 95% of their previous tasks. We employed a 2 (high spoiler intensity vs. low spoiler intensity) × 2 (high uncertainty vs. low uncertainty) between-subjects design.

Participants were shown the title of the movie, the poster, and a synopsis taken from a Google search of the movie's title. We instructed all participants to read two artificial user reviews, one of which rated the movie eight stars, while the other rated the movie two stars out of a ten-star scale. We created the content of the reviews such that we could randomly assign either of the two ratings to participants. The purpose of the conflicting ratings was to induce uncertainty about the quality of the movie. One of the two reviews presented was always a spoiler review, which we tagged using the same warning label used by IMDb. We randomized the order of the movie reviews displayed to each participant to control for potential order effects. We present details on the experimental stimuli in Web Appendix D.1.

We manipulated spoiler intensity by modifying only one sentence in the spoiler review while holding other sentences constant. Participants assigned to the high–spoiler intensity conditions read a spoiler review that gave away slightly more about the movie's narrative than that in the low–spoiler intensity conditions. Using the proposed intensity metric, we verified that the spoiler review in the high condition was indeed more intense than in the low condition (43.93 vs. 40.31).

We manipulated movie uncertainty by varying the content of the video clips that we asked participants to watch before reading the reviews. Participants assigned to the low-uncertainty conditions were provided a short clip of a promotional interview with the star of the movie, Nicole Kidman, in which she discussed her character in the narrative and how she prepared for the role. Participants assigned to the high-uncertainty conditions were instead given behind-the-scenes compilation footage of a similar length, which showed the filming and production on-set but no information related to the actual narrative of the movie. We chose these videos because movie studios commonly use star interviews and behind-the-scenes footage during promotions. Therefore, these videos are unlikely to be out-of-place as a manipulation for the experimental subjects.

Once participants finished reading the two reviews, we asked them to rate on a seven-point scale how much they agreed with the statement, "I want to go see this movie in theaters." This served as the dependent variable for the experiment. At the end of the survey, we asked participants whether they had seen the movie before, as well as their gender and age.

Manipulation check. We recruited 100 participants from MTurk (56.0% female, mean age = 36.7 years) in a pretest to check the validity of manipulations. As in the main experiment, participants were first shown the title, the poster, and the synopsis for the movie. The participants were then randomly assigned to watch either the interview clip shown in the low-uncertainty conditions, or the behind-the-scenes footage shown in the high-uncertainty conditions of the main experiment. We then asked participants to rate on a seven-point scale how much they agreed with the statement, "I feel certain about the quality of this movie."

Table 7. Descriptive Statistics of Experiment Conditions.

	Uncertainty	
DV: Willingness to Watch	**Low**	**High**
Before I Go to Sleep		
Low spoiler intensity	3.94 (1.77) n = 128	3.47 (1.70) n = 129
High spoiler intensity	3.79 (1.77) n = 130	3.94 (1.75) n = 126
Rise of the Guardians		
Low spoiler intensity	3.49 (1.66) n = 59	2.96 (1.53) n = 54
High spoiler intensity	3.41 (1.69) n = 56	3.78 (1.54) n = 51

To check whether movie uncertainty can be successfully manipulated by showing reviews with conflicting ratings, we randomly assigned the participants to read reviews with the conflicting ratings either shown or hidden while keeping the review content constant. For simplicity, we used reviews from the low–spoiler intensity conditions for this manipulation check. We then asked participants to rate on a seven-point scale how much they agreed with the statement, "these reviews make me feel certain about the quality of this movie."

Because of the focus on movie uncertainty, we removed eight participants who indicated that they had seen the movie before. Using a two-sample t-test, we found that those who watched the interview clip indicated higher certainty about the quality of the movie (n = 52, M = 5.08, SD = 1.43) than those who watched the behind-the-scenes footage (n = 40, M = 4.43, SD = 1.52), at a significance level of $p = .039$. We also found that those who did not see conflicting ratings indicated higher feelings of movie certainty (n = 44, M = 4.86, SD = 1.11) than those who saw conflicting ratings (n = 48, M = 4.21, SD = 1.35), at a significance level of $p = .013$.

In a separate pretest, we recruited 200 participants from MTurk (66.0% female, mean age = 40.1 years) to check whether reading the more intense spoiler review reduced uncertainty about movie quality. After showing participants basic information about the movie (i.e., title, poster, and synopsis), we randomly assigned participants to read the spoiler review used in either the high–spoiler intensity conditions or the low–spoiler intensity conditions. We then asked participants to rate on a seven-point scale how much they agreed with the statement, "this review makes me feel certain about the quality of this movie." We excluded 23 participants because of their prior experience watching the movie. We observed that those who read the more intense spoiler review indicated higher certainty about the quality of the movie (n = 84, M = 4.65, SD = 1.41) than those who read the less intense spoiler review (n = 93, M =

Table 8. Experiment Results.

DV: Willingness to Watch	*Before I Go to Sleep* (1)	*Rise of the Guardians* (2)
Intercept	3.938***	3.491***
	(.155)	(.210)
Intensity	−.145	−.081
	(.218)	(.301)
Uncertainty	−.465*	−.529
	(.218)	(.303)
Intensity × Uncertainty	.609*	.902*
	(.309)	(.435)
N	513	220

Notes: Standard errors are in parentheses.
*$p < .05$.
**$p < .01$.
***$p < .001$.

4.15, SD = 1.63), at a significance level of $p = .029$. Thus, results from this pretest support the uncertainty-reduction function of spoiler intensity.

Main results. We dropped 32 participants (6% of the sample) who indicated that they had seen the movie before. We dummy coded the two experimental factors such that Intensity = 1 for participants in the high–spoiler intensity conditions, and Uncertainty = 1 for participants in the high-uncertainty conditions. We present the descriptive statistics (mean, standard deviation, and size) for each condition in Table 7. We estimated the effects of the two factors as well as their interaction using OLS, and we report the results in Column 1 of Table 8 with standard errors in parentheses.

We find a significant and positive interaction effect between Uncertainty and Intensity (.609, $p < .05$). Given the presence of an interaction effect, we further examined the simple effects of spoiler intensity by partitioning the sample into high-uncertainty and low-uncertainty conditions. We then estimated the effect of Intensity on the dependent variable for each sample separately. For those who watched the behind-the-scenes footage (i.e., high-uncertainty condition), reading a more spoiling review increased willingness to watch the movie (.464, $p < .05$). However, for those who watched the interview with the star (i.e., low-uncertainty condition), the effect was statistically nonsignificant (−.145, $p > .10$). These findings are consistent with the proposed uncertainty-reduction mechanism of spoiler reviews. When consumers have low uncertainty about movie quality, they might not benefit much from the additional reduction in uncertainty that results from reading spoiler reviews, which explains the null effect we found in the low-uncertainty condition. It is worth noting that the low-uncertainty condition we used in the experiment is a relatively strong manipulation because most real consumers might not watch an in-depth interview with a movie's actors before making their moviegoing decisions.

Replication using a movie from a different genre. We ran a second experiment using *Rise of the Guardians,* a PG-rated animated movie, to test the replicability of our experimental findings. We recruited 272 participants from MTurk (60.3% female, mean age = 37.9 years) for the second experiment. The experimental design and procedure were identical to the first experiment. We report details of the manipulation checks in Web Appendix D.2. We excluded 52 participants (19% of the sample) who indicated that they had already seen the movie. As before, we report descriptive statistics in Table 6 and report estimation results in column 2 of Table 7. Consistent with the results from the first experiment, we find a significant and positive interaction between Intensity and Uncertainty (.902, $p < .05$). An analysis of simple effects for spoiler intensity again confirmed previous findings: the effect of high-intensity spoiler review is positive and significant (.821, $p < .01$) in the high-uncertainty condition and is statistically nonsignificant (−.081, $p > .10$) in the low-uncertainty condition.

Discussion

Although the relationship between spoilers and media enjoyment has received some academic attention, the relationship between spoilers and demand remains a knowledge gap in the literature. In this research, we show that the degree of plot uncertainty resolved by movie reviews (i.e., spoiler intensity) has a positive and significant association with box office revenue with an elasticity of .06. In addition, we provide evidence that uncertainty reduction is the behavioral mechanism that drives the positive effect of spoiler intensity using various methods (i.e., moderator analysis, event study, online experiments). Our finding of the positive association is novel in the movie industry, in which the conventional knowledge is that spoilers hurt box office revenue. Moreover, our conceptual framework of spoilers can be generalized to other product categories (e.g., television shows, role-playing games, novels, etc.). Although we find a positive net effect of spoiler reviews in the movie context, the relative importance of the positive uncertainty-reduction effect and the negative surprise burst effect of spoilers may vary across product categories and therefore warrants further investigation.

Managerial Implications

Our findings provide important managerial implications for movie studios, theaters, and review platforms. Foremost, our results suggest that online review platforms can potentially increase consumer welfare in the entertainment industry. The uncertainty-reduction mechanism we have uncovered suggests that a spoiler-friendly review platform can provide diagnostic plot-related information through spoiler reviews to help consumers make purchase decisions. Accordingly, we recommend that online review platforms maintain the availability of spoiler reviews, especially plot-intense spoiler reviews for potential consumers. We also recommend that review platforms keep the warning labels of spoiler reviews because of the benefit of allowing consumers to self-select into their exposure to spoilers. These spoiler alerts reduce the search cost for consumers who seek to reduce movie uncertainty while shielding consumers who care about movie enjoyment from the unfavorable effects of spoiler reviews. Furthermore, with advances in information technology, online review platforms can even go one step further to customize the number of displayed spoiler reviews and adjust the prominence of warning labels, catering to an individual consumer's preference revealed by their historical spoiler reading behavior.

Second, movie studios and theaters should actively monitor the content of spoiler reviews to better forecast future box office revenue. To demonstrate the predictive power of spoiler intensity, we randomly split the data into quarters, and then used three-quarters of the data as a training sample and the remaining quarter as a hold-out sample. By adding WOM-related variables individually to the benchmark model without any WOM-related variables, we calculated the predictive power of each WOM-related variable using the lift in the model's R-squared on the holdout sample. We find that the lift in R-squared is .010 for spoiler intensity, .007 for spoiler volume, .011 for WOM volume, and .004 for WOM valence, suggesting that spoiler intensity explains 1% of data variation. More importantly, the predictive power of spoiler intensity is slightly below that of WOM volume and more than twice that of WOM valence. Given the industry routine of monitoring WOM volume and valence in forecasting, we recommend that movie studios and theaters also actively monitor the content of spoiler reviews to improve forecasting performance.

Third, the benefit of monitoring the spoiler intensity of movie reviews, an act of social listening, is greater for movies with less advertising spending. To support this claim, we conducted a spotlight analysis to examine the elasticity of spoiler intensity at different levels of advertising. Specifically, we calculated the elasticity of spoiler intensity for advertising at the 25th ($1,243 per day) and 75th percentiles ($3,452 per day), respectively. We find that for movies with low levels of advertising (25th percentile), the elasticity of spoiler intensity is significant and large (.234, $p < .001$)—almost four times the magnitude of the elasticity for an average movie (.060, $p < .001$). However, the elasticity for movies with high levels of advertising is statistically nonsignificant (.067, $p > .05$). These findings suggest that movies with relatively small advertising budgets (e.g., most movies released by independent and arthouse studios) benefit the most from monitoring the content of spoiler reviews.

Fourth, the decay of the positive effect of spoiler intensity over time suggests that managers should make greater monitoring efforts in the earlier, rather than later, period of a movie's life cycle. To identify the specific window in which it is most beneficial to monitor spoiler reviews, we conducted a spotlight analysis for the elasticity of spoiler intensity at different days after the movie's release. We find that the elasticity is the greatest on the opening day (.149, $p < .01$) and then steadily declines (i.e., Week 1: .129, Week 2: .110, Week 3: .093, all with $p < .05$) until it becomes statistically nonsignificant at the end of the fourth week (.077, $p > .05$).

Finally, we highlight the boundary conditions under which movie studios might benefit from encouraging more intense spoiler reviews that can help reduce uncertainty about movie quality. In particular, our findings suggest that for movies with small advertising budgets and mixed user ratings, marketing managers should place great emphasis on stimulating online WOM, including WOM that might spoil the movie plot. However, for movies with large advertising budgets and extreme user ratings, we do not recommend that movie managers encourage consumers to generate spoiler reviews because of the lack of a significant effect on sales. In addition, the creation of spoiler reviews after three weeks post movie release does not seem to generate an economically meaningful impact on sales either. Although a no-spoiler policy is not recommended, we also caution movie studios that the dissemination of spoilers is sometimes uncontrollable. For example, spoiler reviews on IMDb can spread through social media where warning labels do not exist, which makes consumers more subject to the unfavorable effect of spoilers.

Limitations and Directions for Future Research

We note several limitations of this study, all of which provide promising directions for future research. Although we focused on online movie review platforms as the main source of online WOM, these platforms represent only one source of online WOM—one that consumers must actively seek out. Future research could explore whether our findings can be generalized to spoilers on social media platforms, where users are more likely to read spoilers by chance. Furthermore, we focus on the net effect of spoilers in this research. Future research could test for a parallel mediation of spoilers on movie demand, with a positive path through uncertainty reduction and a negative path through the burst of surprise.

We also note that this research focuses on spoilers that are generated by consumers. While we find a positive net effect of spoiler reviews, our results may not generalize to "leaks." Leaks, unlike spoilers, refer to information that is typically released from the supply side (e.g., movie producers, staff), either accidentally or maliciously prior to a movie's release. Leaks can take many forms but are often disseminated through images (e.g., photos taken on-set, posters, etc.) and videos (e.g., production footage, unedited clips, etc.). Although the effects from leaks would be controlled by movie fixed effects in our model, conceptual questions remain as to how spoilers and leaks differ in affecting ticket sales and whether they operate by the same behavioral mechanism. We leave these questions to future researchers.

Spoilers may also appear in other media, such as images on Pinterest and videos on YouTube. A notable feature of IMDb is that it offers expressive freedom to consumers at a relatively low cost of content generation (i.e., it is free to create an account and write reviews). In contrast, the creation of images and videos often requires skills such as artistic design, content editing, and so on, suggesting a high cost of content generation. Consequently, we suspect that professional content creators are mostly responsible for generating spoilers on platforms that focus on images and videos. With advances in machine learning and unstructured data analysis, future research could examine how user-generated spoilers delivered through media other than text affect consumer demand.

Finally, we use the max function in the specification of spoiler intensity to capture the discounted contribution of spoiling information that has appeared in previous spoiler reviews without the knowledge of individual review-viewing behavior. Because detailed review-viewing data are typically unavailable to movie studios, the proposed spoiler intensity metric should be useful to managers in the movie industry and therefore serves as a first step. Should individual-level data become available, future studies could relax the assumptions we made and extend the spoiler intensity metrics for both academics and practitioners.

Acknowledgments

This article is based on the first author's dissertation. The authors thank Charles Weinberg and Mark Houston for their helpful feedback on previous versions of this article. The authors also thank the participants of 2018 AMA Winter Academic Conference in New Orleans, 2018 INFORMS Marketing Science Conference at Temple University, 2018 Marketing Dynamics Conference at Southern Methodist University, 2018 Conference on Digital Marketing and Machine Learning at Carnegie Mellon University, 2019 INFORMS Annual Meeting in Seattle, and 2020 INFORMS Marketing Science Virtual Conference for all their valuable suggestions and comments.

Associate Editor

Michael Trusov

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This study was supported by Social Sciences and Humanities Research Council of Canada 435-2019-0597.

ORCID iDs

Jun Hyun (Joseph) Ryoo https://orcid.org/0000-0002-9434-1724
Xin (Shane) Wang https://orcid.org/0000-0003-1257-6931
Shijie Lu https://orcid.org/0000-0002-4180-6022

References

Alba, Joseph W. and Elanor F. Williams (2013), "Pleasure Principles: A Review of Research on Hedonic Consumption," *Journal of Consumer Psychology*, 23 (1), 2–18.

Anderson, Theodore W. and Cheng Hsiao (1981), "Estimation of Dynamic Models with Error Components," *Journal of the American Statistical Association*, 76 (375), 598–606.

Anderson, Theodore W. and Cheng Hsiao (1982), "Formulation and Estimation of Dynamic Models Using Panel Data," *Journal of Econometrics*, 18 (1), 47–82.

Babić Rosario, Ana, Francesca Sotgiu, Kristine De Valck, and Tammo HA Bijmolt (2016), "The Effect of Electronic Word of Mouth on Sales: A Meta-Analytic Review of Platform, Product, and Metric Factors," *Journal of Marketing Research*, 53 (3), 297–318.

Bagwell, Kyle (2007), "The Economic Analysis of Advertising," in *Handbook of Industrial Organization*, Mark Armstrong and Robert Porter, eds, 1st ed, Vol. 3, Chap. 28. Amsterdam: Elsevier, 1701–844.

Basuroy, Suman, Kalpesh K. Desai, and Debabrata Talukdar (2006), "An Empirical Investigation of Signaling in the Motion Picture Industry," *Journal of Marketing Research*, 43 (2), 287–95.

Blei, David M. and John Lafferty (2005), "Correlated Topic Models," in *Proceedings of the 18th International Conference on Neural Information Processing Systems,* Cambridge, MA: MIT Press, 147–54.

Blei, David M. and John Lafferty (2007), "A Correlated Topic Model of Science," *The Annals of Applied Statistics*, 1 (1), 17–35.

Blei, David M., Andrew Y. Ng, and Michael I. Jordan (2003), "Latent Dirichlet Allocation," *Journal of Machine Learning Research*, 3, 993–1022.

Blundell, Richard and Stephen Bond (1998), "Initial Conditions and Moment Restrictions in Dynamic Panel Data Models," *Journal of Econometrics*, 87 (1), 115–43.

Blundell, Jonathan, Sean Gerrish, Chong Wang, Jordan L. Boyd-Graber, and David M. Blei (2009), "Reading Tea Leaves: How Humans Interpret Topic Models," in *Advances in Neural Information Processing Systems*, 288–96.

Chevalier, Judith A. and Dina Mayzlin (2006), "The Effect of Word of Mouth on Sales: Online Book Reviews," *Journal of Marketing Research*, 43 (3), 345–54.

Chintagunta, Pradeep K., Shyam Gopinath, and Sriram Venkataraman (2010), "The Effects of Online User Reviews on Movie Box Office Performance: Accounting for Sequential Rollout and Aggregation Across Local Markets," *Marketing Science*, 29 (5), 944–57.

Clemons, Eric K., Guodong Gordon Gao, and Lorin M. Hitt (2006), "When Online Reviews Meet Hyperdifferentiation: A Study of the Craft Beer Industry," *Journal of Management Information Systems*, 23 (2), 149–71.

Deighton, John, Daniel Romer, and Josh McQueen (1989), "Using Drama to Persuade," *Journal of Consumer Research*, 16 (3), 335–43.

Dellarocas, Chrysanthos (2003), "The Digitization of Word of Mouth: Promise and Challenges of Online Feedback Mechanisms," *Management Science*, 49 (10), 1407–24.

Dellarocas, Chrysanthos, Xiaoquan Zhang, and Neveen F. Awad (2007), "Exploring the Value of Online Product Reviews in Forecasting Sales: The Case of Motion Pictures," *Journal of Interactive Marketing*, 21 (4), 23–45.

Duan, Wenjing, Gu Bin, and Andrew B. Whinston (2008), "Do Online Reviews Matter? An Empirical Investigation of Panel Data," *Decision Support Systems*, 45 (4), 1007–16.

Ely, Jeffrey, Alexander Frankel, and Emir Kamenica (2015), "Suspense and Surprise," *Journal of Political Economy*, 123 (1), 215–60.

Fellman, Daniel R. (2006), "Theatrical Distribution," in *The Movie Business Book*, Maidenhead, UK: McGraw-Hill, 362–74.

Gelper, Sarah, Renana Peres, and Jehoshua Eliashberg (2018), "Talk Bursts: The Role of Spikes in Prerelease Word-of-Mouth Dynamics," *Journal of Marketing Research*, 55 (6), 801–17.

Godes, David and Dina Mayzlin (2004), "Using Online Conversations to Study Word-of-Mouth Communication," *Marketing Science*, 23 (4), 545–60.

Godes, David and José C. Silva (2012), "Sequential and Temporal Dynamics of Online Opinion," *Marketing Science*, 31 (3), 448–73.

Goh, Khim-Yong, Cheng-Suang Heng, and Zhijie Lin (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.

Gopinath, Shyam, Jacquelyn S. Thomas, and Lakshman Krishnamurthi (2014), "Investigating the Relationship Between the Content of Online Word of Mouth, Advertising, and Brand Performance," *Marketing Science*, 33 (2), 241–58.

Hennig-Thurau, Thorsten, Caroline Wiertz, and Fabian Feldhaus (2015), "Does Twitter Matter? The Impact of Microblogging Word of Mouth on Consumers' Adoption of New Movies," *Journal of the Academy of Marketing Science*, 43 (3), 375–94.

Hoch, Stephen J. and Young-Won Ha (1986), "Consumer Learning: Advertising and the Ambiguity of Product Experience," *Journal of Consumer Research*, 13 (2), 221–33.

Holbrook, Morris B. (1999), "Popular Appeal Versus Expert Judgments of Motion Pictures," *Journal of Consumer Research*, 26 (2), 144–55.

IMDb (2017), "Updates to User Reviews," (accessed July 17, 2019), https://getsatisfaction.com/imdb/topics/update-to-user-reviews.

Johnson, Benjamin K. and Judith E. Rosenbaum (2015), "Spoiler Alert: Consequences of Narrative Spoilers for Dimensions of Enjoyment, Appreciation, and Transportation," *Communication Research*, 42 (8), 1068–88.

Keele, Luke and Nathan J. Kelly (2006), "Dynamic Models for Dynamic Theories: The Ins and Outs of Lagged Dependent Variables," *Political Analysis*, 14 (2), 186–205.

Kerrigan, Finola (2017), *Film Marketing*. Abingdon, UK: Routledge.

Kim, Youngsoo and Ramayya Krishnan (2015), "On Product-Level Uncertainty and Online Purchase Behavior: An Empirical Analysis," *Management Science*, 61 (10), 2449–67.

Kooser, Amanda (2019), "*Avengers: Endgame* Spoilers? Just Don't, Russo Brothers Say," (accessed July 17, 2019), https://www.cnet.com/news/avengers-endgame-spoilers-just-dont-russo-brothers-say.

Leavitt, Jonathan and Nicholas Christenfeld (2011), "Story Spoilers Don't Spoil Stories," *Psychological Science*, 22 (9), 1152–54.

Leavitt, Jonathan and Nicholas Christenfeld (2013), "The Fluency of Spoilers: Why Giving Away Endings Improves Stories," *Scientific Study of Literature*, 3 (1), 93–104.

Lee, Moontae and David Mimno (2017), "Low-Dimensional Embeddings for Interpretable Anchor-Based Topic Inference," arXiv preprint arXiv:1711.06826.

Li, Xinxin and Lorin M. Hitt (2008), "Self-Selection and Information Role of Online Product Reviews," *Information Systems Research*, 19 (4), 456–74.

Liu, Xiao, Param Vir Singh, and Kannan Srinivasan (2016), "A Structured Analysis of Unstructured Big Data by Leveraging Cloud Computing," *Marketing Science*, 35 (3), 363–88.

Liu, Yong (2006), "Word of Mouth for Movies: Its Dynamics and Impact on Box Office Revenue," *Journal of Marketing*, 70 (3), 74–89.

Lu, Shijie, Xin Wang, and Neil Bendle (2020), "Does Piracy Create Online Word of Mouth? An Empirical Analysis in the Movie Industry," *Management Science*, 66 (5), 2140–62.

Moon, Sangkil, Paul K. Bergey, and Dawn Iacobucci (2010), "Dynamic Effects Among Movie Ratings, Movie Revenues, and Viewer Satisfaction," *Journal of Marketing*, 74 (1), 108–21.

Mudambi, Susan M. and David Schuff (2010), "Research Note: What Makes a Helpful Online Review? A Study of Customer Reviews on Amazon.com," *MIS Quarterly*, 34 (1), 185–200.

Netzer, Oded, Alain Lemaire, and Michal Herzenstein (2019), "When Words Sweat: Identifying Signals for Loan Default in the Text of Loan Applications," *Journal of Marketing Research*, 56 (6), 960–80.

Nickell, Stephen (1981), "Biases in Dynamic Models with Fixed Effects," *Econometrica*, 49 (6), 1417–26.

Olson, Gary M., Robert L. Mack, and Susan A. Duffy (1981), "Cognitive Aspects of Genre," *Poetics*, 10 (2/3), 283–315.

Roberts, Margaret E., Brandon M. Stewart, and Edoardo M. Airoldi (2016), "A Model of Text for Experimentation in the Social Sciences," *Journal of the American Statistical Association*, 111 (515), 988–1003.

Seiler, Stephan, Song Yao, and Wenbo Wang (2017), "Does Online Word of Mouth Increase Demand? (and How?) Evidence from a Natural Experiment," *Marketing Science*, 36 (6), 838–61.

Sun, Monic (2012), "How Does the Variance of Product Ratings Matter?" *Management Science*, 58 (4), 696–707.

Tang, Tanya, Eric Fang, and Feng Wang (2014), "Is Neutral Really Neutral? The Effects of Neutral User-Generated Content on Product Sales," *Journal of Marketing*, 78 (4), 41–58.

Tirunillai, Seshadri and Gerard J. Tellis (2014), "Mining Marketing Meaning from Online Chatter: Strategic Brand Analysis of Big Data Using Latent Dirichlet Allocation," *Journal of Marketing Research*, 51 (4), 463–79.

Trottier, David (1998), *The Screenwriter's Bible: A Complete Guide to Writing, Formatting*. Los Angeles: Silman-James Press.

Wang, Xin, Neil T. Bendle, Feng Mai, and June Cotte (2015), "The Journal of Consumer Research at 40: A Historical Analysis," *Journal of Consumer Research*, 42 (1), 5–18.

Zillmann, Dolf (1995), "Mechanisms of Emotional Involvement with Drama," *Poetics*, 23 (1/2), 33–51.

Zillmann, Dolf, T. Alan Hay, and Jennings Bryant (1975), "The Effect of Suspense and its Resolution on the Appreciation of Dramatic Presentations," *Journal of Research in Personality*, 9 (4), 307–23.