

Opinion Market Model: Stemming Far-Right Opinion Spread Using Positive Interventions

Pio Calderon, Rohit Ram, Marian-Andrei Rizoiu

University of Technology Sydney

piogabrielle.b.calderon@student.uts.edu.au, rohit.ram@student.uts.edu.au, marian-andrei.rizoiu@uts.edu.au

Abstract

Online extremism has severe societal consequences, including normalizing hate speech, user radicalization, and increased social divisions. Various mitigation strategies have been explored to address these consequences. One such strategy uses positive interventions: controlled signals that add attention to the opinion ecosystem to boost certain opinions. To evaluate the effectiveness of positive interventions, we introduce the Opinion Market Model (OMM), a two-tier online opinion ecosystem model that considers both inter-opinion interactions and the role of positive interventions. The size of the opinion attention market is modeled in the first tier using the multivariate discrete-time Hawkes process; in the second tier, opinions cooperate and compete for market share, given limited attention using the market share attraction model. We demonstrate the convergence of our proposed estimation scheme on a synthetic dataset. Next, we test OMM on two learning tasks, applying to two real-world datasets to predict attention market shares and uncover latent relationships between online items. The first dataset comprises Facebook and Twitter discussions containing moderate and far-right opinions about bushfires and climate change. The second dataset captures popular VEVO artists' YouTube and Twitter attention volumes. OMM outperforms the state-of-the-art predictive models on both datasets and captures latent cooperation-competition relations. We uncover (1) self- and cross-reinforcement between far-right and moderate opinions on the bushfires and (2) pairwise artist relations that correlate with real-world interactions such as collaborations and long-lasting feuds. Lastly, we use OMM as a testbed for positive interventions and show how media coverage modulates the spread of far-right opinions.

1 Introduction

Online social media platforms are fertile grounds for deliberation and opinion formation (Gupta, Jain, and Tiwari 2022). Opinions thrive in the *online opinion ecosystem*, interconnected online social platforms where they interact – cooperating or competing for the finite public attention (Wu, Rizoiu, and Xie 2019).

We delineate two types of interventions to mitigate the spread of extremist views. *Negative interventions* aim to subtract attention from the opinion ecosystem by placing

fact-check warnings on postings (Nekmat 2020), shadow-banning (Young 2022) or outright banning extremist social media groups and accounts (Jackson 2019). While negative interventions are effective (Clayton et al. 2020), they are available solely to the social media platforms that tend to use them sparingly (Porter and Wood 2021).

Positive interventions (GIFCT 2021), such as misinformation debunking (Haciyakupoglu et al. 2018; Shu, Wang, and Liu 2019) and increasing media coverage (Horowitz et al. 2022), mitigate extremist views by adding attention to the online opinion ecosystem through informing the public, redistributing attention away from extremist, and toward moderate views. Such interventions are typically in the hands of government and media agencies (Radsch 2016). Testing the viability of positive interventions requires capturing reactions to interventions and inter-opinion interactions.

This work develops a model for the dynamics of the opinion ecosystem and a testbed for evaluating positive interventions. We focus on two open questions. The first question explores the analogy between opinions and economic goods. In a competitive economic market of limited resources, co-existing goods can interact in one of two ways: either they compete for market share (*substitute* brands, like Pepsi and Coke) or reinforce each other (*complementary* items, like bread and butter). We argue that opinions in the online ecosystem behave similarly, allowing us to leverage market share modeling tools (Cooper 1993). The first research question is: **Can we model the online opinion ecosystem as an environment where opinions cooperate or compete for market share?** We propose the Opinion Market Model¹ (OMM), a two-tier model to address this question. Fig. 1 showcases a simple opinion ecosystem under intervention, with two opinions (denoted 0 and 1) on a single social media platform. Each opinion has two polarities: far-right supporters (+) and moderate debunkers (-). Exogenous signals (shown in gray in the top panel of Fig. 1) and interventions (shown in yellow) modulate the dynamics of the opinions' sizes. Exogenous signals are naturally occurring events like bushfires, floods, or political speeches. Interventions (like increased media coverage) are designed to add attention to the opinion ecosystem, increasing the market share of cer-

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹The code and datasets are available at <https://github.com/behavioral-ds/opinion-market-model>.

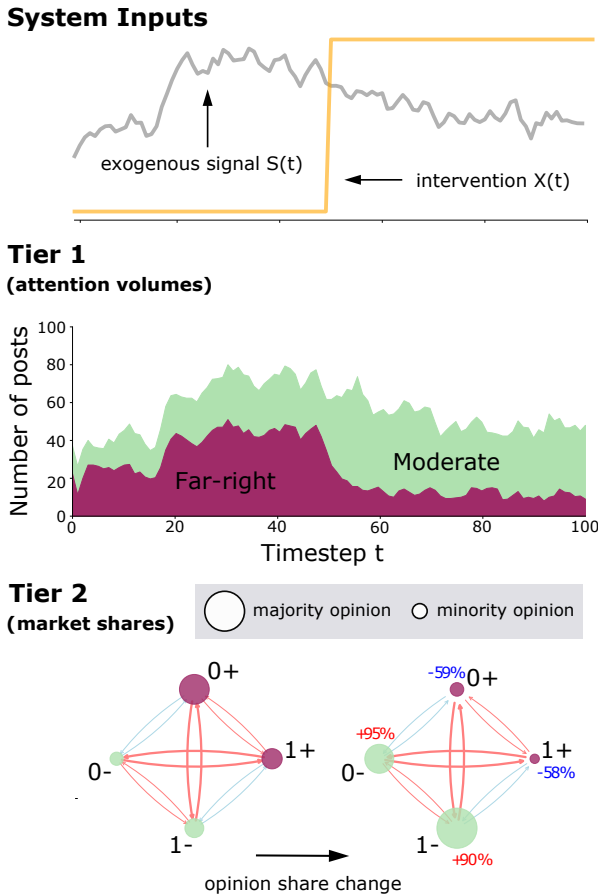


Figure 1: We illustrate how the positive intervention $X(t)$ (in Eq. (9)) suppresses far-right opinions on a simulated toy opinion ecosystem with two far-right ($0+$, $1+$) and two moderate ($0-$, $1-$) opinions. For instance, $0+$ and $1+$ can represent the opinions “Greens policies caused the Australian bushfires” and “mainstream media cannot be trusted,” respectively; $0-$ and $1-$ can be obtained as their negations. Top row: the exogenous signal $S(t)$ (in Eq. (5)) and the intervention $X(t)$. Middle row: total daily opinion market size from our model’s first tier, split into far-right (+) and moderate (-) opinion volumes. Bottom row: market shares and the interactions between the four opinions from our model’s second tier. Nodes are opinions; their sizes indicate market share; edges show exciting (red) and inhibiting (blue) relations. $X(t)$ suppresses far-right opinions for $t > 50$. Shown are average market shares before (left) and after (right) $t = 50$.

tain opinions while suppressing others. The first tier of OMM (middle row in Fig. 1) uses a discrete-time Hawkes process to estimate the size of the opinion attention market – that is, the daily number of postings featuring opinions. The Hawkes process has been widely used to model online attention (Rizoiu et al. 2017; Zarezade et al. 2017) due to its ability to account for exogenous factors and the endogenous “word-of-mouth” through its self- and cross-exciting property. The second tier of OMM (bottom row in Fig. 1) leverages a market share attraction model to capture opinion in-

teractions – we assume that online opinions compete for the users’ limited online attention (Weng et al. 2012; Gelper, van der Lans, and van Bruggen 2021). For the example in Fig. 1, opinions $0-$ and $1+$ have a strong reinforcing relation (shown as red arrows), while $1-$ and $1+$ have a weak competing relation (blue arrows).

We test OMM on two real-world datasets¹. The first contains Facebook and Twitter discussions expressing moderate and far-right opinions on bushfires and climate change (Kong et al. 2022). The second captures the YouTube views and Twitter mentions for the most popular VEVO artists’ songs in 2017 (Wu, Rizoiu, and Xie 2019). We evaluate OMM on two tasks: predicting attention market share and exposing relationships between online items. For the predictive task, OMM outperforms the current state of the art in market share modeling (Correlated Cascades (Zarezade et al. 2017) and Competing Products (Valera and Gomez-Rodriguez 2015)) and predictive baselines on both datasets. For the second task, we leverage the OMM to expose the relations between opinions on the two platforms. For the bushfire case study, no significant interactions occurred on Facebook, as postings were collected from far-right public groups with limited interaction with the opposing side. On Twitter, we observe self-reinforcement behavior of both far-right and moderate opinions, probably due to the *echo chamber effect* (Cinelli et al. 2021) – reinforcing one’s beliefs due to repeated interactions with users sharing similar ideologies on social platforms. Surprisingly, we notice that opposing views reinforce each other, probably due to the deliberative nature of Twitter, where far-right sympathizers and opponents oppose each other. For the VEVO artists case study, we uncover nontrivial pairwise interactions of music artists correlating with real-world relationships – such as Ariana Grande’s and Calvin Harris’ reinforcement relationship due to their collaboration “Heatstroke” and Taylor Swift’s and Justin Bieber’s inhibiting relationship.

Our second research question is: **Can we test the sensitivity of the opinion ecosystem to positive interventions?** OMM accounts for positive interventions – controlled external signals to boost certain opinions. In Fig. 1 an intervention is performed for $t > 50$, which suppresses the far-right opinions (+), leading to the shrinking of their market share. We use OMM for two tasks: first, to estimate whether interventions effectively shape the opinion ecosystem and, second, to construct what-if scenarios as synthetic testbeds for future interventions. For the bushfire case study, we test whether news coverage from reputable and controversial media outlets suppresses or aids the spread of far-right opinions. We fit OMM twice: with and without media coverage. We find a better fit with the intervention, suggesting that media coverage actively shapes the opinion ecosystem. We perform synthetic what-if experiments: we vary the level of media coverage, simulate the system and observe the effect on opinion market shares. On Facebook, reputable media coverage reduces the prevalence of far-right opinions. On Twitter, both reputable and controversial media coverage suppress far-right opinions. However, for some opinions (like “Mainstream media cannot be trusted”), reputable news backfires increasing far-right opinions share.

The main contributions of the work are as follows:

1. A novel two-tier model of the opinion ecosystem that allows studying opinion interactions through an economics-based cooperation-competition lens. We introduce simulation and estimation algorithms and study the convergence with synthetic tests.
2. A synthetic testbed to uncover interactions across sympathizers and opponents of far-right opinions and likely effects of positive interventions via media coverage.
3. A curated dataset of Twitter and Facebook discussions on bushfires/climate change.

Related Work. We focus the discussion of related work on models for cooperative-competitive interaction in a set of co-diffusing online items. These models need to be both *predictive* and *interpretable* (usually generative models). We have observed a lack of recent research in this area, with few works dating after 2017. Closely related to our proposal is the Correlated Cascades (CC) model (Zarezade et al. 2017), a variant of the multivariate Hawkes process to model product adoption across a set of competing products in a social network. It estimates the interaction parameter β , tuning the market cooperation or competition level. A limitation of CC is that all products share a single β value. This simplifies existing asymmetric relationships and assumes that all brands either cooperate or compete. OMM addresses this issue by fully modeling these asymmetric relationships. Another closely related work is the Competing Products (CP) model (Valera and Gomez-Rodriguez 2015), a multivariate Hawkes model for product adoption/use where the frequency of use is affected by the usage of other products. Limitations of the work are the absence of the assumption of limited attention and the possibility of negative intensities since competitive interactions are modeled as negative parameters. OMM avoids the weaknesses of CP by using a multiplicative model, thereby avoiding negative intensities and defining opinion shares as fractions of the total attention volume. The SLANT model (De et al. 2016) and the follow-up SLANT+ (Kulkarni et al. 2017) extend the CP model to differentiate between a user’s latent and expressed opinion and uses a similar Hawkes process to model the intensity. However, SLANT requires fine-grained network information for training, which is prohibitive considering that online platforms are becoming more stringent with fine-grained data access (Venturini and Rogers 2019). On the other hand, OMM requires only opinion counts for training.

Ethics of Opinion Moderation and Broader Perspectives. OMM is intended to model interactions between opinions and be used as a testbed for evaluating positive interventions for opinion moderation. As any tool, OMM is unaware of the intention of its user and, in theory, could be used by oppressive regimes to silence or manipulate the liberal opinions of their opponents (Radsch 2016). In addition, the fundamental value of freedom of speech for democratic societies implies that non-widely accepted opinions also have the right to be expressed. The scientific literature studies this ethical conundrum in the context of Countering Violent Extremism (CVE) initiatives (Betz 2016; Radsch 2016). When viewing OMM as an AI evaluation tool supporting CVE ini-

tatives (Ferguson 2016), these ethical issues can be mitigated using online CVE frameworks in liberal democracies (Henschke and Reed 2021). We argue that the implementing body is responsible for OMM’s ethical usage, and CVE regulations should be leveraged to mitigate malicious intent.

Causal Impact. OMM measures the effect of media coverage on the opinion market shares using a generative model to disentangle endogenous and exogenous factors from observational data, similar to (Rizoiu and Xie 2017; Fujita et al. 2018; Garetto, Leonardi, and Torrisi 2021). Our model works on aggregate observational data (i.e., opinion counts), and it does not prove the causal impact of media coverage on individual opinion formation (i.e., behavior change). We would require a pre-test/post-test control group design to achieve true causal links. Previous work (King, Schneer, and White 2017; Guess et al. 2021; Agovino, Carillo, and Spagnolo 2021) provides evidence of the interventional role of media coverage. In Section 8, we explore this further in a what-if experiment to demonstrate how the level of media coverage affects opinion market shares.

2 Preliminaries

We introduce two classes of models that form the foundation of our approach: (1) the discrete-time Hawkes process (Browning et al. 2021), a model of event counts that display self-exciting behavior, and (2) the market share attraction model (Cooper 1993), a marketing model that uncovers the latent competitive structure of brands and interaction with marketing instruments.

Discrete-time Hawkes Process

The discrete-time Hawkes Process (DTHP) (Browning et al. 2021) is the discrete-time analogue of the popular self-exciting Hawkes process (Hawkes 1971), where instead of modeling the occurrence of events given by $t \in \mathbb{R}^+$, we model the event count $N(t)$ on $[t - 1, t)$ for $t \in \mathbb{N}$.

The DTHP is characterized by the conditional intensity function $\lambda(t)$, defined as the expected number of events that occur at time t , conditioned on the history $H_{t-1} = \{N(s) | s < t\}$. For a DTHP, $\lambda(t)$ is given by

$$\lambda(t) = \mathbb{E}[N(t) | H_{t-1}] = \mu + \sum_{s < t} \alpha \cdot f(t-s) \cdot N(s), \quad (1)$$

where μ is the baseline count of events, α determines the level of self-excitation and is the expected number of events produced by a single event and $f(t)$ is the triggering kernel, which controls the influence of the past events on the present. We specify $f(t)$ with the geometric probability mass function $f(t) = \theta(1 - \theta)^{t-1}, t \in \mathbb{N}$, the discrete-time analogue of the exponential distribution (Browning et al. 2021). Given $\lambda(t)$, model specification is completed by specifying a probability mass function for the count $N(t)$. Following (Browning et al. 2021), we set $N(t) \sim \text{Poi}(\lambda(t))$.

Market Share Attraction Model

In marketing literature, *market share attraction models* (MSAMs) (Cooper 1993) model the competitive structure of a set of M brands in a product category, predict their

market shares, and evaluate how a set of marketing instruments affects resulting market shares. MSAMs assume that the market share s_i of brand $i \in \{1 \dots M\}$ is proportional to consumers' attraction \mathcal{A}_i towards brand i :

$$s_i = \frac{\mathcal{A}_i}{\sum_{j=1}^M \mathcal{A}_j} \in [0, 1]. \quad (2)$$

\mathcal{A}_i is typically modeled as a parametric function of a set of K marketing instruments $\{X_{ki}\}_{k=1}^K \in \mathbb{R}^K$, where X_{ki} gives the value of the k^{th} marketing instrument for brand i . We typically specify \mathcal{A}_i as

$$\mathcal{A}_i = \exp \left(\beta_i + \sum_{k=1}^K \sum_{j=1}^M \gamma_{kij} X_{kj} \right), \quad (3)$$

where β_i measures the inherent attraction of brand i and $\gamma_{kij} \in \mathbb{R}$ measures the effect of the value of the k^{th} marketing instrument for brand j on brand i 's attraction. Whether γ_{kij} is positive (negative) is indicative of the excitatory (inhibiting) relationship from brand j to brand i through marketing instrument X_{kj} .

MSAMs are interpreted via the model elasticity $e(s_i, X_{kj})$, the ratio of the percent change in the market share s_i given a percent change in the value of the k^{th} marketing instrument for brand j . For example, an elasticity of $e(s_i, X_{kj}) = 0.1$ means that a 1% increase in X_{kj} corresponds to a 0.1% increase in s_i . That is,

$$e(s_i, X_{kj}) = \frac{\partial s_i / s_i}{\partial X_{kj} / X_{kj}} = \frac{\partial s_i}{\partial X_{kj}} \cdot \frac{X_{kj}}{s_i}. \quad (4)$$

The elasticity $e(s_i, X_{kj})$ captures the overall effect of brand j 's marketing instrument X_{kj} on brand i 's market share s_i : both the *direct effect* of X_{kj} on s_i , controlled by γ_{kij} , and the *indirect effect* of X_{kj} on s_i through its effect on the attraction of other brands $\{j \neq i\}$.

3 The OMM Model

In this section, we develop a two-tier model of the opinion ecosystem. The first tier models the total size of the opinion attention market on multiple online platforms. The second tier models the market share of opinions on each platform. Next, we introduce a scheme for parameter estimation.

OMM consists of two tiers; the first tier, which we call the *opinion volume model*, tracks the size of the opinion attention market, while the second tier, the *opinion share model*, tracks the market shares of the different opinions. Table 1 summarises the notation for important variables in the OMM. The full table is available in the online appendix (Calderon, Ram, and Rizoiu 2024).

Opinion Volume Model. Suppose our opinion ecosystem consists of P social media platforms. The opinion volume model tracks the attention volume, i.e. the number of opinionated posts $N^p(t)$, on each platform $p \in \{1, \dots, P\}$ and time $t \in \mathbb{N}$. We model $\{N^p(t)\}_p$ as a P -dimensional DTHP (defined analogous to the multivariate Hawkes process (Hawkes 1971)) with conditional intensity $\{\lambda^p(t)\}_p$,

$$\lambda^p(t) = \mu^p \cdot S(t) + \sum_{q=1}^P \sum_{s < t} \alpha^{pq} \cdot f(t-s) \cdot N^q(s). \quad (5)$$

In contrast to Eq. (1), we use a time-varying exogenous signal $S(t)$, which accounts for the baseline volume of events of exogenous origin. The signal $S(t)$ accounts for natural tendencies and events (i.e., epidemics, elections) and typically cannot be controlled. We introduce a scaling term μ^p for each platform p such that $\mu^p \cdot S(t)$ represents the exogenous opinion count for platform p on time t .

Since online platforms are not siloed and have significant user overlap, we allow the P platforms to interact via intra- and inter-platform excitation. The parameter $\alpha^{pq} > 0$ sets the level of intra-platform (for $p = q$) and inter-platform (for $p \neq q$) excitation. Lastly, we set $N^p(t) \sim \text{Poi}(\lambda^p(t))$.

Opinion Share Model. With the attention volumes for each platform p estimated in the first tier, the second tier models the market share $s_i^p(t)$, calculated as the fraction of opinionated posts on platform p conveying opinion i . Given the limited attention market size, opinions compete for attention within each platform.

Suppose that there are M different opinion types. We set $N_i^p(t)$ to be the number of opinionated posts conveying opinion i on platform p on time t , and $\lambda_i^p(t)$ to be its conditional intensity. Using the notion of limited attention (Zaregade et al. 2017), we relate $\lambda_i^p(t)$ to $\lambda^p(t)$ in Eq. (5) by introducing the market share $s_i^p(t) \in [0, 1]$ as the fraction of opinion i posts on platform p . That is,

$$\lambda_i^p(t) = \lambda^p(t) \cdot s_i^p(t), \quad (6)$$

and $\sum_{i=1}^M s_i^p(t) = 1$.

Similar to Eq. (2), we model $s_i^p(t)$ with attraction $\mathcal{A}_i^p(t)$,

$$s_i^p(t) = \frac{\mathcal{A}_i^p(t)}{\sum_{j=1}^M \mathcal{A}_j^p(t)}. \quad (7)$$

Leveraging the MNL form in Eq. (3), we define attraction

$$\mathcal{A}_i^p(t) = \exp \mathcal{T}_i^p(t), \quad (8)$$

where $\mathcal{T}_i^p(t)$ consists of two parts, accounting for *interventions* and *endogenous* dynamics, and is described in detail below,

$$\mathcal{T}_i^p(t) = \underbrace{\sum_{k=1}^K \gamma_{ik}^p \cdot \bar{X}_k(t)}_{\text{interventions}} + \underbrace{\sum_{q=1}^P \sum_{j=1}^M \beta_{ij}^{pq} \cdot \lambda^q(t|j)}_{\text{endogenous}}, \quad (9)$$

$$\bar{X}_k(t) = \sum_{s < t} f(t-s) \cdot X_k(s), \quad \text{and}$$

$$\lambda^p(t|j) = \mu_j^p \cdot S(t) + \sum_{q=1}^P \sum_{s < t} \alpha^{pq} \cdot f(t-s) \cdot N_j^q(s), \quad (10)$$

where $\mu^p = \sum_{j=1}^M \mu_j^p$.

In the first term of Eq. (9), we introduce a set of K positive interventions $\{X_k(t)\}_k$ that modify the opinion market shares in the opinion ecosystem. The interventions $\{X_k(t)\}_k$ have a different to $S(t)$ in Eq. (5), as the latter modifies the attention market size. Parameter $\gamma_{ik}^p \in \mathbb{R}$ measures the direct effect of the k^{th} intervention on the market

Notation	Interpretation
P	number of social media platforms
M	number of opinion types
K	number of positive interventions
T	terminal time
<hr/>	
Variables	
$S(t)$	input signal, volume of exogenous events
$X_k(t)$	input signal, k^{th} positive intervention
$s_i^p(t)$	market share of opin. i on platform p at time t
$\lambda_i^p(t)$	conditional intensity of opinion i
$N_i^p(t)$	#posts with opinion i on platform p at time t
$e(s_i^p(t), \cdot)$	opinion share model elasticity
<hr/>	
Data	
$n_t^p/n_{i,t}^p$	#posts on platform p at time t / with opinion i
$s_{i,t}^p$	fraction of opin. i posts on platf. p at time t
<hr/>	
Parameters	
μ_j^p	exogenous scaling term for opin. j on platf. p
α^{pq}	excitation parameter for intra-platform ($p = q$) and inter-platform (for $p \neq q$) dynamics
θ	memory parameter, describing how fast an event is forgotten, $\theta \in [0, 1]$
γ_{ik}^p	direct effect of the k^{th} intervention on share of opinion i on platform p
β_{ij}^{pq}	direct effect that opinion j on platform q has on share of opinion i on platform p .

Table 1: Summary of important quantities and notations.

share of opinion i on platform p . If γ_{ik}^p is positive (negative), then $X_k(t)$ reinforces (inhibits) opinion i on platform p .

In the second term of Eq. (9) we model the contribution of endogenous dynamics on the attraction of opinion i . To represent the prevalence of opinion j on platform q , we make use of the conditional intensity $\lambda^p(t|j)$ in Eq. (10), which models the dynamics of opinion j independent of other opinions. Parameter $\beta_{ij}^{pq} \in \mathbb{R}$ captures the direct effect that opinion j on platform q has on the market share of opinion i on platform p . Similar to γ_{ik}^p , we allow β_{ij}^{pq} to be positive (negative), representing a reinforcing (inhibiting) relationship from opinion j to i on platform q and p , respectively.

Estimation. Over the observation period $t \in \{1, \dots, T\}$, assume that we observe the exogenous signal $S(t)$, the K interventions $\{X_k(t)\}_k$, and the number $n_{i,t}^p$ of posts conveying opinion i on platform p for each i and p . Our goal is to estimate the parameter set $\Theta = \{\mu_j^p, \alpha^{pq}, \theta, \gamma_{ik}^p, \beta_{ij}^{pq}\}$.

The structure of our two-tier model allows us to cast parameter estimation as a two-tier optimization problem. Let $\Theta_1 = \{\mu^p, \alpha^{pq}, \theta\}$. The key observation here is that the first-tier parameter set Θ_1 can be estimated using only the opinion volume model in Eq. (5), independent of the opinion share model in Eq. (9). By optimizing the likelihood $\mathcal{L}_1(\Theta_1 | \{n_t^p\}_{p,t})$ of the platform-level volumes $\{n_t^p\}_{p,t}$, we can obtain an estimate $\hat{\Theta}_1$ of Θ_1 .

The second-tier parameter set $\Theta_2 = \{\mu_j^p, \gamma_{ik}^p, \beta_{ij}^{pq}\}$ can be obtained by optimizing the likelihood $\mathcal{L}_2(\Theta_2 | \hat{\Theta}_1, \{n_{i,t}^p\}_{i,p,t})$ of the opinion volumes $\{n_{i,t}^p\}_{i,p,t}$, conditioned on our estimate of the first-tier parameters $\hat{\Theta}_1$. Our full estimated parameter set is given by $\hat{\Theta} = \hat{\Theta}_1 \cup \hat{\Theta}_2$. The technical details of the estimation and the derivation of the likelihoods $\mathcal{L}_1(\cdot)$ and $\mathcal{L}_2(\cdot)$ and gradients $\partial_{\Theta_1} \mathcal{L}_1(\cdot)$ and $\partial_{\Theta_2} \mathcal{L}_2(\cdot)$ are available in the online appendix (Calderon, Ram, and RizoIU 2024).

Simulation. Suppose we are given the opinion volume $n_{i,0}^p$ at time $t = 0$ for each platform p and opinion i , such that $n_t^p = \sum_i n_{i,t}^p$. A sample of $n_{i,t}^p$ from OMM can be obtained by calculating the conditional intensity $\lambda_i^p(t)$ using Eq. (6), and then sampling $n_{i,t}^p$ from $\text{Poi}(\lambda_i^p(t))$. We obtain $\{n_{i,t}^p\}_{i,p,t}$ by repeating these steps over $\{1, \dots, T\}$.

Numerical Considerations. To improve model fit in our real-world case studies, we implement three augmentations to the model and estimation method, outlined below and fully detailed in the online appendix. First, we modify the attraction $\mathcal{A}_i^p(t)$ in Eq. (7) to prevent numerical overflow/underflow. Second, we add a regularization term in the second-tier optimization problem in Section 3 to impose structural constraints on $\{\hat{\gamma}_{ik}^p\}$ and improve estimation. Third, we apply log-scaling on $\lambda^q(t|j)$ and standardize both $\lambda^q(t|j)$ and $X_k(s)$ in Eq. (9) to solve scaling issues.

Stability Assumption. We implicitly assume that the opinion attention market is stable over the timeframe of the analysis, in the sense that the parameters Θ governing the behavior of the process stay constant within the timeframe. In situations where this assumption is not expected to hold (e.g. extreme events) and parameters change within the timeframe, a change-point model extension (Browning et al. 2021) of the OMM is necessitated.

4 Learning with Synthetic Data

In this section, we consider the parameter estimation task with synthetic data. First, we discuss our experimental setup and the synthetic dataset. Next, we show that parameter recovery error decreases and stabilizes as we increase the training time T and the number of samples $n_{samples}$.

Experimental Setup. We set $P = M = K = 2$. We set $[\mu_1^1, \mu_2^1, \mu_1^2, \mu_2^2] = [15, 5, 5, 20]$, and $\theta = 0.5$ and draw $\alpha^{pq} \sim \text{Unif}(0, 0.5)$, $\beta_{ij}^{pq} \sim \text{Unif}(0, 0.1)$ and $\gamma_{ik}^p \sim \text{Unif}(0, 0.1)$. The exogenous signals are $S(t) = 1$, $X_1(t) = 5 \sin(0.1x) + 5$, and $X_2(t) = 10 \sin(0.05x + 1.25) + 10$.

We construct our synthetic dataset using the simulation algorithm in Section 3 to get 400 samples of opinion volumes $\{n_{i,t}^p\}_{i,p,t}$ for $t \in \{1, \dots, T = 300\}$. We implement joint fitting (RizoIU et al. 2022): we partition the 400 samples into 20 groups of $n_{samples} = 20$ samples each. The likelihoods $\mathcal{L}_1(\Theta_1)$ and $\mathcal{L}_2(\Theta_2 | \Theta_1)$ of each group are maximised to get an estimate $\hat{\Theta}$, yielding 20 sets of parameter estimates.

Model Evaluation. To study the convergence of our learning algorithm, we compute the root mean-squared error (RMSE) of our estimated $\hat{\Theta} = \{\hat{\mu}_j^p, \hat{\alpha}^{pq}, \hat{\theta}, \hat{\gamma}_{ik}^p, \hat{\beta}_{ij}^{pq}\}$ with respect to the true Θ , following (Valera and Gomez-

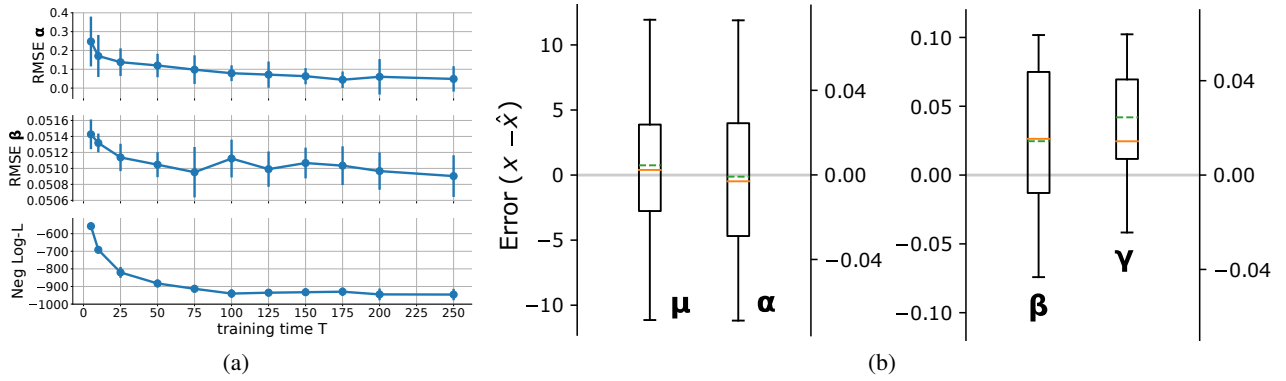


Figure 2: Parameter recovery results on synthetic data. In (a), we show the convergence of the RMSE of the α and β estimates and the negative log-likelihood as we increase the training time T . In (b), we show the difference between our estimates for $\{\mu, \alpha, \beta, \gamma\}$ and the true values. Dashed green lines and orange lines are the mean and median values, respectively.

Rodriguez 2015). We report the average RMSE per parameter type, where the average is taken over the components of the matrix or tensor corresponding to the parameter type.

In Fig. 2(a), we see that training on a longer timeframe leads to lower RMSE for $\hat{\alpha}^{pq}$ and $\hat{\beta}_{ij}^{pq}$ and better model fit measured by the likelihood \mathcal{L}_2 . Results for $\hat{\mu}_j^p$, $\hat{\theta}$ and $\hat{\gamma}_{ik}^p$, and on varying $n_{samples}$ are in the online appendix.

In Fig. 2(b), we plot the difference distribution between our estimates and the true values. We recover first-tier parameters $\{\hat{\mu}_j^p, \hat{\alpha}^{pq}\}$ well, as evidenced by our mean estimates coinciding with the true values. We observe a slight overestimation of $\{\hat{\gamma}_{ik}^p, \hat{\beta}_{ij}^{pq}\}$, given the nonconvexity of \mathcal{L}_2 and the high dimensionality of the second-tier parameter set.

5 Real-World Datasets

This section introduces two real-world datasets we have curated to evaluate the OMM.

Bushfire Opinions Dataset

We construct the *Bushfire Opinions dataset*, containing 90 days of Twitter and Facebook discussions about bushfires and climate change between Nov. 1, 2019 to January 29, 2020. The Facebook postings are a subset of the *SocialSense* dataset (Kong et al. 2022), which was collected with the approval of the Human Research Ethics Committee of the University of Technology Sydney (approval number: ETH19-3877); we select posts and comments about bushfires and climate change (*SocialSense* also contains discussions around COVID-19). Using CrowdTangle², we unobtrusively collected public far-right Australian Facebook discussions, identified via a digital ethnographic study (see (Kong et al. 2022) and (Calderon, Ram, and Rizoiu 2024) for details). We build the Twitter discussions using the Twitter Academic v2 API; we collect tweets emitted between November 1, 2019 to January 29, 2020 that mention bushfire keywords such as *bushfire*, *arson*, *australiaburns* (see the full list in (Calderon, Ram, and Rizoiu 2024)). We use

²<https://www.crowdtangle.com/>

the AWS Location Service³ to geocode users based on their free-text location and description fields and filter only for tweets from Australian users.

Our focus on the 2019-2020 Australian bushfires is motivated by the availability of human-annotated topics, opinions (Kong et al. 2022) and stance classifiers (Ram et al. 2022) trained on the same topic and timeframe. We use these classifiers to filter and label our dataset.

Moderate and Far-Right Opinion Labeling. To filter and label relevant Facebook and Twitter postings, we use the textual topic and opinion classifiers developed by Kong et al. (2022), with a reported 93% accuracy in classifying Facebook and Twitter posts on bushfires and climate change. We select the following most prevalent six opinions, covering 95% of Twitter and 81% of Facebook postings:

0. Greens policies are the cause of the Australian bushfires.
1. Mainstream media cannot be trusted.
2. Climate change crisis is not real / is a UN hoax.
3. Australian bushfires and climate change are not related.
4. Australian bushfires were caused by random arsonists.
5. Bushfires are a normal summer occurrence in Australia.

Furthermore, we deploy the far-right stance detector introduced by Ram et al. (2022) – which leverages a textual homophily measurement to quantify the similarity between Twitter users and known far-right activists. On the *Bushfire Opinions Twitter dataset*, the stance detector achieves a 5-fold CV AUC ROC score of 0.889. An opinion is labeled as *far-right* if the posting agrees with the opinion (denoted as +), or *moderate* if the posting disagrees with the opinion (-). We represent our opinion set as $\{(i-, i+)\mid i \in \{0, \dots, 5\}\}$. In summary, we consider $P = 2$ platforms with 74,461 tweets and 7,974 Facebook postings labeled with $M = 12$ stanced opinions. We aggregate posting volumes by the hour, resulting in $T = 2,160$ time points over 90 days from Nov 1, 2019, to Jan 29, 2020.

Exogenous Signal S and Intervention X. The exogenous signal $S(t)$ (Eq. (5)) modulates the total size of the attention

³<https://aws.amazon.com/location/>




Figure 3: Fitting and predicting with OMM on the Bushfire Opinions dataset. We train OMM on the first 1800 timesteps and predict on timesteps 1801 to 2160 (shaded area). We show results for Facebook (top row) and Twitter (bottom row). (a) Actual (dashed blue lines) vs. fitted/predicted (orange lines) volumes; (b) Actual (left panels) and fitted during training and predicted during testing (right panels) opinion market shares on Facebook and Twitter. We aggregate the far-right and moderate opinions.

market in the first tier of OMM. We use the 5-day rolling average of the Google Trends query *bushfire+climate change* in Australia, normalized to a max value of 1. Google Trends captures the baseline interest on topics (Sheshadri and Singh 2019) and is a proxy for offline events (ex. actual bushfires and government measures) (Milinovich et al. 2014).

The interventions $\{X_k(t)\}$ modulate the market share of far-right and moderate opinions. Our interventions consist of two sources of news coverage: reputable (R) mainstream Australian publishers (e.g., The Sydney Morning Herald, Canberra Times, Crikey) and controversial (C) international publishers (e.g., Sputnik News, Breitbart, Red State). For each opinion $i \in \{0, \dots, 5\}$, we consider a pair of interventions $(R_i(t), C_i(t))$, consisting of reputable and controversial daily news volumes discussing opinion i . We assemble the intervention set $\{X_k(t)\}$ ($K = 12$) so that the first six interventions correspond to $\{R_0(t), \dots, R_5(t)\}$ while the last six correspond to $\{C_0(t), \dots, C_5(t)\}$.

We sourced reputable Australian news publishers from the Reputable News Index (RNIX) (Kong, RizoIU, and Xie 2020). We query Factiva (Johal 2009) to obtain the daily news volume of these outlets for each of the six opinions using a keyword search. We similarly obtain the news volumes from controversial international publishers from NELA-GT-2019 (Gruppi, Horne, and Adalı 2020) using a keyword search. We subtract the Google Trends signal from the news volumes for each intervention. We compute the standardized form of $X_k(t)$ as $\hat{X}_k(t) = \text{news}_k(t) - \frac{\max_t \text{news}_k(t)}{\max_t S(t)} S(t)$. For brevity, in the bushfire case study, we denote $\hat{X}_k(t)$ as $X_k(t)$ (i.e., always in standardized form). After standardiza-

tion, $X_k(t)$ indicates whether reputable or controversial media over- or under-reports relative to the public’s attention.

VEVO 2017 Top 10 Dataset

We assemble the *VEVO 2017 Top 10* dataset by aligning artist-level time series of YouTube views and Twitter post counts ($P = 2$) for the top $M = 10$ VEVO-affiliated artists over $T = 100$ days from Jan 2, 2017 to Apr 11, 2017.

The YouTube time series are obtained from the *VEVO Music Graph dataset* (Wu, RizoIU, and Xie 2019), containing daily view counts for music videos posted by verified VEVO artists in six English-speaking countries (USA, UK, Canada, Australia, New Zealand, and Ireland). We combine the view counts for all music videos that belong to a given artist to obtain artist-level YouTube view time series. For Twitter, we leverage the Twitter API to get daily counts of posts with text containing an input query. We obtain the artist-level Twitter post time series using the artist’s name as the input query.

Unlike the single exogenous signal $S(t)$ in the Bushfire Opinions dataset, we use a different exogenous signal $S_i(t)$ for each artist i – the Google Trends for each artist i . Using the set $\{S_i(t)\}$ instead of a single $S(t)$ requires several small changes to Eq. (5), Eq. (10), and the model gradients. We fully detail these changes in the online appendix. We do not consider any interventions $\{X_k(t)\}$ as we seek to uncover endogenous interactions across artists.

6 Predictive Evaluation

This section evaluates the OMM’s predictive capabilities on two real-world datasets. We introduce our prediction task,

evaluation metrics and baselines, then present the results.

Model Setup. We use a temporal holdout strategy similar to prior literature (Rizoïu et al. 2017, 2018; Kong, Rizoïu, and Xie 2020): we fit OMM on \mathcal{T}_{obs} and evaluate performance on \mathcal{T}_{pred} . Backtesting is another viable alternate evaluation approach; however, it is significantly more computationally intensive, and we prefer the temporal holdout. For the bushfire case study, $\mathcal{T}_{obs} = \{1, \dots, 1800\}$ where time is in hours (i.e., days 1-75 of our period of interest) and $\mathcal{T}_{pred} = \{1801, \dots, 2160\}$ (i.e., days 76-90). For the VEVO case study, $\mathcal{T}_{obs} = \{1, \dots, 75\}$ and $\mathcal{T}_{pred} = \{76, \dots, 100\}$.

We consider two tasks: (1) opinion volume prediction and (2) opinion share prediction. For the first task, we predict the total volume of opinionated posts on the P platforms during the evaluation period. We measure performance using the platform-averaged symmetric mean absolute percentage error (SMAPE) of predicted volumes $\{\bar{n}_t^p | t \in \mathcal{T}_{pred}\}$ on platform p relative to the actual volumes $\{n_t^p | t \in \mathcal{T}_{pred}\}$,

$$\text{SMAPE} = \frac{1}{P} \sum_{p=1}^P \left(\frac{100\%}{360} \sum_{t=1801}^{2160} \frac{|\bar{n}_t^p - n_t^p|}{|\bar{n}_t^p| + |n_t^p|} \right). \quad (11)$$

The predicted opinion volumes $\{\bar{n}_t^p\}$ are obtained using the OMM simulation algorithm. We (1) condition on $\{n_{i,t}^p | t \in \mathcal{T}_{obs}\}$, (2) run the algorithm to sample $\{n_{i,t}^p\}$ on \mathcal{T}_{pred} , then (3) sum over opinion types $\{i\}$ to get predicted opinion volumes $n_t^p = \sum_i n_{i,t}^p$. We repeat $R = 5$ times, and average over the samples to obtain $\{\bar{n}_t^p | t \in \mathcal{T}_{pred}\}$.

For opinion share prediction, we predict the opinion market shares $\{s_{i,t}^p\}$ for each platform p on the evaluation period. To evaluate how well we predict opinion market shares, we calculate the KL divergence of predicted market shares $\{\bar{s}_t^p | t \in \mathcal{T}_{pred}\}$ (obtained similar to $\{\bar{n}_t^p\}$ described above) relative to actual market shares $\{s_t^p | t \in \mathcal{T}_{pred}\}$,

$$\text{KL}^p(t) = \sum_{i=1}^M s_{i,t}^p \log \frac{\bar{s}_{i,t}^p}{s_{i,t}^p}. \quad (12)$$

Baselines. We compare OMM with the discretized versions of the Correlated Cascades (CC) model (Zarezade et al. 2017) and Competing Products (CP) model (Valera and Gomez-Rodriguez 2015) – the current state-of-the-art models in product share modeling, covered in related works. For the bushfire study, we test the effectiveness of interventions by fitting OMM without $\{X_k(t)\}$ (indicated as OMM\X).

We also consider a feature-based predictive baseline – the multivariate linear regression (MLR), used previously for online popularity prediction (Pinto, Almeida, and Gonçalves 2013; Rizoïu et al. 2017). We build MLR with a one-week sliding window of three types of features: the previous event counts, exogenous signal $S(t)$ and interventions $\{X_k(t)\}$. The predictive targets are the event counts $\{n_{i,t}^p\}$ for each point on \mathcal{T}_{pred} . Analogous to OMM fitted without interventions $\{X_k(t)\}$, we additionally train MLR without $\{X_k(t)\}$ (indicated as MLR\X) for the bushfire case study.

OMM, CC and CP are generative models typically designed for explainability and are known to be suboptimal for prediction (Mishra, Rizoïu, and Xie 2016). In contrast,

feature-driven approaches (e.g., MLR) use machine learning to predict using training features. Such approaches are designed mainly for prediction and have weaker explainability since they do not model the data-generation process (Mishra, Rizoïu, and Xie 2016). In this work, we are interested in the dual tasks of predicting and explaining opinion market shares, hence our focus on generative approaches.

Predict Opinion Volumes. Fig. 3(a) showcases the observed (blue line) and modeled (orange line) opinion volumes for the bushfire dataset. We visually observe that OMM achieves a tight fit on both the training and the prediction period (hashed area). The VEVO dataset results are shown in the online appendix. We further compare OMM’s predictive performances against baselines. The top row of boxplots in Figs. 4(a) and 4(b) shows the platform-averaged SMAPE of predicted volumes for the bushfire and VEVO datasets, respectively. We make two observations. First, in both case studies, OMM outperforms all baselines on opinion volume prediction. Second, OMM outperforms OMM\X, indicating the role of media coverage in shaping attention.

Predict Opinion Market Share. Fig. 3(b) visualizes the observed (left column) and fitted during training and predicted during testing (right column) opinion market shares for the bushfire dataset. We see that the opinion distribution on Twitter has significantly more variation than on Facebook, and that OMM closely captures the trend in opinion shares on both platforms. The VEVO dataset results are in the online appendix. Figs. 4(a) and 4(b) show the KL-divergence of predicted market shares for the bushfire (Facebook and Twitter) and VEVO (YouTube and Twitter) datasets, respectively. We make several observations. First, on the bushfire dataset, performance is better for Twitter than Facebook ($\text{KL}^{TW}(t) < \text{KL}^{FB}(t)$) due to Facebook having lower opinion counts than Twitter. Similarly, on the VEVO dataset $\text{KL}^{YT}(t) < \text{KL}^{TW}(t)$. Second, OMM consistently outperforms all baselines on both datasets, except for Twitter on bushfires, where CP and OMM are comparable. CC performs poorly since it does not model asymmetric opinion interactions and assumes all opinions reinforce or inhibit one another. CP performs poorly on Facebook (Twitter) for the bushfire (VEVO) dataset due to CP not having the notion of limited total attention. Due to higher bushfire postings on Twitter, CP pays more attention to Twitter. Lastly, OMM with $\{X_k(t)\}$ outperforms OMM without $\{X_k(t)\}$ on the bushfire dataset, suggesting that mainstream and controversial media effectively shape the opinion ecosystem.

7 Interpreting OMM Elasticities

In this section, we leverage the fitted OMM to uncover interactions across opinions and platforms in the bushfire dataset and artists in the VEVO dataset.

Uncovering Opinions Interactions. To study opinion interactions in the bushfire dataset, we calculate the opinion share model elasticities (see Eq. (4)) accounting for the endogenous volume $\lambda^p(t|j)$ and the intervention $\bar{X}_k(s)$ (see Eq. (9)). The endogenous elasticities $e(s_i^p(t), \lambda^q(t|j))$ quantify the competition-cooperation interactions across opinions. The intervention elasticity $e(s_i^p(t), \bar{X}_k(t))$ quanti-

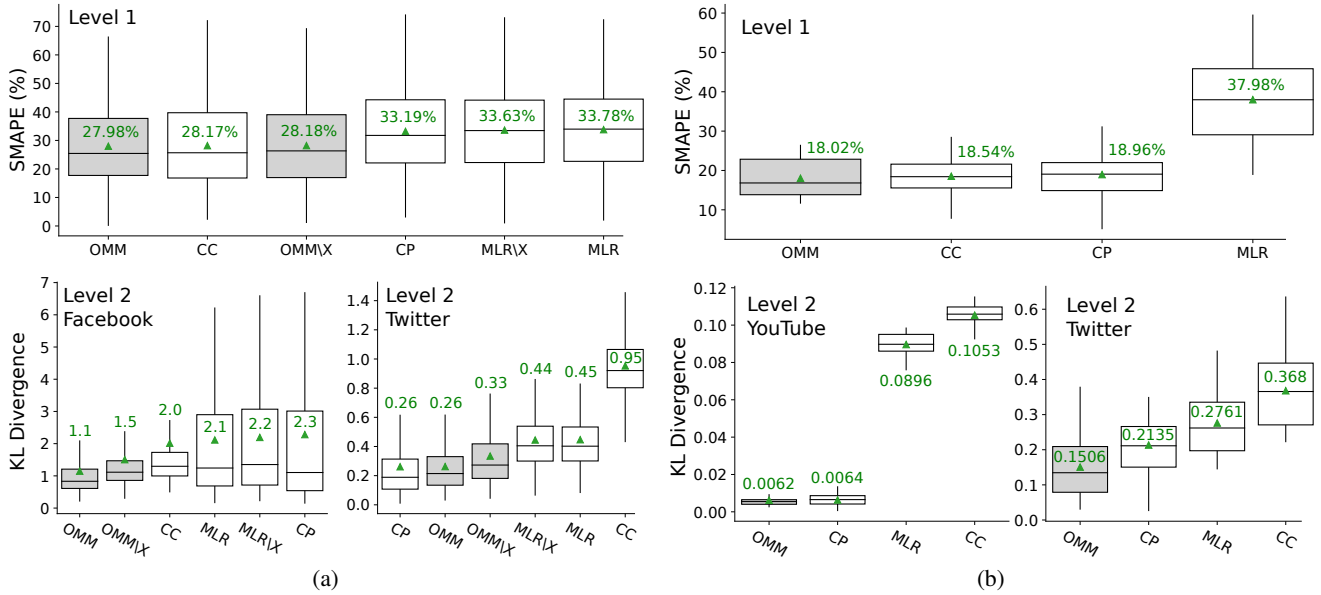


Figure 4: Predictive evaluation of OMM on (a) Bushfire Opinions and (b) VEVO 2017 Top 10 datasets. Boxplots are sorted left to right by the mean (shown with green triangle). Shaded boxplots correspond to versions of OMM. The top panels show the platform-averaged SMAPE of volumes on \mathcal{T}_{pred} . Bottom panels plot the KL divergence of predicted and actual market shares.

ifies the sensitivity of opinion market shares to intervention $X_k(t)$. We derive the elasticities and show results for $e(s_i^p(t), X_k(t))$ in the online appendix. Fig. 5(a) reports the time averages of $e(s_i^p(t), \lambda^q(t|j))$.

First, we study intra-platform reinforcement (top-left & bottom-right in Fig. 5(a)). We see different behaviors for Facebook and Twitter. For Twitter, we have two observations (i.e., main diagonal), indicative of the echo chamber effect (Cinelli et al. 2021). Second, there is significant cross-reinforcement among far-right sympathizers and opponents (i.e., diagonals on the upper-right & lower-left submatrices), implying exchanges or arguments between opposing camps. For Facebook, OMM detects little interaction among opinions, aside from the generally inhibitory effect of the opinions “Australian bushfires and climate change are unrelated” (3+) and “Bushfires are a normal summer occurrence” (5+) on other opinions. This is because Facebook far-right groups have limited interaction with the opposing side.

How to Effectively Suppress Far-Right Opinions. The above implies that confrontation is not the most effective method to suppress far-right opinions, as it has the potential to backfire by bringing even more attention to them. A more effective method is boosting related counter-arguments; for instance, to suppress “Australian bushfires were caused by random arsonists” (4+) on Twitter, OMM indicates to promote “Climate change is real” (2-) and “Greens are not the cause of the bushfires” (0-). Boosting the opposite argument, i.e., “Australian bushfires were not caused by random arsonists” (4-), would backfire. The opinion “Bushfires are a normal summer occurrence in Australia” (5+) shows a different behavior: it reinforces most moderate opinions and

inhibits far-right opinions. In particular, the “Bushfires are normal” opinion (5+) appears to trigger “Climate change is real” (2-), probably due to the diametric opposition nature of these opinions. The effect of 5+ on 2- holds across every pair of platforms. Additionally, on Facebook, “Australian bushfires and climate change are not related” (3+) has a similar effect on other opinions as the “Bushfires are normal” opinion (5+), probably due to the similarity of their topic content.

Cross-Platform Reinforcement is generally weak due to the Facebook far-right groups acting as a filter bubble. Apart from the effect of “Bushfires are normal” (5+) (see above), there is little cross-reinforcement among opinions from Twitter to Facebook. In the bottom-left matrix of Fig. 5(a), we see that “Australian bushfires and climate change are not related” (3+) affects other opinions in a similar way to “Bushfires are normal” (5+); furthermore, “Climate change is real” (2-) triggers “Australian bushfires were caused by arsonists” (4+).

Interactions Across VEVO Artists. Lastly, in Fig. 5(b), we shift our attention to the VEVO dataset and look at the YouTube-to-YouTube elasticities $e(s_i^{YT}(t), \lambda^{YT}(t|j))$ across our set of artists. The Twitter and cross-platform elasticities are available in the online appendix.

We highlight three key observations. First, there is strong self-reinforcement for most artists (i.e., the main diagonal), an intuitive result reflecting these popular artists’ strong fan-base. Second, OMM picks up non-trivial artist interactions that correspond with real-world events – the animosity and friendship relations show up in their popularity dynamics. For instance, we see that Calvin Harris inhibits both Taylor Swift (the two broke up in 2016⁴) and Katy Perry (the

⁴people.com/celebrity/taylor-swift-calvin-harris-breakup-timeline/

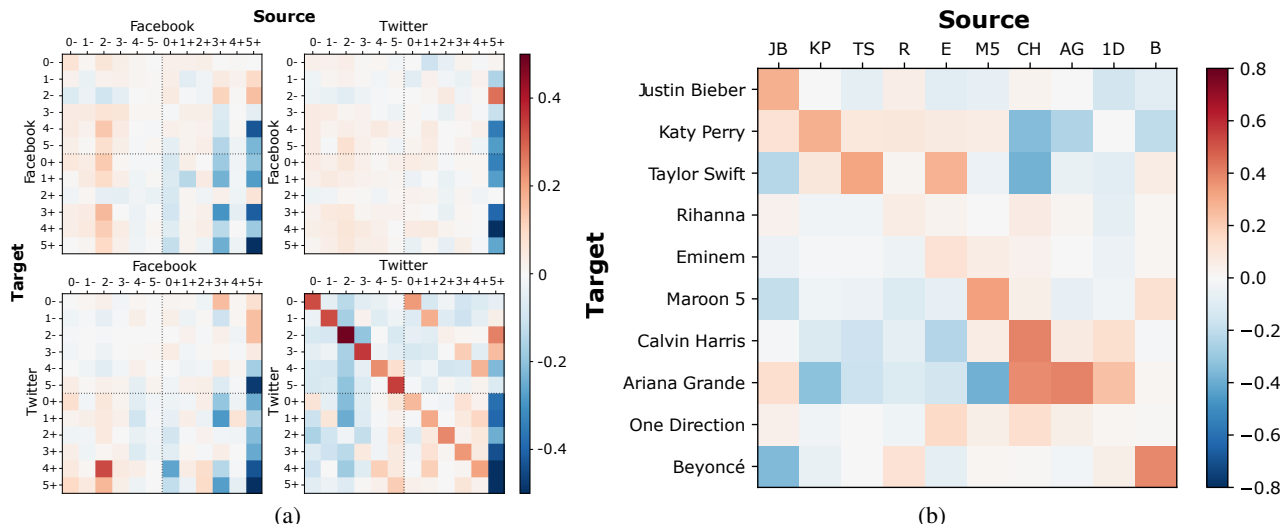


Figure 5: Interpretability of OMM. (a) Endogenous elasticities $e(s_i^p(t), \lambda^q(t|j))$ across opinion pairs (i, j) on respective platforms (p, q) in the bushfire dataset. Elasticities have direction and should be read from column (source) to row (target) for the platform and within each matrix. For example, the bottom-right matrix corresponds to influences from Twitter to Twitter; the cell $\{4-, 4+\}$ ($\{\text{row, column}\}$) is the influence of opinion $4+$ on $4-$, positive and large meaning that $4+$ has a strong reinforcing effect on $4-$. (b) YouTube elasticities $e(s_i^{YT}(t), \lambda^{YT}(t|j))$ across artist pairs (i, j) in the VEVO 2017 Top 10 dataset.

two had a long-lasting feud⁵, due to Harris pulling out of Perry’s 2011 tour last minute). Similarly, Taylor Swift and Justin Bieber have a mutually inhibiting relationship. The two have a well-known uneasy relationship⁶ since Justin Bieber and Selena Gomez used to date and the latter is one of Taylor Swift’s close friends. Meanwhile, Calvin Harris and Ariana Grande have a reinforcing relationship, correlating with their collaboration “Heatstroke” released in March 2017. OMM picks up these relationships because we fit on online popularity driven by audience response. Fans of a given artist can choose to support or not support another artist based on real-world interactions, as indicated by the results above. Our third observation relates to the complexity of fanbase support for artists occupying the same genre: similar artists do not all just cooperate or compete for market share but can have unique pairwise relationships. For instance, Katy Perry, Taylor Swift and Ariana Grande occupy a similar niche (mainstream pop). However, our model uncovers that Taylor Swift and Katy Perry reinforce each other, while these two inhibit (and are inhibited by) Ariana Grande.

8 OMM as a Testbed for Interventions

The interventions $\{X_k(t)\}$ can lead to delayed effects in the opinion ecosystem due to the opinion dependency structure. For example, if an intervention is designed to boost a target opinion, it will indirectly boost all other opinions with a cooperative relationship with the target opinion. Furthermore, it will inhibit opinions with a competitive relationship with the target. Since elasticities only inform us of the *instantaneous* effect on opinion market shares, we perform a what-if

exercise to study the role of interventions in the bushfire case study. We vary the size of the intervention and synthetically sample outcomes to observe the long-term effects of media coverage on the opinion ecosystem.

What-if Can Inform A/B Test Design. We train OMM on observational data; therefore, the inferred effects of interventions $\{X_k(t)\}$ are not causal impact estimates but rather evidence of causal effects. However, the previous section demonstrates that OMM can uncover complex relationships across opinions, providing compelling evidence that OMM is also able to uncover relationships between opinions and interventions. Therefore, the what-if exercise in this section showcases OMM as a testbed for interventions, usable for designing A/B testing that determines true causal effects. The OMM informs us of the effectiveness of interventions, allowing us to prioritize which specific interventions to test.

“What-if” Setup. We test the effect of interventions by synthetically increasing or decreasing their volumes past a given time point (see top panel of Fig. 1) and measuring the percentage change in far-right opinions. Let $k^* \in \{1, \dots, K\}$ be the index of the modulated intervention. We modulate $X_{k^*}(t)$ as $X_{k^*}^{(r)}(t) = X_{k^*}(t) + r \cdot \mu_{X_{k^*}} \cdot \mathbb{1}_{(t > 1800)}$, where $\mathbb{1}_{(\cdot)}$ is the indicator function and $\mu_{X_{k^*}}$ is the mean volume of $X_{k^*}(t)$ on \mathcal{T}_{obs} . The parameter r controls the percent increase ($r > 0$) or decrease ($r < 0$) in media coverage beyond the change point $t = 1800$; $r = 0$ is the original $X_{k^*}(t)$. We run OMM with $X_{k^*}^{(r)}(t)$ for various r , and keep $X_k(t)$ fixed for $k \neq k^*$. We quantify the effects of intervention $X_{k^*}(t)$ as the average percent change (relative to $r = 0$) in the opinion market shares after the change point, i.e., \mathcal{T}_{pred} . We perform this procedure for all $k^* \in \{1, \dots, K\}$.

How News Influences Far-Right Opinions. Fig. 6 shows

⁵ nme.com/news/music/katy-perry-ends-six-year-beef-calvin-harris-2128100

⁶ people.com/music/justin-bieber-selena-gomez-relationship-look-back/




Figure 6: We modulate the volume of reputable (R) and controversial (C) news for each opinion (in $\{0, 1, 2, 3, 4, 5\}$) from -100% to 100% of the mean volume and simulate OMM to see the percent change in the far-right (+) opinion market shares on Facebook (left) and Twitter (right).

the average percent changes in the market share of far-right opinions when modulating the interventions $\{R_i(t), C_i(t)\}$ one at a time for various r over 50 simulations. On Facebook, far-right opinions are suppressed by reputable news and reinforced by the majority of controversial news, except for news concerning “Greens policies are the cause of the Australian bushfires” (R_0) and “Australian bushfires were caused by arsonists” (R_4). On Twitter, both reputable and controversial news suppress far-right opinions, except for reputable news concerning “Australian bushfires/climate change are unrelated” (R_3), “Australian bushfires were caused by arsonists” (R_4) and to a lesser extent “Mainstream media cannot be trusted” (R_1).

We have two key insights. First, we see that the effect of the news on Facebook is modest compared to Twitter since the far-right public groups on Facebook behave as almost perfect filter bubbles in which news has little penetration. Second, indiscriminately increasing reputable news is not an effective strategy for suppressing far-right opinions on Twitter (see R_3 and R_4). Doing so can backfire since it brings even more attention to far-right users and their narratives (Peucker, Fisher, and Davey 2022).

How to Effectively Use the Testbed. Assuming that A/B testing is performed by an entity in control of reputable news coverage (R_i here above), the results above indicate that the test should mainly concentrate on the effects of increasing R_1 (on Facebook), increasing R_0 and decreasing R_3 and R_4 (on Twitter). We leave as future work the design and execution of such an experiment. Our analysis in this paper focuses on mitigating far-right opinions with media coverage. However, OMM can be leveraged as an intervention evaluation tool for information operations in other domains and fighting mis- & disinformation and online propaganda.

9 Summary and Discussion

This work introduces the Opinion Market Model (OMM), a novel two-tier model of the dynamics of the online opin-

ion ecosystem. The first tier models the size of the attention market, and the second tier models opinions competing or cooperating for limited public attention under the influence of positive interventions. We develop algorithms to simulate and estimate OMM, showing the convergence using synthetic data. We demonstrate real-world applicability on a dataset of Facebook and Twitter discussions containing moderate and far-right opinions on bushfires and climate change (Kong et al. 2022) and a dataset of YouTube and Twitter attention volumes for popular artists on VEVO (Wu, Rizoio, and Xie 2019). We show OMM predicts opinion market shares better than state-of-the-art baselines (Valera and Gomez-Rodriguez 2015; Zarezade et al. 2017) and uncovers latent competitive and cooperative interactions across opinions: self-reinforcement attributable to the echo chamber effect and interactions between far-right sympathizers and opponents. Lastly, we quantify the effect of reputable and controversial media coverage on Facebook and Twitter.

Scope of Study. This work focuses on the manifestation of far-right opinions in the context of the 2019-2020 Australian bushfires. Note that far-right ideology manifests in other political issues (e.g., gun control, LGBT rights, xenophobia), which we do not tackle here. Moreover, we do not focus on the general political science of far-right ideology since we are projecting onto a specific context.

Acknowledgments

This work was partially funded by the Australian Department of Home Affairs, the Defence Science and Technology Group, the Defence Innovation Network and the Australian Academy of Science.

References

- Agovino, M.; Carillo, M. R.; and Spagnolo, N. 2021. Effect of Media News on Radicalization of Attitudes to Immigration. *Journal of Economics, Race, and Policy*.
- Betz, M. 2016. Constraints and opportunities: what role for media development in countering violent extremism?
- Browning, R.; Sulem, D.; Mengersen, K.; Rivoirard, V.; and Rousseau, J. 2021. Simple discrete-time self-exciting models can describe complex dynamic processes: A case study of COVID-19. *PLoS ONE*, 16.
- Calderon, P.; Ram, R.; and Rizoio, M.-A. 2024. Online Appendix: Opinion Market Model. <https://arxiv.org/pdf/2208.06620.pdf#page=13>. Accessed: 2024-04-11.
- Cinelli, M.; De Francisci Morales, G.; Galeazzi, A.; Quattrociocchi, W.; and Starnini, M. 2021. The echo chamber effect on social media. *PNAS*, 118(9).
- Clayton, K.; et al. 2020. Real solutions for fake news? Measuring the effectiveness of general warnings and fact-check tags in reducing belief in false stories on social media. *Political Behavior*, 42(4): 1073–1095.
- Cooper, L. G. 1993. Chapter 6 Market-share models. In *Marketing*, volume 5, 259–314. Elsevier.
- De, A.; Valera, I.; Ganguly, N.; Bhattacharya, S.; and Gomez-Rodriguez, M. 2016. Learning and Forecasting Opinion Dynamics in Social Networks. *NIPS’16*, 397–405.

- Ferguson, K. 2016. Countering violent extremism through media and communication strategies. *Reflections*, 27: 28.
- Fujita, K.; Medvedev, A.; Koyama, S.; Lambiotte, R.; and Shinomoto, S. 2018. Identifying exogenous and endogenous activity in social media. *Phys. Rev. E*, 98: 052304.
- Garetto, M.; Leonardi, E.; and Torrisi, G. L. 2021. A time-modulated Hawkes process to model the spread of COVID-19 and the impact of countermeasures.
- Gelper, S.; van der Lans, R.; and van Bruggen, G. 2021. Competition for attention in online social networks: Implications for seeding strategies. *Management Science*.
- GIFCT. 2021. Content-Sharing Algorithms, Processes, and Positive Interventions Working Group.
- Gruppi, M.; Horne, B. D.; and Adalı, S. 2020. NELA-GT-2019: A Large Multi-Labelled News Dataset for The Study of Misinformation in News Articles. arXiv:2003.08444.
- Guess, A. M.; Barberá, P.; Munzert, S.; and Yang, J. 2021. The consequences of online partisan media. *PNAS*, 118(14).
- Gupta, S.; Jain, G.; and Tiwari, A. A. 2022. Polarised social media discourse during COVID-19 pandemic: evidence from YouTube. *Behaviour & Information Technology*, 1–22.
- Hacıyakupoglu, G.; Hui, J. Y.; Suguna, V.; Leong, D.; and Rahman, M. F. B. A. 2018. Countering fake news: A survey of recent global initiatives.
- Hawkes, A. G. 1971. Spectra of Some Self-Exciting and Mutually Exciting Point Processes. *Biometrika*, 58(1).
- Henschke, A.; and Reed, A. 2021. Toward an Ethical Framework for Countering Extremist Propaganda Online. *Studies in Conflict & Terrorism*, 1–18.
- Horowitz, M.; Cushion, S.; Dragomir, M.; Gutiérrez Manjón, S.; and Pantti, M. 2022. A framework for assessing the role of public service media organizations in countering disinformation. *Digital Journalism*, 10(5).
- Jackson, S. 2019. The double-edged sword of banning extremists from social media.
- Johal, R. 2009. Factiva: Gateway to Business Information. *Journal of Business & Finance Librarianship*, 15(1): 60–64.
- King, G.; Schneer, B.; and White, A. 2017. How the news media activate public expression and influence national agendas. *Science*, 358: 776–780.
- Kong, Q.; Booth, E.; Bailo, F.; Johns, A.; and Rizoïu, M.-A. 2022. Slipping to the Extreme: A Mixed Method to Explain How Extreme Opinions Infiltrate Online Discussions. In *AAAI ICWSM*, volume 16, 524–535.
- Kong, Q.; Rizoïu, M.-A.; and Xie, L. 2020. Describing and predicting online items with reshare cascades via dual mixture self-exciting processes. In *29th ACM CIKM*, 645–654.
- Kulkarni, B.; Agarwal, S.; De, A.; Bhattacharya, S.; and Ganguly, N. 2017. SLANT+: A nonlinear model for opinion dynamics in social networks. In *ICDM*, 931–936. IEEE.
- Milinson, G. J.; Williams, G. M.; Clements, A. C. A.; and Hu, W. 2014. Internet-based surveillance systems for monitoring emerging infectious diseases. *The Lancet Infectious Diseases*, 14(2): 160–168.
- Mishra, S.; Rizoïu, M.-A.; and Xie, L. 2016. Feature driven and point process approaches for popularity prediction. In *CIKM*, 1069–1078.
- Nekmat, E. 2020. Nudge effect of fact-check alerts: source influence and media skepticism on sharing of news misinformation in social media. *Social Media+ Society*, 6(1).
- Peucker, M.; Fisher, T. J.; and Davey, J. 2022. Mainstream media use in far-right online ecosystems. Technical report.
- Pinto, H.; Almeida, J. M.; and Gonçalves, M. A. 2013. Using early view patterns to predict the popularity of youtube videos. In *WSDM*, 365–374.
- Porter, E.; and Wood, T. J. 2021. Fact checks actually work, even on Facebook. But not enough people see them. *The Washington Post*.
- Radsch, C. 2016. Media Development and Countering Violent Extremism: An Uneasy Relationship, a Need for Dialogue. *Center for International Media Assistance*. (2016).
- Ram, R.; Thomas, E.; Kernot, D.; and Rizoïu, M.-A. 2022. Detecting Extreme Ideologies in Shifting Landscapes: an Automatic & Context-Agnostic Approach. arXiv:2208.04097.
- Rizoïu, M.; Mishra, S.; Kong, Q.; Carman, M.; and Xie, L. 2018. SIR-Hawkes: Linking epidemic models and Hawkes processes to model diffusions in finite populations. In *WWW*, 419–428.
- Rizoïu, M.; Xie, L.; Sanner, S.; Cebrian, M.; Yu, H.; and Van Hentenryck, P. 2017. Expecting to be HIP. In *WWW 2017*.
- Rizoïu, M.-A.; Soen, A.; Li, S.; Calderon, P.; Dong, L.; Menon, A. K.; and Xie, L. 2022. Interval-censored Hawkes processes. *Journal of Machine Learning Research*, 23(338).
- Rizoïu, M.-A.; and Xie, L. X. 2017. Online Popularity Under Promotion: Viral Potential, Forecasting, and the Economics of Time. *ICWSM*, 11(1): 182–191.
- Sheshadri, K.; and Singh, M. P. 2019. The public and legislative impact of hyperconcentrated topic news. *Science Advances*, 5(8).
- Shu, K.; Wang, S.; and Liu, H. 2019. Beyond News Contents: The Role of Social Context for Fake News Detection. *WSDM '19*, 312–320.
- Valera, I.; and Gomez-Rodriguez, M. 2015. Modeling Adoption and Usage of Competing Products. In *ICDM*.
- Venturini, T.; and Rogers, R. 2019. “API-based research” or how can digital sociology and journalism studies learn from the Facebook and Cambridge Analytica data breach. *Digital Journalism*, 7(4): 532–540.
- Weng, L.; Flammini, A.; Vespignani, A.; and Menczer, F. 2012. Competition among memes in a world with limited attention. *Scientific reports*, 2(1): 335.
- Wu, S.; Rizoïu, M.-A.; and Xie, L. 2019. Estimating Attention Flow in Online Video Networks. *CSCW*, 3: 1–25.
- Young, G. K. 2022. How much is too much: the difficulties of social media content moderation. *Information & Communications Technology Law*, 31(1): 1–16.
- Zarezade, A.; Khodadadi, A.; Farajtabar, M.; Rabiee, H. R.; and Zha, H. 2017. Correlated cascades: Compete or cooperate. *AAAI 2017*, 238–244.

Paper Checklist

1. For most authors...

- (a) Would answering this research question advance science without violating social contracts, such as violating privacy norms, perpetuating unfair profiling, exacerbating the socio-economic divide, or implying disrespect to societies or cultures? **Yes, we have designed the OMM as a tool to promote social good, functioning as a testbed to evaluate positive interventions in opinion moderation.**
- (b) Do your main claims in the abstract and introduction accurately reflect the paper's contributions and scope? **Yes, the main claims in the abstract and introduction reflect the scope of the paper.**
- (c) Do you clarify how the proposed methodological approach is appropriate for the claims made? **Yes, in Section 1 we discuss how the OMM models the online opinion ecosystem using an analogy between opinions and economic goods. In Sections 4 and 6 we show how our approach predicts better than competing models. In Sections 7 and 8 we show how the OMM unveils latent relationships and acts as a testbed for positive interventions.**
- (d) Do you clarify what are possible artifacts in the data used, given population-specific distributions? **Yes, to curate the Bushfire Opinions dataset we performed a filtering procedure which we describe in detail in Section 5 of the main text and Section G of the Appendix.**
- (e) Did you describe the limitations of your work? **Yes, we discuss the limitations of the model in the 'Causal Impact' subsection of Section 1 and the limitations of the scope of study in Section 9.**
- (f) Did you discuss any potential negative societal impacts of your work? **Yes, we discuss potential negative societal impacts (i.e. misuse by oppressive regimes) in the 'Ethics of Opinion Moderation and Broader Perspectives' subsection of the Introduction.**
- (g) Did you discuss any potential misuse of your work? **Yes, we discuss potential misuse in the 'Ethics of Opinion Moderation and Broader Perspectives' subsection of the Introduction.**
- (h) Did you describe steps taken to prevent or mitigate potential negative outcomes of the research, such as data and model documentation, data anonymization, responsible release, access control, and the reproducibility of findings? **Yes, we described how to mitigate ethical issues when deploying the OMM in the real world (i.e. online CVE frameworks) in the 'Ethics of Opinion Moderation and Broader Perspectives' subsection of the Introduction.**
- (i) Have you read the ethics review guidelines and ensured that your paper conforms to them? **Yes, we have completed the Ethics checklist and included a discussion on ethical implications of the work in the 'Ethics of Opinion Moderation and Broader Perspectives' subsection of the Introduction.**

2. Additionally, if your study involves hypotheses testing...

- (a) Did you clearly state the assumptions underlying all theoretical results? **Yes, the OMM model is fully specified and developed in Section 3. We discuss stability assumptions at the end of Section 3. Full implementation details are provided in the Appendix.**
 - (b) Have you provided justifications for all theoretical results? **Yes, proofs of theoretical results are contained in the Appendix.**
 - (c) Did you discuss competing hypotheses or theories that might challenge or complement your theoretical results? **Yes, we consider alternative models in the Related Works section of Section 1. In Section 6 we compare against other models (Correlated Cascades and Competing Products) and show that the OMM fits and predicts better.**
 - (d) Have you considered alternative mechanisms or explanations that might account for the same outcomes observed in your study? **Yes, we consider alternative models based on different assumptions (Correlated Cascades and Competing Products) and show that the OMM fits and predicts better.**
 - (e) Did you address potential biases or limitations in your theoretical framework? **Yes, we discuss how the OMM does not prove causal impact in the 'Causal Impact' subsection in Section 1.**
 - (f) Have you related your theoretical results to the existing literature in social science? **Yes, in Section 7 we connect results from the OMM with the echo chamber effect discussed in prior literature.**
 - (g) Did you discuss the implications of your theoretical results for policy, practice, or further research in the social science domain? **Yes, we discuss how the OMM can be used as an intervention evaluation tool for information operations in Section 8.**
- ### 3. Additionally, if you are including theoretical proofs...
- (a) Did you state the full set of assumptions of all theoretical results? **Yes, model assumptions are discussed in Section 3, particularly the Stability Assumption subsection.**
 - (b) Did you include complete proofs of all theoretical results? **Yes, proofs of theoretical results are contained in the Appendix.**
- ### 4. Additionally, if you ran machine learning experiments...
- (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **Yes, the Appendix contains implementation details and we have linked the code repository as a URL.**
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **Yes, training details are discussed in the Experimental Setup subsection in Sections 4 and 6.**
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **Yes, error bars are indicated in Figures 2 and 4.**

- (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **No, but we discussed implementation details in the Appendix.**
 - (e) Do you justify how the proposed evaluation is sufficient and appropriate to the claims made? **Yes, we demonstrate how the OMM outperforms the state-of-the-art in predicting opinion volumes and opinion market shares in Section 6. In Section 7 we discuss interpretation of results and connect it with existing literature on echo chambers in social media.**
 - (f) Do you discuss what is “the cost“ of misclassification and fault (in)tolerance? **No, our work deals with regression (i.e. predicting opinion volumes and market shares). Fault tolerance is not relevant at the moment as we are proposing a model, but would be if the OMM is integrated as a component of an automated system.**
5. Additionally, if you are using existing assets (e.g., code, data, models) or curating/releasing new assets, **without compromising anonymity...**
- (a) If your work uses existing assets, did you cite the creators? **Yes, datasets are described and attributed in Section 5. Comparison models are discussed and attributed in the Baselines subsection in Section 6.**
 - (b) Did you mention the license of the assets? **No, but we have appropriately cited all datasets we used and models we compared against.**
 - (c) Did you include any new assets in the supplemental material or as a URL? **Yes, we included the code and data used in the experiments as a URL: <https://github.com/behavioral-ds/opinion-market-model>.**
 - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? **No, because the data that we used (SocialSense) only consists of publicly available social media posts.**
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? **Yes, we describe the SocialSense dataset that we use in Section G of the Appendix, and discuss how it consists of online content expressing problematic opinions on climate change.**
 - (f) If you are curating or releasing new datasets, did you discuss how you intend to make your datasets FAIR? **No, but we intend to do it as a next step.**
 - (g) If you are curating or releasing new datasets, did you create a Datasheet for the Dataset? **No, but we intend to do it as a next step.**
6. Additionally, if you used crowdsourcing or conducted research with human subjects, **without compromising anonymity...**
- (a) Did you include the full text of instructions given to participants and screenshots? **NA**
 - (b) Did you describe any potential participant risks, with mentions of Institutional Review Board (IRB) approvals? **NA**
- (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? **NA**
 - (d) Did you discuss how data is stored, shared, and de-identified? **NA**