

Clickbait News and Algorithmic Curation

A Game Theory Framework of the Relation between Journalism, Users, and Platforms

Juliane Lischka, *University of Hamburg*

Marcel Garz, *Jönköping International Business School*

June 2021

Abstract

Algorithmic curation of social media platforms is considered to create a clickbait media environment. Although clickbait practices can be risky especially for legacy news outlets, clickbait is widely applied. We conceptualize clickbait content supply as a revision game with an unknown threshold. Combining supervised machine learning with time series analysis of Facebook posts and Twitter messages of 37 German legacy news outlets over 54 months, we observe outlets' behavior following algorithm changes. Results reveal (1) an infrequent use of clickbait with few heavy-using outlets and (2) turning points of clickbait performance as clickbait supply and user interaction form a reversed *U*-shaped relationship. News outlets (3) collectively adjust towards an industry clickbait standard. While we (4) cannot prove that algorithmic curation increases clickbait, (5) Facebook's regulative intervention to decrease clickbait disperses heterogeneous tendencies in clickbait supply. We contribute to an understanding of editorial decision-making in competitive environments facing platforms' regulative intervention.

Keywords: Algorithms, Digital Journalism, Facebook, Game Theory, Legacy Media, News, Social Media Platforms, Supervised Machine Learning, User Interaction, Twitter

Clickbait News Posting on Facebook and Twitter

Social media platforms provide an infinitely scrollable and personalized news feed intending to tie users in. To be part of the feed, Caplan and boyd (2018) argue that news outlets must construct their content with algorithmic curation in mind. Therefore, social interaction metrics such as likes or shares denote an online currency (Myllylahti, 2020) and serve as popularity cues for editors and users alike (Haim et al., 2018; Munger, 2020). By rewarding popularity, algorithmic curation can lead to conformity in the news industry (Caplan and boyd, 2018). News outlets collectively design “addictive distractions” (Naughton, 2018: 387) to grab user attention, maximize popularity and ultimately profits (Sismeiro and Mahmood, 2018).

Clickbait headlines signify one example of such addictive distractions created to succeed in a high-choice algorithmic-curation environment of social media. Clickbait represents a linguistic strategy to articulate a message in a curiosity-arousing way that entices readers to click on the referring article (Potthast et al., 2018). Clickbait headlines originate from new entrants to journalism such as Upworthy or BuzzFeed (Tandoc, 2018). Meanwhile, they are used by tabloids (Blom and Hansen, 2015), broadsheets (Potthast et al., 2018; Rony et al., 2017), and by hoax outlets to spread disinformation (Braun and Eklund, 2019). Particularly legacy news outlets find themselves in a dilemma of jeopardizing their reputation and user trust when providing too much clickbait (Molyneux and Coddington, 2020; Zhang et al., 2020).

Against the background of algorithmic curation and competition for user attention, we ask, *What are determinants of legacy news outlets’ clickbait headline supply on social media platforms over time?* Going beyond previous research focusing on clickbait performance, this study sheds light on the collective clickbait provision behavior of legacy news outlets facing algorithmic curation. Our goals are (a) to contribute to an understanding of journalistic decision-making in competitive social media environments and (b) to empirically assess “how algorithms and data-driven technologies, enacted by an organization like Facebook, can induce similarity across an industry” (Caplan and boyd, 2018: 1).

Empirically, we focus on the news industry level and analyze 856,252 Facebook posts and 1,305,297 tweets of $n = 37$ national and regional legacy news outlets in Germany, a country with a strong ethical-professional news industry (Hanitzsch et al., 2011) and comparatively

CLICKBAIT NEWS

modest social media use for news (Newman et al., 2018), representing a conservative example for the present research goals. The observation period begins in 2013, when the German news industry was widely active on social media and ends in 2017, when social media use for news had stagnated in Germany (Newman et al., 2018). During this period, Facebook provided an algorithmic curation environment and intended to reduce clickbait supply twice. Twitter shifted from non-curation to algorithmic curation, potentially inviting clickbait.

Literature Review

Clickbait News Making

Journalists consider various aspects for news selection and editing, including audience interaction (Ferrer-Conill, 2017), competitor content (Boczkowski, 2010), social media and third-party content (Cagé et al., 2020), and outlet reputation (Andina-Díaz and García-Martínez, 2020). Social media editors or members of the editorial team are involved in editing news for social media (Cornia et al., 2018). These editors develop conceptions of social media algorithms and “feel the need to optimize their content” for platforms (Peterson-Salahuddin and Diakopoulos, 2020: 27). Consequently, social media editors recognize that certain topics, e.g., health, work better on Facebook than foreign politics (Lischka, 2021). On social media, editors depart from the traditional text norms and add subjective views (Hågvar, 2019) or highlight surprise and emotions (Peterson-Salahuddin and Diakopoulos, 2020). Therefore, news outlet’s social media teasers are more tabloid-like (Magin et al., 2021) and include emojis and more question marks (Haim et al., 2021) than news on their websites. Peterson-Salahuddin and Diakopoulos (2020) conclude that social media editors constantly negotiate newsworthiness with interactionworthiness to be promoted through curation algorithms.

Clickbait represents one way to edit clickworthy teasers. Clickbait headlines highlight gossip, sensational, or provoking aspects or create curiosity gaps by withholding information, which includes listicle, forward-referencing, and question-based headlines (Scacco and Muddiman, 2020). Traditional (summary) headlines also intend to attract readers’ attention by giving more weight to certain aspects of a story while remaining precise (Andrew, 2007). In contrast, clickbait headlines exaggerate arousal and lack precision (Scacco and Muddiman, 2020). Clickbait is regarded as deviating from professional traditions (Hågvar, 2019). Hence,

journalists employ clickbait even if they do not favor it (Potthast et al., 2018). Since clickbait originates from digital-native news outlets (Tandoc, 2018), especially legacy news editors may face a dilemma when writing clickbait headlines.

Game-Theoretic Framework to Clickbait News Making

Clickbait headlines could be conceptualized resulting from heteronomous market forces in journalism based on field theory (Bourdieu, 1983), which deviate from the dominant view, the journalistic “doxa”. Field theory is applied to explain new entrances in journalism (Tandoc, 2018) and changes in journalistic culture and practices (Schultz, 2007). To explain institutional pressures and homogeneity within journalism facing algorithmic curation, the institutionalist perspective of isomorphism (DiMaggio and Powell, 1983) is useful, as adopted in Caplan and boyd (2018). The strength of field theory and isomorphism lies in connecting actors to their environment. Game theory focuses on decision-making while considering the environment of actors. Game theory provides a suitable framework for describing clickbait news supply, acknowledging the interests and actions of editors, competitors, users, and platforms.

A game is an abstract, usually mathematical, description of possible choices that players can make given their information about available actions and consequences (Osborne and Rubinstein, 1994). Players can be cooperative or do not exchange information (Nash, 1951; Neumann and Morgenstern, 1944). Game theoretic approaches are rarely applied to journalism. One exception demonstrates a leader-follower game, in which reputational news outlets withhold questionable scoops fearing reputation loss, while less reputable outlets cover the scoop (Andina-Díaz and García-Martínez, 2020). We describe the roles of journalists, users, competitors, and social media platforms for clickbait news supply in four scenes that can occur subsequently or simultaneously.

Scene 1: Game invitation.

The online environment of “clickbait media” (Munger, 2020), enabled by interaction-rewarding algorithmic curation (Caplan and boyd, 2018), signals cues to news outlets and hoax publishers alike to join an attention maximization game on social media platforms (Braun and Eklund, 2019; Myllylahti, 2020). Previous research describes platforms as powerful intermediaries (Nielsen and Ganter, 2017) while placing journalism into dilemmas with

professional standards (Peterson-Salahuddin and Diakopoulos, 2020). Journalists regard Facebook as a central distributor they invest resources into and depend on (Meese and Hurcombe, 2020). Nielsen and Ganter (2017: 1613) argue that “digital intermediaries act, leaving news media organizations to react.” If social media platforms efficiently initiate an attention maximization game by algorithmic curation, clickbait supply of news outlets should increase.

H1. Clickbait supply increases in algorithmic curation but not in a non-curative setting.

Scene 2: Collective outcomes through stigmergy.

While competing news organizations rarely coordinate news supply, news outlets observe each other’s news coverage leading to greater news homogeneity online (Boczkowski, 2010). Through mutual monitoring, news formats and style collectively adopt over time across the news industry, especially under fierce competition (Arbaoui et al., 2020), and even if individual journalists disapprove (Potthast et al., 2018). Accordingly, Haim et al. (2021) show that news outlets’ language is more homogeneous on social media than across their websites. Collective outcomes in journalism thus may involve coordination via environmental cues, which can also explain the universal usage of clickbait.

The concept of stigmergy describes this kind of coordination. Originally proposed to explain collective behavior of termites (Grassé, 1959), stigmergy is a way of indirect communication to coordinate actions through changes to the environment, which actors mutually perceive and respond to (Vrancx et al., 2010). While termites leave pheromone trails that coordinate nest building, human actors’ communicative cues within networks of users instigate stigmergy (Cimino et al., 2018). If clickbait dissemination in journalism is stigmergic, the following should hold.

H2. Outlets adjust their clickbait supply to an industry standard.

Scene 3: Playing against an unknown threshold.

Whether clickbait effectively maximizes interaction is questioned by studies showing that clickbait may either receive more (Rony et al., 2017) or less clicks (Kuiken et al., 2017; Mourão and Robertson, 2019; Scacco and Muddiman, 2020) than non-clickbait posts. Zhang et al. (2020) find an inverted *U*-shaped relationship between the number of clickbait headlines

CLICKBAIT NEWS

and page views. That is, traffic increases up to certain level of clickbait, but decreases if used too extensively.

An inverted *U*-shape can be justified with regret game theories (Starmer, 2000). Accordingly, the decision of whether to click on a piece of content depends on users' expected utility. If users frequently experience disappointment with clickbait from any source, because of the lower quality these articles have than summary headlines (Scacco and Muddiman, 2020), regretful users may avoid the type of content or source for future decisions. In this regard, Molyneux and Coddington (2020) argue that negative perceptions of clickbait headlines may aggregate over time, whenever a clickbait headline disappoints users, and ultimately risks losing their trust. Interaction with clickbait may decrease considerably, as reflected by lower popularity cues, including likes and shares (Haim et al., 2018). Disappointment follows a steep convex negative function (little disappointment is related to strong utility loss), whereas exceeding expected utility is characterized by a flatter concave positive utility function (a lot more utility than expected is required for extra utility gain) (Starmer, 2000). The combination of concave and convex utility functions provides an explanation for the reverse *U*-shaped relation between clickbait and clicks.

To learn about the optimal amount of clickbait, editors need to experiment. Chen's (2020) revision game of experimentation on an unknown threshold represents a risk-control game to explore an optimum. Chen (2020) uses the example of a food maker that tries to find out how much of a cheaper substance can be added to a food product without decreasing its quality or, worse, causing health issues. Similarly, news organizations may post clickbait headlines to increase traffic. But if a news outlet adds too many clickbait headlines over time, users' utility and the outlet's reputation may be compromised (Molyneux and Coddington, 2020). However, news outlets do not exactly know how much of their own and other clickbait is too much for users. By revising their own actions while observing user behavior, news outlets can incrementally approach the unknown clickbait threshold. Meese and Hurcombe (2020: 5) describe such behavior as "riding the waves of engagement." We propose,

H3. The more user interaction with past clickbait, the more clickbait is supplied at present.

Scene 4: Interfering in the game.

Outlets benefit from platforms with audience traffic (Sismeiro and Mahmood, 2018) and platforms from engaging content with user loyalty (Chen and Pain, 2019), suggesting a strategic alliance between platforms and outlets. In an alliance, the threat of opportunistic behavior of each party is small (Hansen et al., 2008). Yet, a downward slope of utility initiated by user regret can be harmful for social media platforms due to decreasing user loyalty—and it implies opportunistic behavior by content providers. Being providers for the outlet-user interaction, platforms can transform from an ally to a regulator. Their curation algorithms serve as bureaucratic mechanisms (Caplan and boyd, 2018) to enforce regulation. Platforms will be interested to maximize utility for users and punish outlets for crossing users' clickbait threshold. Facebook has intended to penalize clickbait in 2014 and 2016 due to displeased users (Peysakhovich and Hendrix, 2016), revealing a form of responsive regulator interference (Etienne, 2013). To outlets, regulation implies opportunistic behavior of the outlet's former ally and increases uncertainty towards the future behavior of the platform. Since journalistic routines are used to adapt to frequent algorithm shifts (dos Santos et al., 2018), news outlets may adhere to platform's algorithm changes, as shown for increases in ideological news diversity (Garz and Szucs, 2021) and social video supply (Tandoc and Maitra, 2018). Alternatively, outlets may not comply facing unpredictable algorithm changes (Cornia et al., 2018; Meese and Hurcombe, 2020). We ask,

RQ 1. How does clickbait supply change after a platform's regulative algorithm change?

Method

Sample

Platform selection.

Facebook and Twitter were selected as information-rich cases stratified regarding news outlets' dependency (high versus low), algorithmic curation (permanent versus incremental), intention to reduce clickbait (twice versus none), and a rather general versus information-elitist audience, respectively (Cornia et al., 2018; Myllylahti, 2020; Newman et al., 2018). Given

CLICKBAIT NEWS

Facebook's and Twitter's accessibility, the platform selection strategy combines stratified purposeful and convenience sampling (Suri, 2011).

Regarding Facebook's News Feed, two algorithm updates explicitly targeted clickbait. In August 2014, Facebook announced that its news feed would "weed out" clickbait (El-Arini and Tang, 2014). In August 2016, Facebook refined its clickbait detection using supervised learning techniques to curb the distribution of "misleading and spammy" clickbait more effectively (Peysakhovich and Hendrix, 2016).

Twitter users were exposed to tweets in reverse chronological order until 2016. In January 2015, Twitter announced to start testing features of algorithmic content selection to "surface a few of the best Tweets" (Rosiana, 2015). These features were gradually rolled out on an opt-in basis. In March 2016, Twitter switched to algorithmic curation as users' default option (e.g., Kamen, 2016).

Country and news outlet selection.

Germany is a country with a strong ethical-professional conduct in journalism (Hanitzsch et al., 2011), strong in print revenues compared to slowly growing online revenues (Newman, 2016), comparatively low social media use for news but increasing mobile news use during the observation period (Newman et al., 2018). Facebook was used for news by 25% and Twitter by 4% of users (Newman et al., 2018).

The sample of news outlets comprises 10 national and 27 regional news outlets with a legacy in print following a traditional, two-sided model of advertising and audience revenue, which had a Facebook and Twitter account during the study's observation period (see Figures 3 and 4). The sample thus stands for a potentially maximized tension between following traditional professional standards and adopting to algorithmic curation—while news outlets' revenue dependency on platforms is comparatively low.

Data Collection

We collected 856,252 posts and 1,305,297 tweets published between January 2013 and June 2017 by the official Facebook Pages and Twitter accounts of 37 German news outlets, which comprises the total posting activity of these outlets. The data include the date and time of publication, the post message, and the number of likes and shares. The dataset was compiled

using the Facebook Graph Application Programming Interface (API) combined with the mining application Netvizz. For Twitter, we used the Python library “GetOldTweets3.” This library offers an automated approach to access and parse the browser-based Twitter Search. We recovered each tweet’s date, time, message, and number of likes and retweets.

Supervised Machine Learning Approach

We used supervised learning to obtain clickbait classifications of all content in our dataset, based on the characteristics of the posts and tweets rated by human coders.

Human coding.

We drew random samples of 20,000 posts and 20,000 tweets and had humans evaluate whether these messages classify as clickbait. German-speaking coders were recruited via Fiverr.com, an online platform for freelance services. To verify the coding quality, we had each potential coder evaluate a test batch of 500 messages, which we had coded ourselves. We paid potential coders 20 Euros for rating these 500 messages and offered further work if their coding “fulfilled our quality standards”. That is, we compared their evaluation of the test batch with our own coding and selected only coders whose kappa-statistic (Landis and Koch, 1977) indicated that the ratings are systematically related to the gold standard (at the 1% level). The employed coders achieved inter-rater agreement rates between 94.8% and 98.0%, corresponding to kappa-statistics between 0.61 and 0.77. The coders received further batches of randomly selected post messages, each batch for 20 Euros incentive. We assured ongoing coding quality by spot checking the results. Importantly, we kept repeating our offer to ask for more work “as long as our quality standards were met,” which arguably created a strong incentive for the coders to take the task seriously. Each batch included a nondescript identifier, the post or tweet message, and one blank column for the evaluation. The files did not include date and outlet, to avoid priming the coders.

The coding instructions asked to evaluate whether a given post message qualifies as clickbait (“yes”, “no”, and “maybe”), considering the following instructions: “Clickbait refers to headlines designed to maximize clicks, often for commercial interests such as ad revenues. For instance, clickbait exaggerates or withholds information to arouse the attention and curiosity of the reader, so that the reader clicks on the headline and opens the linked story.” The

instructions included specific examples of non-clickbait and clickbait, including forward-references, questions, listicles, and sensationalization (see Table 1).

[Table 1 about here]

Clickbait classification.

The classification was implemented using the R package “caret – Classification And REgression Training” (Kuhn, 2008). To evaluate the power of different classification methods, we randomly divided the samples of coded messages into training (70%) and testing observations (30%). Given the potential differences between formats, content, and outlets’ social media strategies across platforms, we used separate classifiers for each platform (see Appendix A).

Following the literature on automated clickbait classification (Biyani et al., 2016; Chakraborty et al., 2016; Rony et al., 2017), we considered the following predictors to train the models. First, we created count variables for each 1-, 2-, and 3-gram in the messages. For that purpose, we removed punctuation and numeric characters and tokenized the texts. We did not remove stop words and refrained from stemming words to their root because differences between clickbait and non-clickbait could be subtle. The messages consist of ca. 540,000 n-grams, all of which were considered as potential predictors of clickbait. In addition, we constructed various variables based on quantifiable characteristics of the messages, including the number of words, the mean word length, as well as the number of periods, question marks, exclamation marks, colons, and numerical characters. We also constructed binary variables indicating if a message starts with a number or a 5W1H word (i.e., what, why, when, who, which, and how), contains a third person personal pronoun (i.e., he, she, his, her, him), or includes a demonstrative (this, that, those, these). Finally, we evaluated the predictive power of meta-level characteristics, including outlet, day of the week, day of the month and hour of publication, and whether the post/tweet included a link embedded in the text.

Given the large number of potential predictors, we pre-screened the variables based on their bivariate correlation with the observed clickbait status of a post or tweet to speed up training. We retained the 500 predictors with the largest correlation coefficients to train various classification models, including (boosted) logistic regression, (penalized) discriminant analysis,

generalized linear models, gradient boosting, and random forests. All models were trained using 20-fold cross validation. We used penalized discriminant analysis, because this approach outperformed the other approaches.ⁱ We also compared the performance for a 3-point (i.e., “yes”, “maybe”, “no”) versus a 2-point clickbait scale (i.e., “yes” and “no”; combining initial “yes” and “maybe” into one category). We decided to proceed with the 2-point scale because the coding instructions induced conservative ratings and the classification models performed better.

Analysis

Data preparation.

Model-based classification does not involve a “hard” assessment of whether a headline is clickbait or not. Instead, for each post and tweet in our dataset, we obtained the *probability* that the corresponding headline classifies as clickbait.ⁱⁱ For the regression analyses, the individual probabilities were then aggregated at the outlet-month level. We created two panel datasets (one for each platform), consisting of 37 outlets and 54 months. We follow Zhang et al. (2020) and use a monthly frequency, because quarterly or yearly observations would imply an unnecessary loss of information, whereas a higher frequency (e.g., weekly or daily) would make the analyses vulnerable to the influence of posts/tweets with extreme interaction metrics.

Estimation.

To test the hypotheses, we estimated regression models of the following form:

$$cb_{i,t} = \beta_1 + \beta_2 interaction_{i,t-1}^{cb} + \beta_3 above_{i,t-1} + \beta_4 T_t + \beta_5 D_{i,t}^{algo} + \beta_6 D_{i,t}^{algo} \times T_t + \beta_7 D_{i,t}^{algo} \times above_{i,t-1} + \beta_8 cb_{i,t-1} + u_i + \varepsilon_{i,t} \quad (1)$$

where cb denotes the average clickbait probability of outlet i in month t . User $interaction^{cb}$ is defined as the sum of all likes and shares/retweets of posts/tweets with a clickbait probability > 0.5 . The binary variable $above$ takes the value 1 if the outlet’s average clickbait probability in the previous month was above the median clickbait probability (over all outlets) in that month. Using the lagged values $t - 1$ of both variables allows us to evaluate the clickbait response of social media editors to the behavior of users and competitors, rather than the other way around. We also included dummy variables to capture the algorithm changes of Facebook

and Twitter ($D_{i,t}^{algo}$) and a linear time trend (T_t). The interaction terms $D^{algo} \times T$ and $D^{algo} \times above$, respectively, can be used to verify if the clickbait probability and outlets' tendency to adjust to the industry standard varied across different algorithm phases.

We estimated Equation (1) as a fixed effects panel model via ordinary least squares, using u_i to denote the outlet fixed effect. As shown below, some outlets were generally more likely to provide clickbait than others, presumably due to different audiences, business models, and social media strategies. The fixed effects model makes it possible to account for those factors and any time-invariant differences in clickbait supply across outlets. Equation (1) also includes the outlet's clickbait probability in the previous month ($cb_{i,t-1}$), to control for routine behavior of social media editors.ⁱⁱⁱ

Results

Clickbait Probability

The average probability that a post or tweet is clickbait was 5.9% (SD = 4.6, Min = 0, Max = 33.5) on Facebook during our observation period and 2.8% (SD = 2.0, Min = 0, Max = 33.3) on Twitter. On average, the monthly like and share interaction with clickbait summed up to 4,712 likes and shares on Facebook and 133 on Twitter per outlet. On Facebook, the clickbait probability followed an upward trend, starting at around 5% in 2013 and reaching a peak of about 7% towards the end of 2016. Afterwards, the use of clickbait seems to have stagnated or even declined (Figure 1). The clickbait probability on Twitter appears quite stable over time (Figure 2). On Facebook, news outlets ranged in their average clickbait probability between 2% (*Freie Presse*) and 28% (*Focus*); see Figure 3. On Twitter, the clickbait probability ranged between 1% and 3% for most outlets; but four outlets stand out, with probabilities around 6% (Figure 4).

[Figures 1 to 4 about here]

Reverse U-Shape

We assume a threshold of clickbait that news outlets can supply before users grow tired. Our data reveal a reverse U-shaped relationship between clickbait supply and subsequent user interaction (see Figures 5 and 6). That is, user interaction increases when the outlets provide more clickbait—but only until a certain point, after which more clickbait results in decreasing

interaction. We calculate this point at 21% clickbait probability on Facebook and 6% on Twitter.

[Figures 5 and 6 about here]

Scene 1: Game Invitation

To evaluate if clickbait supply increases in algorithmic curation but not in a non-curative setting (H1), we compare the development of the clickbait probability on Twitter for the time before and after the platform rolled out its algorithm. The raw probability plotted in Figure 2 did not change visibly when Twitter started testing its algorithm in 2015, and there does not seem to be any increase after the global roll-out in 2016 either. To evaluate H1 formally, we consult the regression results in Table 2, Column (2), especially the coefficients on the interaction terms between the algorithm dummies and the time trend variable. The coefficient on *Jan 2015 to Feb 2016 × time trend* captures the development of the clickbait probability after the start of Twitter’s testing phase (relative to the pre-algorithm period from Jan 2013 to Dec 2014), whereas *After Feb 2016 × time trend* measures the development after the global roll-out. As the estimates show, neither the testing phase nor the implementation phase were associated with increases in clickbait probability. In contrast, we find that this probability decreased during the testing phase between Jan 2015 and Feb 2016 – the corresponding coefficient is negative and significant at the 5% level. In addition, the *Time trend* variable – which captures the development between Jan 2013 to Dec 2014 – indicates that clickbait had been growing in a non-curative setting, before Twitter started testing its algorithm. Hence, we reject H1.

Scene 2: Stigmergy

H2 is evaluated by consulting the coefficient on *Above industry level, t-1* in Table 2. For Facebook, we obtain a negative significant value (Column 1). The estimate thus implies that those outlets that have a higher clickbait probability than the median outlet adjust their clickbait supply downwards in the following month. The coefficient also implies that outlets below the industry level increase their clickbait supply. Thus, the monthly adjustment process can be described as one of convergence between outlets. For Twitter, the coefficient is also negative,

but not significant at conventional levels. Hence, we confirm H2 for Facebook but reject it for Twitter.

Scene 3: Playing Against a Threshold

Results in Table 2 confirm that higher amounts of past user interaction with clickbait are associated with greater levels of current clickbait supply (H3). On Facebook, this relationship is statistically significant at the 5% level. According to the point estimate of 0.0141, a one-unit increase in *Combined interaction with clickbait* (= 100,000 additional likes/shares) results in an increase in clickbait probability by 1.4 percentage points in the next month. To put this effect size in perspective, the combined interaction with clickbait on Facebook averaged 4,712 likes/shares per month. If this number was doubled, the outlets' clickbait probability would increase by $4,712 \times 1.4 / 100,000 = 0.066$ percentage points. Thus, the outlets reacted rather carefully to changes in user interaction. We also obtain a positive and significant coefficient for Twitter. The point estimate of 0.4321 suggests that 100,000 additional likes/shares lead to subsequent increase in clickbait probability by 43 percentage points. A change in the magnitude of the mean value of user interaction with clickbait on Twitter affects the clickbait probability by $133 \times 43.21 / 100,000 = 0.058$ percentage points, which suggests that social media editors on Twitter adjust their clickbait supply in very small steps, too. We accept H3.

Scene 4: Interfering in the Game

To answer RQ1, we first test if Facebook's modifications of the news feed algorithm to curb the proliferation of clickbait worked as intended. The interaction terms *Aug 2014 to Jul 2016* \times *time trend* and *After Jul 2016* \times *time trend* allow us to compare the development of clickbait supply over time after each modification of the algorithm in comparison to the period from Jan 2013 to Jul 2014. Accordingly, the time after the platform's first algorithm update (Aug 2014 – Jul 2016) was not characterized by a different trend in clickbait supply. However, the coefficient on *After Jul 2016* \times *time trend* is negative and significant at the 1% level, which implies that Facebook's second modification initiated a downward trend in clickbait. Yet, this development did not (immediately) result in a lower level of clickbait. As the coefficient on the

After Jul 2016 dummy indicates, the overall clickbait supply remained on a higher level than between Jan 2013 to Jul 2014, despite the initiation of a downward trend.

Second, the interaction terms *Aug 2014 to Jul 2016* × *above industry* and *After Jul 2016* × *above industry* have a significantly positive sign. Thus, after Facebook's modification, the news outlets were less inclined to revert to the industry standard, compared to the time before the platform tried to lessen the distribution of clickbait. Hence, the algorithm updates also reduced the outlets' tendency to align their behavior with their competitors.

[Table 2 about here]

Discussion

This study explores the clickbait posting behavior of legacy news outlets in the German news industry. Overall, we observe a low level of clickbait during the observation period (about 6% on Facebook and 3% on Twitter). The variations in the use of clickbait across platforms as well as between news outlets could be driven by differences in audiences' news consumption preferences. That is, Twitter's information-elitist audience of "politicians, journalists, and news lovers" (Cornia et al., 2018: 7) could be generally less attracted to clickbait because of a greater demand for quality news. Since social media use is related to an incidental news consumption during leisure time (Boczkowski et al., 2018), consuming sensational news on Facebook may be more valued than on Twitter. Moreover, the outlets with the highest clickbait probabilities in our sample tend to be tabloids (e.g., Express, B.Z. Berlin, and Hamburger Morgenpost), whose readers may have a greater preference for sensational content.

Beyond these structural differences, news outlets appear to play a stigmergic revision game for finding the optimal amount of clickbait. Clickbait supply reveals characteristics of collective behavior that is coordinated via environmental cues. With clickbait headlines, news outlets leave each other communicative trails that explain an industry-wide diffusion of clickbait as shown in previous research (Blom and Hansen, 2015; Rony et al., 2017). News outlets also adopt towards an industry-wide clickbait supply standard as predicted by stigmergy games (Vrancx et al., 2010).

However, although platforms represent a powerful force for journalism (Caplan and boyd, 2018; Nielsen and Ganter, 2017), we cannot prove that introducing algorithmic curation

increases clickbait—at least not on Twitter. In contrast, the clickbait game was already being played before Twitter introduced algorithmic curation, as we find a slight upward trend in clickbait during that period. Thus, we observe a stigmergic tendency towards interaction-cumulating content within the news industry across platforms, which may reflect a universal appeal of clickbait in an environment of competing for fleeting user attention. In addition, our results show that Facebook’s attempts to reduce clickbait were not initially successful. The platform’s first modification of its algorithm used reading time and interaction with a piece of content and thus focused on indicators for user regret. This approach may not have been sophisticated enough to identify clickbait and outlets may not have experienced lower visibility. Only the second update in 2016, in which Facebook used a supervised approach with humans identifying clickbait to train a classifier, resulted in a sustainable decrease in clickbait growth. That is, news outlets adjusted to this algorithm change but did not reduce clickbait to the lower 2013 level. Hence, the universal appeal of clickbait represents a sustainable force in the news industry whose effect is hardly reversible.

Still, regulative intervention has an indirect effect on news outlets’ clickbait supply. Outlets’ tendency to adjust to the average industry clickbait supply declined after both algorithm changes of Facebook. That is, the news outlets’ clickbait supply deviated from each other, which indicates a replacement of stigmergic orientation with opportunistic behavior. This result complements the findings of Meese and Hurcombe (2020) and Cornia et al. (2018), who show that algorithm changes induced outlets to turn away from Facebook as a central distributor. Thus, algorithm changes may have unintended consequences leading to a decrease of news homogeneity online.

User behavior determines clickbait supply as well, and perhaps more so than algorithm changes. We show that likes and shares of clickbait predict subsequent clickbait supply on Facebook and Twitter. However, the outlets’ reactions to user interaction are relatively small, which implies cautious supply adjustment in every step. News outlets continue their upward search for a threshold as long as the user interaction with clickbait content grows—or else they decrease their supply accordingly, indicating a careful ride of the waves of interaction.

Regarding a fatigue threshold, our results show that a moderate level of clickbait performs better than low or high levels. This result is in line with previous research pointing at an inverted *U*-shaped relationship between clickbait and user interaction (Zhang et al., 2020) and suggesting a wear-out effect of these headlines (Molyneux and Coddington, 2020). We specify a wear-out threshold of clickbait at 21% probability on Facebook and 6% on Twitter over all observed news outlets, with a flat increase and a steeper decrease of clickbait performance for Twitter. The lower threshold, flat increase, and steeper decrease on Twitter may be due to an information-elitist audience whose experienced utility of clickbait news barely exceeds their expected utility but is easily disappointed. Moreover, the steepness of the interaction decline on Twitter may be fueled by decreasing popularity cues signaling low utility to users and lowering stigmergic user behavior as suggested by Dipple et al. (2012).

Regarding limitations, first, our data capture social media news and thus ignore the larger news environment that publishers consider. Future research on clickbait provision outside social platforms could reveal whether clickbait supply collectively changes as found here. Also, our data do not include non-journalistic content suppliers, which are however relevant as clickbait fatigue is related to users' total clickbait exposure. Second, our sample period ends in 2017 due to Facebook's restrictions, not including Facebook's 2018 algorithm changes that broadly address news content. More recent data could also reveal whether legacy news outlets abstain from using clickbait since the rise of hoax outlets, as predicted by reputation games. Third, our investigation is limited to legacy media in Germany, implying a conservative context with relatively low online advertising revenue growth and modest social media news use. Future research could evaluate if our results hold in countries with more strongly growing online revenues and larger social media audiences, possibly exhibiting greater responsiveness to platforms' regulative intentions. Fourth, likes and shares are certainly of interest to social media editors—although these popularity cues cannot be equated with engagement with journalism (Haim et al., 2018). User interaction in the form of referral traffic might have an even larger impact on media outlets' decisions because this traffic can be directly monetized. Sismeiro and Mahmood (2018) show that social media interaction with news stories is a strong predictor of news outlets' website traffic, which gives us confidence that the likes

and shares considered here are relevant metrics. Fifth, supervised machine learning makes it possible to detect clickbait in a large corpus of text. The automated classification may not reach the same accuracy as human coding though. As shown in Appendix A, our clickbait classifiers achieved a satisfactory performance—especially in comparison to previous clickbait detection studies (Biyani et al., 2016; Chakraborty et al., 2016; Potthast et al., 2018)—but methodological advances are desirable to improve the accuracy of automated text classification. Lastly, as our data include only the editorial output of the news organizations under study, it remains unclear to what extent players consciously follow an established strategy. While audience metrics seem to gamify journalism to some extent (Ferrer-Conill, 2017), follow-up research is required to understand whether social media editors make decisions with a revision game on an unknown threshold in mind.

Conclusions

We explore the clickbait supply trends and patterns of legacy news outlets in an altering algorithmic-curation platform environment with utility-seeking users. This study is the first to determine the relation between news outlets and platforms by observing collective long-term news supply behavior during several curation algorithm changes.

Legacy outlets collectively supply interaction-inducing content on social media platforms following a revision game. They tend to underuse clickbait, except for a few heavy-using outlets with a tabloid focus, suggesting that professional norms endure on social media. An adjustment towards the optimal clickbait threshold takes place in small steps, which implies uncertainty despite the availability of detailed audience metrics. This uncertainty may lead to sharing audience metrics across outlets—because scale effects allow to better exploit audience metrics and thus reduce decision uncertainty more efficiently (Claussen et al., 2019).

Results endorse an increasing audience focus in online environments. Audience interaction serves as utility cues for users, revision cues for news outlets, and loyalty cues for platforms. In line with Myllylahti (2020), present results indicate that news outlets embrace a universal and almost irreversible focus on audience attention. In a competitive environment, even legacy news outlets become “clickbait media” (Munger, 2020) that design “addictive distractions” (Naughton, 2018: 387), at least on a low level.

Lastly, our findings contribute to the understanding of the relationship between social platforms and legacy media, especially concerning implications of platform policy for editorial decisions. News outlets can adhere to (Garz and Szucs, 2021; Tandoc and Maitra, 2018) and, as in our study, be indirectly responsive to changes in algorithmic curation. Both reactions acknowledge the central role of platforms not only in news distribution but also for news supply. Hence, curation algorithms serve as bureaucratic mechanisms (Caplan and boyd, 2018) that can enforce news similarity (Boczkowski, 2010; Haim et al., 2021). Our results advance these findings suggesting that algorithm changes reverse such heterogeneity tendencies and unintentionally may increase news diversity online.

References

- Andina-Díaz A and García-Martínez JA (2020) Reputation and news suppression in the media industry. *Games and Economic Behavior* 123: 240–271.
- Andrew BC (2007) Media-generated shortcuts: Do newspaper headlines present another roadblock for low-information rationality? *Harvard International Journal of Press/Politics* 12(2): 24–43.
- Arbaoui B, Swert K de and van der Brug W (2020) Sensationalism in News Coverage: A Comparative Study in 14 Television Systems. *Communication Research* 47(2): 299–320.
- Biyani P, Tsioutsoulouklis K and Blackmer J (eds) (2016) "8 Amazing Secrets for Getting More Clicks": *Detecting Clickbaits in News Streams Using Article Informality*.
- Blom JN and Hansen KR (2015) Click bait: Forward-reference as lure in online news headlines. *Journal of Pragmatics* 76: 87–100.
- Boczkowski PJ (2010) *News at Work: Imitation in an Age of Information Abundance*. University of Chicago Press.
- Boczkowski PJ, Mitchelstein E and Matassi M (2018) "News comes across when I'm in a moment of leisure": Understanding the practices of incidental news consumption on social media. *New Media & Society* 20(10): 3523–3539.
- Bourdieu P (1983) The field of cultural production, or: The economic world reversed. *Poetics* 12(4-5): 311–356.
- Braun JA and Eklund JL (2019) Fake News, Real Money: Ad Tech Platforms, Profit-Driven Hoaxes, and the Business of Journalism. *Digital Journalism* 7(1): 1–21.
- Cagé J, Hervé N and Mazoyer B (2020) *Social Media and Newsroom Production Decisions*. Working Paper.
- Caplan R and boyd d (2018) Isomorphism through algorithms: Institutional dependencies in the case of Facebook. *Big Data & Society*. DOI: 10.1177/2053951718757253.
- Chakraborty A, Paranjape B, Kakarla S, et al. (2016) Stop Clickbait: Detecting and preventing clickbaits in online news media. In: *2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*: San Francisco, CA, USA, 18 - 21 August 2016, pp. 9–16. IEEE.

CLICKBAIT NEWS

- Chen VY and Pain P (2019) News on Facebook: How Facebook and Newspapers Build Mutual Brand Loyalty Through Audience Engagement. *Journalism & Mass Communication Quarterly*. DOI: 10.1177/1077699019876634.
- Chen Y (2020) A revision game of experimentation on a common threshold. *Journal of Economic Theory* 186: 104997.
- Cimino MG, Lazzeri A, Pedrycz W, et al. (2018) Using Stigmergy to Distinguish Event-Specific Topics in Social Discussions. *Sensors* 18(7): 2117.
- Claussen J, Peukert C and Sen A (2019) *The Editor vs. the Algorithm: Returns to Data and Externalities in Online News*(8012). Available at: https://ideas.repec.org/p/ces/ceswps/_8012.html.
- Cornia A, Sehl A, Levy DAL, et al. (2018) Private Sector News, Social Media Distribution, and Algorithm Change. Available at: https://reutersinstitute.politics.ox.ac.uk/sites/default/files/2018-10/Cornia_Private_Sector_News_FINAL.pdf.
- DiMaggio PJ and Powell WW (1983) The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. *American Sociological Review* 48(2): 147–160.
- Dipple A, Raymond K and Docherty M (2012) Stigmergy in web 2.0. In: *Proceedings of the 3rd Annual ACM Web Science Conference on - WebSci '12*: (eds N Contractor, B Uzzi, M Macy and W Nejdl), Evanston, Illinois, 22 - 24 June 2012, pp. 86–94. New York, New York, USA: ACM Press.
- dos Santos MA, Lycarião D and Aquino JA de (2018) The virtuous cycle of news sharing on Facebook: Effects of platform affordances and journalistic routines on news sharing. *New Media & Society* 21(2): 398–418.
- El-Arini K and Tang J (2014) Click-baiting. Available at: <https://about.fb.com/news/2014/08/news-feed-fyi-click-baiting/>.
- Etienne J (2013) Ambiguity and relational signals in regulator-regulatee relationships. *Regulation & Governance* 7(1): 30–47.

CLICKBAIT NEWS

- Ferrer-Conill R (2017) Quantifying Journalism? A Study on the Use of Data and Gamification to Motivate Journalists. *Television & New Media* 18(8): 706–720.
- Garz M and Szucs F (2021) *Algorithmic Selection and Ideological Diversity of News on Facebook*. Working Paper.
- Grassé P-P (1959) La reconstruction du nid et les coordinations interindividuelles chez *Bellicositermes natalensis* et *Cubitermes* sp. La théorie de la stigmergie: Essai d'interprétation du comportement des termites constructeurs. *Insectes sociaux* 6(1): 41–80.
- Hågvar YB (2019) News Media's Rhetoric on Facebook. *Journalism Practice* 13(7): 853–872.
- Haim M, Karlsson M, Ferrer-Conill R, et al. (2021) You Should Read This Study! It Investigates Scandinavian Social Media Logics ♪. *Digital Journalism*: 1–21.
- Haim M, Kümpel AS and Brosius H-B (2018) Popularity cues in online media: A review of conceptualizations, operationalizations, and general effects. *Studies in Communication / Media* 7(2): 186–207.
- Hanitzsch T, Hanusch F, Mellado C, et al. (2011) Mapping journalism cultures across nations: A comparative study of 18 countries. *Journalism Studies* 12(3): 273–293.
- Hansen MH, Hoskisson RE and Barney JB (2008) Competitive advantage in alliance governance: Resolving the opportunism minimization–gain maximization paradox. *Managerial and Decision Economics* 29(2-3): 191–208.
- Kamen M (2016) Twitter's non-chronological timeline is now opt-out. *WIRED*, 17 March.
- Kuhn M (2008) Building Predictive Models in R Using the caret Package. *Journal of Statistical Software* 28(5).
- Kuiken J, Schuth A, Spitters M, et al. (2017) Effective Headlines of Newspaper Articles in a Digital Environment. *Digital Journalism* 5(10): 1300–1314.
- Landis JR and Koch GG (1977) The Measurement of Observer Agreement for Categorical Data. *Biometrics* 33(1): 159.
- Lischka JA (2021) Logics in social media news making: How social media editors marry the Facebook logic with journalistic standards. *Journalism* 22(2): 430–447.
- Magin M, Steiner M, Häuptli A, et al. (2021) Is Facebook driving tabloidization. In: Conboy M and Eldridge SA (eds) *Global Tabloid: Culture and technology*. [S.l.]: Routledge.

CLICKBAIT NEWS

- Meese J and Hurcombe E (2020) Facebook, news media and platform dependency: The institutional impacts of news distribution on social platforms. *New Media & Society*. DOI: 10.1177/1461444820926472.
- Molyneux L and Coddington M (2020) Aggregation, Clickbait and Their Effect on Perceptions of Journalistic Credibility and Quality. *Journalism Practice* 14(4): 429–446.
- Mourão RR and Robertson CT (2019) Fake News as Discursive Integration: An Analysis of Sites That Publish False, Misleading, Hyperpartisan and Sensational Information. *Journalism Studies* 20(14): 2077–2095.
- Munger K (2020) All the News That's Fit to Click: The Economics of Clickbait Media. *Political Communication* 37(3): 376–397.
- Myllylahti M (2020) Paying Attention to Attention: A Conceptual Framework for Studying News Reader Revenue Models Related to Platforms. *Digital Journalism* 8(5): 567–575.
- Nash J (1951) Non-cooperative games. *Annals of mathematics*: 286–295.
- Naughton J (2018) Platform Power and Responsibility in the Attention Economy. In: Moore M and Tambini D (eds) *Digital Dominance: The Power of Google, Amazon, Facebook, and Apple*. Oxford: Oxford University Press, Incorporated, pp. 371–395.
- Neumann J von and Morgenstern O (1944) *Theory of games and economic behavior*. Princeton: Princeton University Press.
- Newman N (2016) Journalism, media, and technology: Trends and predictions 2016. Digital News Project. Available at: <https://reutersinstitute.politics.ox.ac.uk/sites/default/files/2017-06/Journalism%2C%20media%20and%20technology%20predictions%202016.pdf>. (accessed 11 January 2017).
- Newman N, Fletcher R, Kalogeropoulos A, et al. (2018) Reuters Institute Digital News Report 2018. Available at: media.digitalnewsreport.org/wp-content/uploads/2018/06/digital-news-report-2018.pdf?x89475 (accessed 22 June 2017).
- Nielsen RK and Ganter SA (2017) Dealing with digital intermediaries: A case study of the relations between publishers and platforms. *New Media & Society* 20(4): 1600–1617.
- Osborne MJ and Rubinstein A (1994) *A course in game theory*. MIT Press.

- Peterson-Salahuddin C and Diakopoulos N (2020) Negotiated Autonomy: The Role of Social Media Algorithms in Editorial Decision Making. *Media and Communication* 8(3): 27–38.
- Peysakhovich A and Hendrix K (2016) Further Reducing Clickbait in Feed. Available at: <https://about.fb.com/news/2016/08/news-feed-fyi-further-reducing-clickbait-in-feed/>.
- Potthast M, Gollub T, Komlossy K, et al. (2018) Crowdsourcing a large corpus of clickbait on twitter. In: *Proceedings of the 27th international conference on computational linguistics*: pp. 1498–1507.
- Rony MMU, Hassan N and Yousuf M (2017) Diving Deep into Clickbaits. In: *Proceedings of the 2017 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2017*: (eds J Diesner, E Ferrari and G Xu), Sydney Australia, 31 07 2017 03 08 2017, pp. 232–239. New York, NY, USA: ACM.
- Rosiana P (2015) While You Were Away. Available at: https://blog.twitter.com/official/en_us/a/2015/while-you-were-away-0.html.
- Scacco JM and Muddiman A (2020) The curiosity effect: Information seeking in the contemporary news environment. *New Media & Society* 22(3): 429–448.
- Schultz I (2007) The Journalistic Gut Feeling: Journalistic doxa, news habitus and orthodox news values. *Journalism Practice* 1(2): 190–207.
- Sismeiro C and Mahmood A (2018) Competitive vs. Complementary Effects in Online Social Networks and News Consumption: A Natural Experiment. *Management Science* 64(11): 5014–5037.
- Starmer C (2000) Developments in Non-Expected Utility Theory: The Hunt for a Descriptive Theory of Choice under Risk. *Journal of Economic Literature* 38(2): 332–382.
- Suri H (2011) Purposeful Sampling in Qualitative Research Synthesis. *Qualitative Research Journal* 11(2): 63–75.
- Tandoc EC (2018) Five ways BuzzFeed is preserving (or transforming) the journalistic field. *Journalism* 19(2): 200–216.
- Tandoc EC and Maitra J (2018) News organizations' use of Native Videos on Facebook: Tweaking the journalistic field one algorithm change at a time. *New Media & Society* 20(5): 1679–1696.

CLICKBAIT NEWS

Vrancx P, Verbeeck K and Nowé A (2010) Analyzing the dynamics of stigmergetic interactions through pheromone games. *Theoretical Computer Science* 411(21): 2116–2126.

Zhang W, Du W, Bian Y, et al. (2020) Seeing is not always believing: An exploratory study of clickbait in WeChat. *Internet Research* 30(3): 1043–1058.

CLICKBAIT NEWS

Table 1: Examples of non-clickbait and clickbait headlines

| Non-clickbait | Clickbait |
|---|---|
| <i>Frau Merkel denkt noch lange nicht ans Aufhören</i> (Ms. Merkel is nowhere near considering leaving) | <i>Traurig – und jetzt auch noch unfassbar tragisch!</i> (Sad – and now also unbelievably tragic!) |
| <i>Polizei sucht weiter nach den Buntmetall-Dieben</i> (Police continues to look for the non-ferrous metal thieves) | <i>Dieser Trick ist goldwert! Ihr solltet ihn unbedingt kennen</i> (This trick is worth a mint! You definitely need to know it) |
| <i>Opernhaus des Jahres steht in Köln</i> (Opera house of the year is in Cologne) | <i>Diese 10 Apps helfen euch beim Sparen</i> (These 10 apps will save you money) |

CLICKBAIT NEWS

Table 2: Determinants of clickbait supply

| | (1) Clickbait probability (Facebook) | (2) Clickbait probability (Twitter) |
|--|---|---|
| Avg. clickbait probability (t-1) | 0.6521*** (0.0977) | 0.1904** (0.0870) |
| Combined interaction with clickbait (t-1) | 0.0141** (0.0055) | 0.4321*** (0.1161) |
| Above industry level, t-1 (binary) | -0.0147*** (0.0025) | -0.0015 (0.0018) |
| Time trend | 0.0002** (0.0001) | 0.0002** (0.0001) |
| <i>Facebook algorithm changes (reference category: Jan 2013 to Jul 2014)</i> | | |
| -Aug 2014 to Jul 2016 | 0.0655 (0.0698) | |
| -Aug 2014 to Jul 2016 × time trend | -0.0001 (0.0001) | |
| -Aug 2014 to Jul 2016 × above industry (t-1) | 0.0100*** (0.0018) | |
| -After Jul 2016 | 0.3153*** (0.1016) | |
| -After Jul 2016 × time trend | -0.0005*** (0.0002) | |
| -After Jul 2016 × above industry (t-1) | 0.0104*** (0.0019) | |
| <i>Twitter algorithm changes (reference category: Jan 2013 to Dec 2014)</i> | | |
| -Jan 2015 to Feb 2016 | | 0.3474* (0.1730) |
| -Jan 2015 to Feb 2016 × time trend | | -0.0005* (0.0003) |
| -Jan 2015 to Feb 2016 × above industry (t-1) | | 0.0012 (0.0018) |
| -After Feb 2016 | | 0.1724 (0.1255) |
| -After Feb 2016 × time trend | | -0.0003 (0.0002) |
| -After Feb 2016 × above industry (t-1) | | 0.0035 (0.0023) |
| Constant | -0.1271* (0.0672) | -0.1287* (0.0678) |
| Adj. R ² | 0.5890 | 0.0661 |

Notes: Estimates based on fixed effects panel data models. Robust standard errors (clustered by outlet) in parentheses.

* p<0.10, ** p<0.05, *** p<0.01

CLICKBAIT NEWS

Fig. 1: Use of clickbait over time, Facebook




Fig. 2: Use of clickbait over time, Twitter




Fig. 3: Use of clickbait across outlets, Facebook

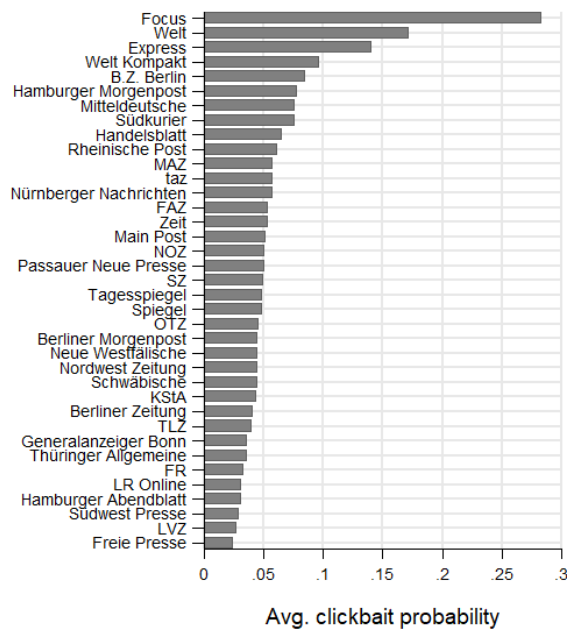


Fig. 4: Use of clickbait across outlets, Twitter

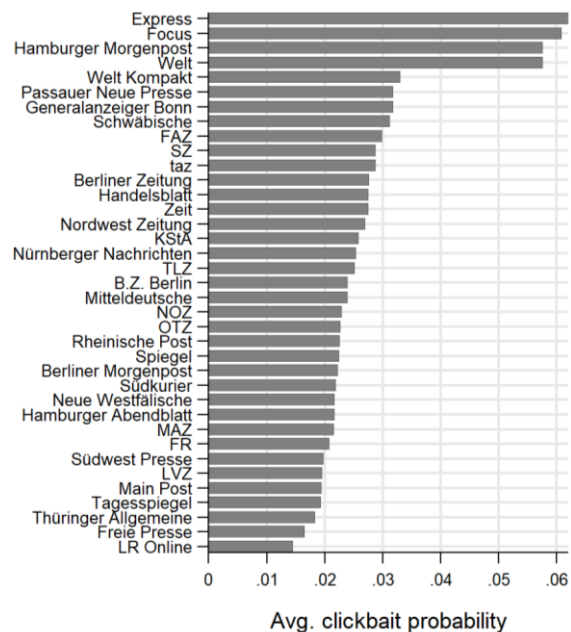


Fig. 5: Clickbait and Facebook user interaction

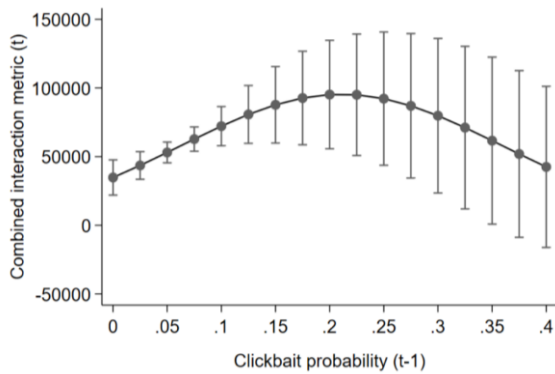



Fig. 6: Clickbait and Twitter user interaction



Notes: The figures show an inverse *U*-shaped relationship between the average clickbait probability and model-based predictions of user interaction with the postings of the news outlets in our sample. The predictions were obtained by fitting models of the form: $interaction_{i,t} = \beta_1 + \beta_2 cb_{i,t-1} + \beta_3 cb_{i,t-1}^2 + \beta_4 X_{i,t} + \varepsilon_{i,t}$, where *interaction* captures the sum of likes and shares/retweets of outlet *i* in month *t*, *cb* denotes the (squared) clickbait probability in the previous month, and the variable vector *X* includes outlet dummies, the number of worldwide monthly active users, the total number of posts/tweets, and a time trend polynomial of order 3. The estimated models – detailed results are available on request – suggest a turning point of 0.211 clickbait probability on Facebook and 0.063 on Twitter, using the formula $-\beta_2/(2 \times \beta_3)$. The error bars represent the 95% confidence interval.

Appendix A: Classifier comparison

| Classifier | Accuracy | Balanced accuracy | Recall | Precision |
|---------------------------------|----------|-------------------|--------|-----------|
| Facebook | | | | |
| Boosted logistic regression | 0.885 | 0.621 | 0.942 | 0.934 |
| Generalized linear model | 0.843 | 0.621 | 0.918 | 0.935 |
| Gradient boosting | 0.901 | 0.606 | 0.961 | 0.931 |
| Penalized discriminant analysis | 0.911 | 0.638 | 0.968 | 0.936 |
| Random forest | 0.905 | 0.588 | 0.971 | 0.927 |
| Twitter | | | | |
| Boosted logistic regression | 0.970 | 0.505 | 0.998 | 0.972 |
| Generalized linear model | 0.955 | 0.590 | 0.712 | 0.976 |
| Gradient boosting | 0.966 | 0.528 | 0.995 | 0.973 |
| Penalized discriminant analysis | 0.955 | 0.600 | 0.976 | 0.977 |
| Random forest | 0.963 | 0.538 | 0.987 | 0.973 |

ⁱ We evaluate the models' performance by comparing the clickbait predictions obtained for the hold-out sample with its actual values. Specifically, we selected the approach where standard performance measures computed by the caret package—including accuracy, balanced accuracy, recall, and precision—yield the highest values (see Appendix A).

ⁱⁱ It would also be possible to use a binary clickbait classification (yes/no) by distinguishing between headlines above and below a probability of 0.5, but this approach involves an unnecessary loss of detail.

ⁱⁱⁱ Fixed effects panel data estimates can be biased if the model includes lags of the dependent variable (here $cb_{i,t-1}$). However, this bias disappears if the time dimension of the panel is large enough, as in our case ($T = 54$ months).