Simulating the Impact of Personality on Fake News

Annabel Coates
University of Nottingham
annabelcoates@gmail.com

Tim Muller
University of Nottingham
tim.muller@nottingham.ac.uk

Sean Sirur
University of Oxford
sean.sirur@gmail.com

Abstract

Fake news is a key issue for social networks. We use an agent-based network simulation to model the spread of (fake) news. The agents' behaviour captures the OCEAN ("big five") personality trait model. The network is homophilic for political preference, analytical thinking and emotion. We studied the system with personality traits and homophily each turned on/off. Personality traits and homophily exhibited a statistically significant but typically minimal impact. Ignoring personality traits when modelling fact-checking can overestimate its effectiveness.

1 Introduction

The spread of false information or "fake news" is a key issue on online social networks (OSNs). With 55% and 49% of US and UK adults now getting news from social networks [SG19, Ofc19], the effects are broad and potent. For example, fake news has been shown to spread more effectively than true news on social media [VRA18]. In the medical domain, misinformation was present in 40% of top links relating to common diseases [WKWK18]. This fuels dangerous ideas among the public [CMP19] and leads to avoidable public health crises [Hot16].

To counter fake news, it is essential to investigate both the nature of its spread as well as the OSN users who create and share it, at times unintentionally. From a user's perspective, there is an inherent uncertainty in the verity of a news item. Clearly, trust is an important factor in the spread of news on OSNs, both in terms of shares and views, and an honest user will not share a news item they distrust. The trustworthiness of a news item is also deeply related to its believability: with users' susceptibility to fake news varying with e.g. analytical thinking skills [PR19b]. Still, it is important to recognize that whether a news item is true or fake is not the only factor in a user's decision to share it.

We hypothesise that users' *personalities* affects how and why (fake) news spreads. Various prior works study the effects of personality on OSN usage [CCTS14, ROS+09, AV10]. As shown by [CCT17], understanding how personality affects the spread of news is important. Agent-based models can be used to capture rich user properties such as detailed internal states, allowing for behaviours of greater depth. As such, we use these results to implement personality effects in the users.

Network models are used in a range of disciplines including sociology, computer science and economics. As such, their use extends naturally to the study of OSNs. The presence of homophily, the tendency for individuals to connect to those similar to themselves, is common to social networks [DJH+16, MSLC01]. Simulations of social networks that do not take homophily into account cannot have accurate network structures [AMS13]. Yet, its effects on (fake) news spreading is relatively unexplored. We implement a network that is homophilic for political leaning, analytic thinking skills and emotional state.

In: R. Falcone, J. Zhang, and D. Wang (eds.): Proceedings of the 22nd International Workshop on Trust in Agent Societies, London, UK on May 3-7, 2021, published at http://ceur-ws.org

The aims of this study are to: 1) Account for personality traits in user behaviour 2) Account for homophily in the network structure 3) Determine the impact of these features on the spread of fake news.

2 Relevant literature

In order to accurately simulate user behaviour in OSNs, it is important to consider social and psychological studies. Specifically, how users interact with OSNs [AM15], how personality traits affect use of OSNs [CCTS14, AV10, ROS+09] and how individuals perceive fake news [PR19b, SCM20, PCR18].

We also discuss similar agent-based studies and their findings. Including the use of agent-based modelling as an effective way of studying the spread of fake news, and comparison of simulation methods.

2.1 Social Science

One study explored user interactions within OSNs [AM15]. Descriptive data from surveys were used to identify motivations for using different social networks and the interactions within them. It found that, for Facebook use, entertainment and information sharing were the most valued reasons. I was also found that content is more likely to be "liked" than shared on Facebook. Usage motivations may also be affected by personality which must be considered when interpreting such results. The authors provide a quantitative analysis of users' motivations for using Facebook and partaking in certain interactions, this behavioural perspective is valuable for simulating how users interact with other users and posts. However, relationships and data obtained in this study are representative of a relatively small demographic, as the participants of the surveys were all college aged. Therefore the study is limited in representing older users that are less computer literate [MP12] and can play a large role in the spread of fake news [Cho19].

[PR19b] study the cause of susceptibility to fake news. In particular, they study the extent to which susceptibility lies in motivated reasoning (wanting to believe something is true) or failing to identify that something is false (lack of analytical thinking). Participants completed cognitive reasoning tests and rated their perceived accuracy of fake and true news headlines. From this it was found that analytic thinking gave participants a higher ability to differentiate between true and fake news. This led the authors to conclude that failing to discern between fake and true news was the more significant factor than partisanship, in believing it.

In line with [PR19b], [SCM20] find that analytical thinking is negatively related to susceptibility to fake news. An additional finding by [SCM20], is that participants believe the news consistent with their own political beliefs to be more accurate. Therefore the combined findings of [PR19b] and [SCM20], tell us that while analytic thinking is the primary factor in susceptibility to sharing fake news, (politically) motivated reasoning must also be taken into account. Our model uses these idea in having agents assess the news.

Caci et al [CCTS14] provide a detailed analysis on how the OCEAN personality variables [Gol93] affect Facebook use (e.g. the amount of connections made or the session frequency and length). Paired with the distribution of personality traits in the population, this represents an accurate model of the population's use of Facebook. We incorporate the following relationships from Caci et al in our simulation:

- Conscientiousness (B=-0.18) and agreeableness (-0.21) are predictors of lower frequency of use.
- Extroversion (B=0.12) and neuroticism (B=0.14) are predictors of high frequency of use.
- Conscientiousness (B=-0.16) is a predictor of low session length.
- Extroversion (B=0.24) and neuroticism (B=0.14) are predictors of high session length.
- Conscientiousness (B=-0.28) and agreeableness (B=-0.28) are predictors of a lower number of connections.
- Openness (B=0.24) and extroversion (B=0.47) are predictors of a higher number of connections.

Given that Caci et al studied only the population of Italy, our results apply primarily to that population. However, the OCEAN model aims to provide an exhaustive, universally understood representation of personality. As such, our results help to lay a foundation upon which future research may generalize.

Similarly to [CCTS14], [ROS+09] use self-report personality tests and a questionnaire of Facebook use, to gather data on the link between the OCEAN personality traits and Facebook use. They find that an individual's 'motivation to communicate' was largely influential on how they used Facebook. They also make predictions relating a willingness to share personal information to neuroticism, and the use of a range of communication features with extroversion. These findings have been used within this study to determine how frequently different users share posts in general (share probability).

Based on the work of [ROS+09], [AV10] seek to determine relationships between personality and Facebook use. They do this using "more objective" data than [ROS+09], gathering data from participants' Facebook profiles. In agreement with [CCTS14] they find that highly extroverted individuals have a "signficantly higher number of friends". However some of their findings disagree with those of [CCTS14], including that: 1) Higher agreeableness is not a predictor more Facebook friends, and 2) conscientiousness is a predictor of more Facebook friends. In agreement with [ROS+09]'s predictions, [AV10] find highly neurotic individuals post more personal information, providing confidence in this result's accuracy.

2.2 Data Analyses

In [GRG+20], the aim is to determine the link between personality and linguistics of users, and their likelihood to share fake news, or fact-check. That research question is similar to ours. They create a convolutional neural network (CheckerOrSpreader) that distinguishes between users based on linguistic patterns and features that represents users' personality traits within tweets. The results of the study show that the model has improved performance at identifying spreaders and fact-checkers when personality and linguistic features are taken into account. These results can be directly compared to the results of this study in reference to whether personality affects the spread of fake news, despite the differences in methodology. As well as using real data instead of simulated behaviour, Gianchanou et al.'s study focuses on spreading behaviour on an individual level, whereas this study investigates the effect of personality on the spread throughout a network.

The authors of [VRS18] aim to determine how and why fake news spreads differently from true news. The authors analyse the properties of fake news content, the spread of fake news and properties of susceptible users. They find results applicable to this model, such as that users that spread fake news tend to follow fewer people, and have fewer followers. This provides insight into other network properties that could affect the spread of fake news.

Through content and linguistic analysis, [VRS18] find that certain categories of fake news – politics and urban legends – are particularly susceptible to large spread. The importance of the category of news found provides motivation to take into account individual interests, as is done so in this model through the political leaning of users and news. They find a higher emotional content in fake news, with replies and comments expressing greater levels of disgust and surprise than true news. This supports the use of emotional level as an important attribute of news articles, as in this model.

Key findings on homophily in social networks are found by [DJH+16], who focus on homophily in twitter users of similar moral purity, determined by analysing tweet content. They find that social networks reflect moral selection and that offline and online differences in moral purity are particularly predictive of distance between two individuals in the social network. This supports the modelling of homophily in this study, in particular, homophily around political leaning as this is contextually similar to moral purity.

2.3 Other simulation studies

Jin et al. create a epidemiological model of the spread of true and fake news on Twitter based on data analysis of 8 viral stories [JDS+13]. The population is modelled with 4 possible states for users to be in (Susceptible, Exposed, Infected and Sceptical), when in certain states users can 'infect' other users and spread fake news. They prove that this model is accurate in representing the spread of both true and fake news. As well as this, they generate specific results such as that users obtain true news from many sources versus limited or single sources for fake news, causing an initial spike in shares when true news is first made public, which is not present for fake news. Our agent-based approach investigates whether such a model of uniform agents is realistic; specifically, whether it is fine to ignore personality.

In [BFP11], the authors create a study simulating an OSN (Facebook) in an agent-based approach (as we do). They model Facebook by defining 4 groups of users (typical user, power-user, town crier, deep-end diver) with different behavioural characteristics. These groups are then given representative weights within the population. They create a model of Facebook that reproduces the real world phenomena of increased connectivity in the community and increased heterogeneity of the population with time. Whereas they use 4 categories of users with their own characteristics, we use personality traits as hidden variables determining characteristics.

3 Model description

The simulation model for this study uses an agent-based modelling approach. *Agents* are discrete entities with their own goals and behaviours, their defining quality is making decisions based not only on their environment but also their own *internal* state. Of particular importance to our study are personality traits of agents.

In this model, the agents are users of an online social network based on Facebook. It is effective for users to be modelled as agents, as their actions depend on both the environment they are in (such as the posts they view and other users) and their internal state. As explained, we suppose that people with similar political leaning, analytical skills and emotional state tend to cluster together – a network property known as homophily [New18]. As our study focuses on the effects of personality traits and homophily, our simulation can run with personality traits and homophily individually being turned on, down or off.

3.1 Modelling Entities

In order to portray an individual's behaviour towards news seen online, the effect of personality [CCTS14, AV10], analytic thinking [PR19b] and interests (i.e political leaning) [PR19b] need to be considered.

Personality traits. In this model, personality is represented using the OCEAN (or "big five") model of personality traits as defined in [Gol93]: *openness, conscientiousness, extroversion, agreeableness and neuroticism.* We implement the results of Caci et al [CCTS14] into the model directly.

For simulation purposes, we introduce a personality factor, P, a value between 0 and 1, determined the weighting that the different personality traits have in the model. We set all other parameters, such that their distribution remains unchanged as we remove the effect of personality traits. This way any effect resulting from changing P must be due to personality traits interacting with other variables.

When creating a population of people within the model, a normal distribution of the prevalence of each personality trait within the 2015 UK population was used [RJL15] to ensure the prevalence within the model reflected that of the real world.

The weights from Caci et al.'s path analysis [CCTS14] are taken directly to determine an individual Facebook user's number of friends, frequency of use and session duration within the model. If $P = 1$, these values are based purely on openness (o), conscientiousness (c), extroversion (e), agreeableness (a), neuroticism (n). A user's tendency to have a large network, C is defined as:

$$C = 0.24o - 0.28c + 0.47e - 0.28a + 0.2d_1 \tag{1}$$

A users frequency of use, f is defined as:

$$f = -0.18c + 0.12e - 0.21a + 0.14n, \tag{2}$$

A user's session length,L is defined as:

$$L = -0.16c + 0.24e + 0.14n \tag{3}$$

A user's sharing tendency,w is defined as (after [AV10]):

$$w = \frac{n + e}{2} \tag{4}$$

The variables were diluted and normalised as required, depending on the personality factor P. In the case that $P = 0$ these variables are randomly distributed along the same prior distribution.

Analytic thinking. Each individual also had an analytic thinking value, representing a user's general likelihood of believing fake news. This captures analytic thinking as described in the literature [PR19b, SCM20]. While some studies find a link between personality traits and analytic thinking [AR20, Yam20, CBK04], it was decided that defining the traits independently was of more value to this model.

Political leaning The political leaning attribute is used to determine the level of relevancy that a specific news article has to an individual. Although this attributed is named political leaning, it encompasses all interests and opinions that a user holds. The magnitude of the difference between the political leaning of an individual and a news article they are viewing, determines the relevancy to an individual. A higher relevancy leads to a higher probability of sharing.

Emotional state Due to the finding that news that invokes more of an emotional response receives more shares [VRS18], each user in this model has an attribute, emotional state, to represent their susceptibility to emotional content. This was independent of personality within this model.

Assessing News When viewing a true or fake news article online, an individual's likelihood of sharing the news depends on a number of factors. In this model three factors are used to determine this likelihood:

1. Believability: the extent to which the person believes the news is true. This factor has been found to have a large effect on the sharing of fake news [PR19b]

2. Relevancy: how much the news aligns with the person's (political) beliefs. Modelled as 1 dimensional.
3. Emotional level: the level of emotional response that the news invokes in the viewer, for example Vosoughi et al. link the likelihood of sharing fake news articles with their emotional content.

Believability is calculated from:

$$B = a \cdot (b_{news} - 1) + 1 = a \cdot b_{news} + (1 - a) \tag{5}$$

where b_{news} measures the *intrinsic* believability of the news and a measures the analytic thinking of the viewer. For high analytical thinkers $B \approx b_{news}$, meaning they estimate the believability of the news well. For low analytical thinkers, $B \approx 1$, meaning they tend to believe most news, regardless of its veracity.

The emotional response invoked by a piece of news, is related to both the emotional content of the news and how susceptible the viewer is to emotional content. One measure of this is a person's level of neuroticism, as it is an indicator of emotional instability [CM99, ORR12, Gol93]. Due to this, the emotional factor, E, in this model was determined as:

$$E = n \cdot m_{news}, \tag{6}$$

where n is the individual's neuroticism and m_{news} is the emotional level of the news (for $P = 1$).

Finally, the relevancy factor of the news exists to measure the degree to which the news relates to a person and their life. Therefore the relevancy factor is negatively dependent on how different the political leaning of the news is and the person's political leaning of topics are. It is therefore calculated from:

$$R = \max(0.4 - 0.4 \cdot |l_{news} - l_{person}|, 1 - 5 \cdot |l_{news} - l_{person}|) \tag{7}$$

where l_{news} and l_{person} are the political leaning of the news and of the person respectively. The formula is sensitive to nearby leanings, but still allows for non-zero values for further news.

These factors are used to determine the share probability, S, for a specific person and news article as follows:

$$S = w \cdot \frac{2 \cdot B + R + 1.8 \cdot E \cdot s}{4}, \tag{8}$$

where w is a person's share tendency, R is the relevancy factor, E is the emotional factor, s is the person's emotional state and B is the believability factor.

[PR19a] and [PR19b] found it is a belief or lack of belief in a news article that most significantly limits the spread of news in OSNs, hence, in the formula B is weighted 2. The factor 1.8 is to compensate for the fact that s is a fractional value between 0 and 1 – on average $1.8 \cdot s \approx 1$.

News Differences in the content and perception of news was modelled via attributes measuring *emotional level*, *believability* and *political leaning*. These attributes were given normally distributed values between 0 and 1. *Believability* is a measure of the intrinsic credibility of a news item. Belief in a news item depends on its believability and the individual's online literacy. It was found that, on average, social media users are "quite good" at distinguishing between news sources of different quality [PR19a]. As such, the believability of fake and true news were each generated from a low distribution with $\mu = 0.2, \sigma = 0.25$ and a high distribution with $\mu = 0.8, \sigma = 0.25$, respectively. *Political leaning* determines the news' level of relevancy to an individual viewer. Political leaning is modelled neutrally and symmetrically along a "left-right" axis.

3.2 Modelling the Network and Homophily

OSN structure is captured through mathematical graphs or *networks*. Such networks are composed of *nodes*, representing the OSN users, and *links*, representing two users mutually "following" each other.

One crucial activity in the study of real-world networks is the development of algorithms that can *generate* realistic simulated networks. The properties of such networks depend on the algorithm as different ones each capture different properties that are believed to be common to social networks. Two foundational algorithms are described here:

Barabasi-Albert (BA) model: This algorithm models "preferential attachment". That is, nodes are more likely to be attached to a node that already has many other attachments.

Connected Watts-Strogatz (CWS) model: This model ensures the network exhibits high clustering (the number of *actual* links between a node's *neighbours* divided by the number of *potential* links between them) and small-world effects (the average number of hops between any two nodes grows "slowly" with the network size).

We ran our simulations across both these types of networks, finding similar results. Except Figure 4, all figures are computed using the Barabasi-Albert model with degree 30. Figure 4 is included as a typical comparison.

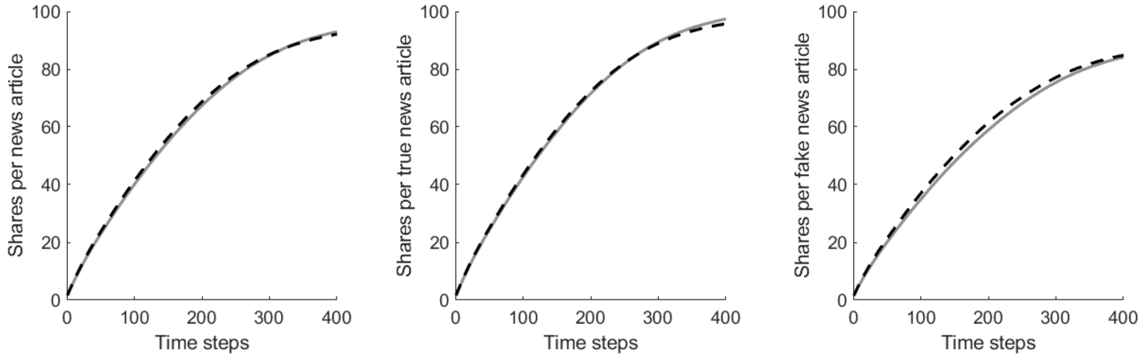

(a) All news (b) True news (c) Fake news

Figure 1: Plots of how often news is shared over time (topology BA).

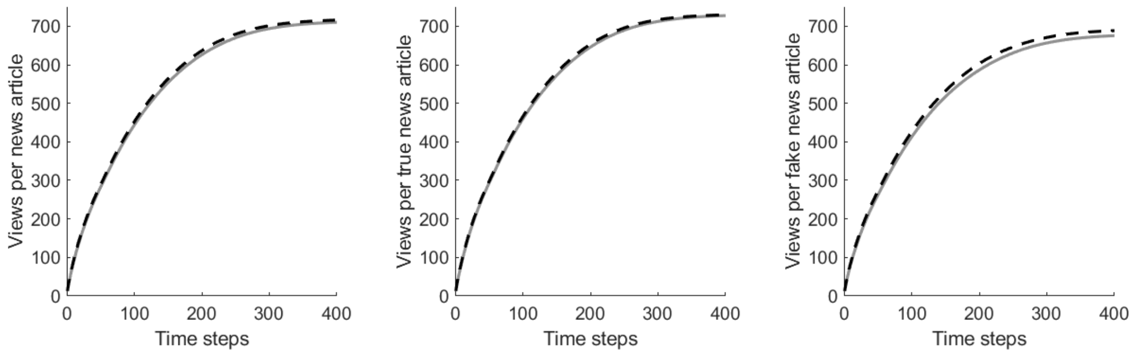

(a) All news (b) True news (c) Fake news

Figure 2: Plots of how often news is viewed over time (topology BA).

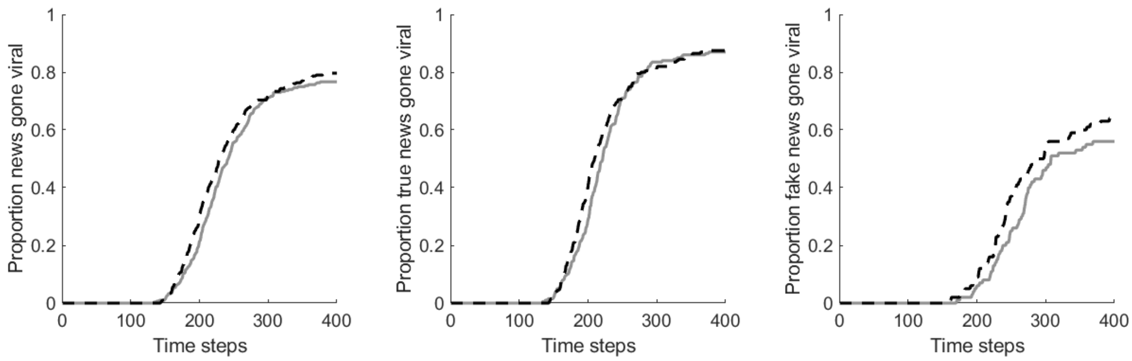

(a) All news (b) True news (c) Fake news

Figure 3: Plots of how often news goes viral over time (topology BA).

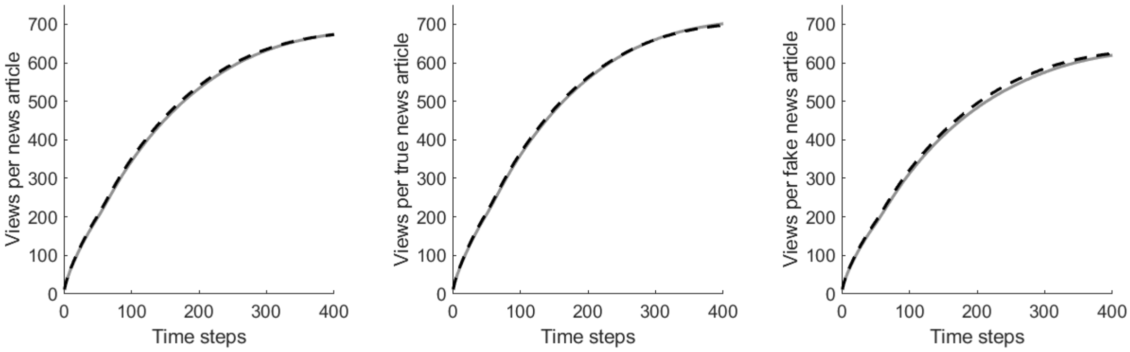

(a) All news (b) True news (c) Fake news

Figure 4: Plots of how often news is viewed over time (topology CWS).

We implement a degree of homophily in the network structure. Neighbours are likely to share political leaning, analytic thinking and emotional state. We do not assume homophily of personality traits directly. The degree of homophily can be varied. At full homophily, we generate the attributes based on neighbours, ensuring that the result will be distributed according to a specific distribution. As such, homophily does not skew the distributions of political leanings, analytical thinking and emotion.

4 Findings

4.1 News Spreading over Time

In this section, we look at the effect that personality traits have on the spread of news. We can consider true news, fake news, and all news. We found that personality traits have some limited effect on the spread of news.

There are various ways to measure the spread of news e.g. how often news is shared. It is possible for some news to spread via a few influential nodes to a disproportionately large audience. As such, the average size of the audience may be a good measure. In practice, a lot of news either fails to reach a significant audience, or it plateaus at a point where nearly everyone has seen it. We refer to the later as *going viral*, and we can have a look at what proportion of news goes viral as a measure of spread too. We present three measures in Figures 1-4, discussed below.

The graphs in Figures 1-4 are plotted against time steps. Agents in the system can act (e.g. view and/or share news) every time step. After enough time steps, the system stops. We have chosen to cut off after 400 time steps, as most things happen in 300 time steps.

In Figures 1, 2 and 4, each data point is the average of 50 runs with 200 true news articles and 100 fake news articles each (so 15,000 news articles of which 10,000 true and 5,000 fake). As a result lines have little random variation. In Figure 3, however, each datapoint represents the average of 50 runs, of the proportion of news articles that goes viral. The effective sample size is much smaller, and thus it has some random variation.

Figures 1-3 are based on random Barabasi-Albert networks. We have run the same experiments on Connected Watts-Strogatz as well, and the results are similar, even if the exact numbers differ a bit. To illustrate, we present the statistics for the number of views in Figure 4. We focus on BA topology only, for brevity sake.

Each of the graphs in Figures 1-4 has 2 plots, a gray (solid) plot, and a (black) dashed plot. The difference is in whether *personality traits* are used in the simulation. The gray plot does not use personality traits, and the dashed plot does. The pairs of plots in Figures 1 and 2 are very similar. But in subfigures (c), there is a noticeable, albeit minor difference; this difference is not random. In Figure 3, the difference between the pairs of plots is larger, but there is randomness in the plots. The difference in Figure 3(c) is not caused by randomness. Personality traits do have a noticeable effect (12% difference) on the rate with which fake news is going viral. If an estimation error of 12% for a simulation or data analysis is not acceptable, then the model must take into account personality traits.

4.2 Personality Traits

In this section, we have a further look at the effect of personality traits. Recall that we can gradually adjust the effect of personality traits from none to full. We refer to this as the personality factor P. Recall that the personality factor only affects how much impact the personality traits have. The distribution of the other

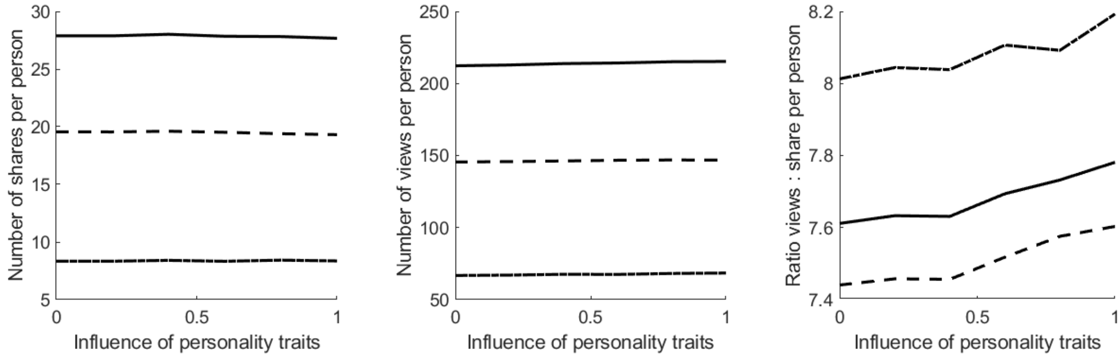

(a) Shares per person (b) Views per person (c) Ratio views per share

Figure 5: Plots of shares/views relative to the influence of personality traits.

variables (share tendency; time and frequency of Facebook use; and the way news is assessed) remains the same – but loses its correlation with personality as P goes to 0. Since these distributions remain constant over P, any effect of adjusting P on the spread can be expected to be subtle.

Here, we have run 100 runs, for the values $P \in \{0, 0.2, 0.4, 0.6, 0.8, 1\}$, each data point is the mean amount of shares/views per person of any/true/fake news. There are 1000 people to take the average from. However, if one news article more or less goes viral, this will result in 100's of people viewing/sharing the article. The average of the 1000 participants is therefore relatively volatile. Since we expect any effect to be subtle, we need the larger quantity of runs to get accurate results. The current amount of runs (100) keeps fluctuations well under 1%, anything over that is significant as a result. The solid plots represent the total number of shares/views/ratios per person, the dashed plots of true news, and the dash-dotted plots of fake news.

From Figure 5, "subtle effects" is an understatement. Figure 5(a) represents the sharing of news per person. Figure 5(a) looks flat, but all three plots have a slight decrease. Only the plot for true news is (barely) significant, decreasing by 0.24 shares of true news per person, from 19.55 to 19.31. Personality traits cause a 1.2% decrease in true news shares. Figure 5(b) represents the views of news per person. Figure 5(b) has a marginally visible increase. The plots for all news and for fake news are significant – again barely. The views of fake news increases 2.48% (from 66.77 to 68.47); the views of all news increases 1.41% (from 212.26 to 215.28). It is interesting to note that the shares tend to decrease whereas views tend to increase.

We created Figure 5(c) to more clearly represent this diverging trend, by plotting the number of views per numbers of shares (per person). With visual inspection, it's clear that there is a statistically significant trend upwards, of around 2.2% per plot. Since shares are decreasing, and views are increasing, it is fair to say that the model is appropriately balanced. Meaning that increases or decreases cannot be explained by other variables being skewed when adjusting personality traits. The only reasonable explanation for the increase in ratio, is the effect of personality traits having a positive impact on that ratio.

The effect of personality traits exists – i.e. there are statistically significant effects. But the effects are minimal. For a basic model of social networks, arguably, a 2.2% error is acceptable. However, there may be hidden effects of personality traits that do not directly affect the spread of fake news, but that may interfere constructively or destructively with other mechanisms or phenomena. As we can see in the next two sections.

4.3 Analytic Thinking

Here, we discuss how analytic thinking (AT) interacts with personality and homophily. Overall, it can be seen that low AT is highly related to greater sharing and viewing of fake news but no decrease in that of true news. Low AT individuals are prone to believe articles with little credibility. Personality variation, however, primarily affects frequency of use, session length, sharing frequency and sensitivity to emotional news.

In figure 6, four investigations are presented: with personality traits turned on/off and with homophily turned on/off. The personality traits are controlled by the personality factor P. Homophily controls whether personal variations between neighbours correlate or not. In every investigation, the weighting factor on analytic thinking was varied from 0 to 1. As the AT weighting decreases, both average AT and AT variation per user decrease. As the weighting factor increases, average AT and AT variation grow. The model performed 300 runs per weighting and the average taken to reduce the impact of random factors.

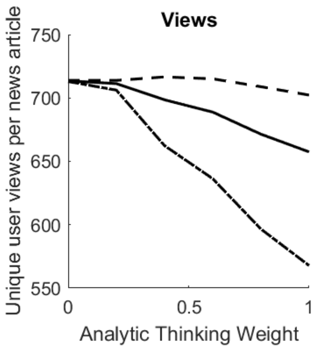

(a) Personality off, homophily off

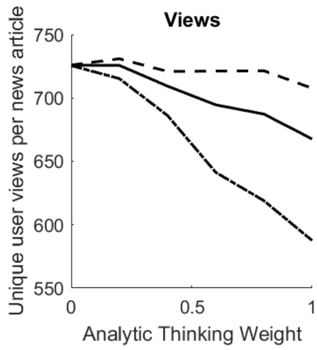

(b) Personality on, homophily off

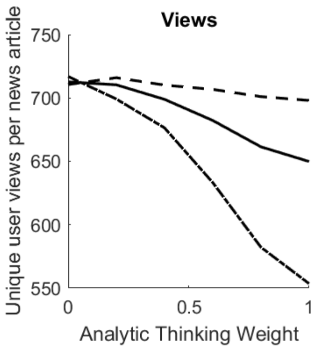

(c) Personality off, homophily on

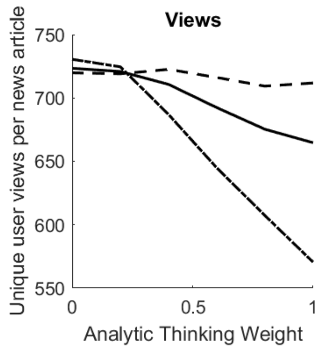

(d) Personality on, homophily on

Figure 6: Plots of how often news is viewed over the degree of analytic thinking on CWS graphs.

It can be seen that true news views remain relatively high throughout. This is because the bottleneck for sharing true news is not related to credibility but instead to emotional engagement or political leaning. This is in line with prior studies [PR19b, SCM20, VRS18]. In Figure 6, a dashed line is true news, a dash-dotted line is fake news and a solid line is the average. The x-axis is the amount of analytic thinking that average people are assigned in a simulation. On the left hand side, few people are analytic thinkers, and on the right side, nearly all people are analytical thinkers. The y-axis is the number of unique readers of an article on average.

Figure 6(a): A user's ability to assess news as false is dependent solely on AT, with low AT users being predisposed to believe all news. Both the high AT weight and low AT weight systems share and view true news at a high frequency. However, as AT weights decrease, users become increasingly unable to distinguish low-believability fake news from high-believability true news. As such, the sharing of fake news rises somewhat above that of true news (as fake news is more likely to be emotionally engaging) and the viewership of fake news rises to meet that of true news.

Figure 6(b): Here, the impact of personality traits (and therefore the importance of their inclusion in such models) is highlighted. As each personality trait only affects a well-defined set of system variables, the traits that play a crucial role can be inferred from the system behaviour.

Firstly, it can be seen that the sharing of fake news has increased at all AT weights. The directly relevant factors (other than AT) are user sharing frequency, user connectivity and news article emotionality. These are related only to the extroversion and neuroticism traits. Extroversion contributes to both the sharing frequency *and* the connectivity of a user. Highly extroverted users will share more and have more followers. Neuroticism contributes positively to both sharing frequency and sensitivity to emotional news. At low AT weights, the average user becomes poor at distinguishing fake and true news. However, users high in neuroticism will respond to the emotionality of fake news and share it with an increased frequency.

Figure 6(c): Homophily can present itself through grouping high/low AT individuals, left/right individuals and high/low emotion individuals. With respect to the impact on this graph, it seems that grouping by AT seems to have the largest impact: fake news shares at high AT weights are lower than for when homophily is not present. Firstly, AT is distributed along a bell curve, so the majority of the network has a relatively high AT. Secondly, due to homophily, these high AT agents are also more likely to connect. Therefore, the core cluster(s) of the network will be mostly composed of high AT individuals who are relatively good at not sharing fake news

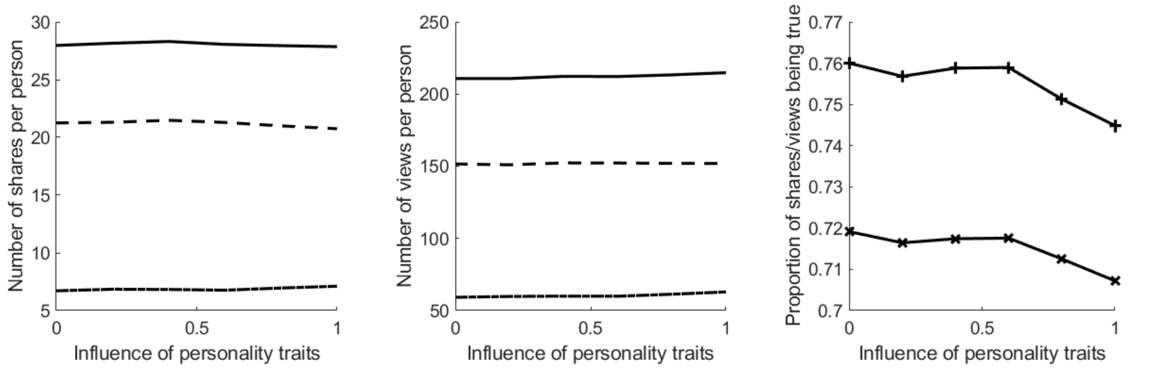

(a) Shares per person (b) Views per person (c) Prop. of views/shares being true

Figure 7: Plots of shares/views relative to the influence of personality traits.

Fake shares will likely occur in smaller fringe clusters of lower AT individuals. They are unlikely to spread into the core cluster and so are unlikely to be picked up by other lower AT clusters, further lowering the spread of fake news. The effect of AT on fake news in Figure 6(c) is therefore exaggerated compared to 6(a).

Figure 6(d): For this figure, the same effects as in Figure 6(b) and Figure 6(c) apply, but they are combined. Notice the effect of homophily on the influence of AT on the spread of fake news is clearly the strongest effect.

Comparing the pairs of graphs of personality traits, we see that the effect of personality traits is minimal, although it is statistically significant. Comparing the pairs of graphs of homophily, we see that the effect is more pronounced for fake news. The effect is still limited in scope.

4.4 Fact Checking News

In this section, we implement a fact-checking mechanism. If the news is true, then users will be informed that the news is verified true. If the news is fake, then users will be warned that the news is verified false. Of course, this only accomplishes that users' beliefs may shift, but not necessarily to 100% or 0%. For most people, the belief will increase if an article is verified to be true, and decrease if it is verified to be fake. But for some people, the opposite may be true.

If we consider personality traits, then personality traits will influence the probability of believing and valuing the fact-checker's judgement. Our implementation is that conscientiousness alone determines the direction and magnitude of the effect of the fact-checker's judgement. If personality traits are off, a similarly distributed random value will take its place. The mean and variance of the impact of fact-checkers remains the same, for all personality factors $0 \leq P \leq 1$. The belief in an article before fact-checking is used as a baseline, and the belief will shift according to the impact of the fact check.

Graphs (a) and (b) in Figure 7 are the same as in Figure 5, but with the fact-checking mechanism turned on. There are clearly somewhat more shares/views of true news, and somewhat less shares/views of fake news. This is the intended effect of fact-checking. There are various reasons why the effect is modest: first, there are other factors such as emotional engagement or political leaning at play, and second, the impact of fact-checking may be small or negative for some people. As before, the influence of personality traits is plotted on the x-axis.

The influence of personality traits is different when fact-checking is introduced; compare Figures 5 and 7. In Figure 7, the true shares decreases by 2.3% (from 21.25 to 20.75), but the fake shares increases by 5.6% (from 6.71 to 7.11). For views, in Figure 7(b), true news does not change significantly, and fake news is seen 5.9% more. In both cases, introducing personality traits lowers the proportion of true news versus fake news.

We plot this explicitly in Figure 7(c). The line with + is the proportion of shares that are of true news, and the line with $\times$ is the proportion of views. Recalling that fact-checking was a measure to raise the proportion of true news being spread, it shows that personality traits are lowering the effectiveness of the fact-checking. The effectiveness is lowered by a percentage point. However, the total effectiveness of the fact-checking is only about 5 percentage points. That is a fairly big relative difference.

5 Conclusion

We built an agent-based simulation of the spread of fake news. The agents were modelled to have personality traits and homophilic distribution of attributes. The numbers used to model personality traits are based on psychology research. The simulation provides results that are in line with the findings in the literature.

Our simulation has the ability to turn personality traits and homophily off, without affecting the general distribution of all the characteristics of the people. We have looked at the effects of doing to on the spread of news – and fake news in particular. It turns out that there is a statistically significant effect in all configurations we have tried. But that the effect is limited in size. In particular, personality traits increase the number of views per share, irrespective of whether the news is true or not. Homophily exaggerates the effects of analytic thinking on the spread of fake news. When implementing a fact-checker, the personality traits of people work *against* the countermeasures. Not modelling personality can lead a specialist to overestimate the effectiveness of fact-checking.

References

[AM15] S Alhabash and A McAllister. Redefining virality in less broad strokes: Predicting viral behavioural intentions from motivations and uses of facebook and twitter. *SAGE New Media and Society*, 17:1317–1339, 9 2015.

[AMS13] Sinan Aral, Lev Muchnik, and Arun Sundararajan. Engineering social contagions: Optimal network seeding in the presence of homophily. *Network Science*, 1(2):125–153, 2013.

[AR20] Shamshad Ahmed and Tariq Rasheed. Relationship between personality traits and digital literacy skills: a study of university librarians. *Digital Library Perspectives*, 2020.

[AV10] Y Amichai-Hamburger and G Vinitzky. Social network use and personality. *Computers in human behaviour*, 26(6):1289–1295, 2010.

[BFP11] D. Blanco-Moreno, R. Fuentes-Fernández, and J. Pavón. Simulation of online social networks with krowdix. In *2011 International Conference on Computational Aspects of Social Networks (CASoN)*, pages 13–18, 2011.

[CBK04] Jennifer S Clifford, Magdalen M Boufal, and John E Kurtz. Personality traits and critical thinking skills in college students: Empirical tests of a two-factor theory. *Assessment*, 11(2):169–176, 2004.

[CCT17] Alina Campan, Alfredo Cuzzocrea, and Traian Marius Truta. Fighting fake news spread in online social networks: Actual trends and future research directions. In *2017 IEEE International Conference on Big Data (Big Data)*, pages 4453–4457. IEEE, 2017.

[CCTS14] B. Caci, M. Cardaci, M.E. Tabacchi, and F. Scrima. Personality variables as predictors of facebook usage. *SAGE Psychological Reports*, 114:528–539, 04 2014.

[Cho19] N Chokshi. Older people shared fake news on facebook more than others in 2016 race, study says. `https://www.nytimes.com/2019/01/10/us/politics/facebook-fake-news-2016-election.html`, 2019.

[CM99] Paul T Costa and Robert R McCrae. A five-factor theory of personality. *Handbook of personality, 2nd edn. Guilford Press, New York*, pages 139–153, 1999.

[CMP19] Vincenzo Carrieri, Leonardo Madio, and Francesco Principe. Vaccine hesitancy and (fake) news: Quasi-experimental evidence from italy. *Health economics*, 28(11):1377–1382, 2019.

[DJH+16] Morteza Dehghani, Kate Johnson, Joe Hoover, Eyal Sagi, Justin Garten, Niki Jitendra Parmar, Stephen Vaisey, Rumen Iliev, and Jesse Graham. Purity homophily in social networks. *Journal of Experimental Psychology: General*, 145(3):366, 2016.

[Gol93] L Goldberg. The structure of phenotypic personality traits. *American Psychologist*, 48:26–34, 1993.

[GRG+20] Anastasia Giachanou, Esteban A Ríssola, Bilal Ghanem, Fabio Crestani, and Paolo Rosso. The role of personality and linguistic patterns in discriminating between fake news spreaders and fact checkers. In *International Conference on Applications of Natural Language to Information Systems*, pages 181–192. Springer, 2020.

[Hot16] Peter J Hotez. Texas and its measles epidemics. *PLoS medicine*, 13(10):e1002153, 2016.

[JDS+13] F. Jin, E. Dougherty, P. Saraf, Y. Cao, and N. Ramakrishnan. Epidemiological modelling of news and rumors on twitter. In *Proceedings of the 7th Workshop on Social Network Mining and Analysis*, 08 2013.

[MP12] S Mamedova and E Pawlowski. A description of u.s. adults who are not digitally literate. *National Centre for Education Statistics*, page 9, 2012.

[MSLC01] Miller McPherson, Lynn Smith-Lovin, and James M Cook. Birds of a feather: Homophily in social networks. *Annual review of sociology*, 27(1):415–444, 2001.

[New18] Mark Newman. *Networks*. Oxford University Press, 2 edition, 2018.

[Ofc19] Ofcom. Half of people now get their news from social media. `https://www.ofcom.org.uk/about-ofcom/latest/features-and-news/half-of-people-get-news-from-social-media`, 2019.

[ORR12] J Ormel, H Riese, and J Rosmalen. Interpreting neuroticism scores across the adult life course: immutable or experience-dependent set points of negative affect? *Clinical Psychology Review*, 32(1):71–79, 2012.

[PCR18] Gordon Pennycook, Tyrone D Cannon, and David G Rand. Prior exposure increases perceived accuracy of fake news. *Journal of experimental psychology: general*, 147(12):1865, 2018.

[PR19a] G Pennycook and D Rand. Fighting misinformation on social media using crowdsourced judgments of news source quality. *Proceedings of the National Academy of Sciences*, 116(7):2521–2526, 2 2019.

[PR19b] G Pennycook and D Rand. Lazy, not biased: Susceptibility to partisan fake news is better explained by lack of reasoning than by motivated reasoning. *Cognition*, 188:39–50, 7 2019.

[RJL15] P Rentfrow, M Jokela, and M Lamb. Regional personality differences in great britain. *PloS one*, 10(3):e0122245, 2015.

[ROS+09] Craig Ross, Emily S Orr, Mia Sisic, Jaime M Arseneault, Mary G Simmering, and R Robert Orr. Personality and motivations associated with facebook use. *Computers in human behavior*, 25(2):578–586, 2009.

[SCM20] C Sindermann, A Cooper, and C Montag. A short review on susceptibility to falling for fake political news. *Current Opinion in Psychology*, 2020.

[SG19] Elisa Shearer and Elizabeth Grieco. Americans are wary of the role social media sites play in delivering the news. *Pew Research Center*, 2, 2019.

[VRA18] Soroush Vosoughi, Deb Roy, and Sinan Aral. The spread of true and false news online. *Science*, 359(6380):1146–1151, 2018.

[VRS18] S. Vosoughi, D. Roy, and S.Aral. The spread of true and false news online. *Science*, 359:1146–1151, 03 2018.

[WKWK18] Przemyslaw M Waszak, Wioleta Kasprzycka-Waszak, and Alicja Kubanek. The spread of medical fake news in social media–the pilot quantitative study. *Health policy and technology*, 7(2):115–118, 2018.

[Yam20] Metin Yaman. Examining media literacy levels and personality traits of physical education and sports students according to certain demographic variables. *Turkish Online Journal of Educational Technology-TOJET*, 19(1):1–8, 2020.