Written evidence submitted by Dr. Ansgar Koene, Horizon Digital Economy Research Institute, University of Nottingham (FNW0116)

1.   Horizon[1] is a Research Institute at The University of Nottingham and a Research Hub within the RCUK Digital Economy programme[2]. Horizon brings together researchers from a broad range of disciplines to investigate the opportunities and challenges arising from the increased use of digital technology in our everyday lives. Dr. Koene is a Senior Research Fellow at Horizon and is co-investigator on the EPSRC funded UnBias[3] project within Horizon which is studying issues related to non-operationally justified bias in algorithmic systems that control access to information online (e.g. search engines, recommender systems, news feeds). Dr. Koene conducts research as part of the UnBias project. An important part of this work includes the facilitation of multi-stakeholder workshops with industry, civil-society organizations, academics and teachers designed to identify experiences, concerns and recommendations information mediating algorithms. Dr. Koene is chair of the IEEE working group for the development of a Standards on Algorithm Bias Considerations[4], and member of the Internet Society (ISOC UK[5]). I would be willing to give verbal evidence if so desired.

Questions

1. What is 'fake news'? Where does biased but legitimate commentary shade into propaganda and lies?

2.   The term 'fake news' as it is currently used in the media is somewhat problematic due to the wide range of content that is covered by it. The article "Fake news. It's complicated" in FirstDraftnews by Claire Wardle[6] presented a useful classification of mis- and disinformation types:

- Satire or Parody (No intention to cause harm but has potential to fool)
- False Connection (When headlines, visuals or captions don't support the content)
- Misleading Content (Misleading use of information to frame an issue of individual)
- False Context (When genuine content is shared with false contextual information)
- Imposter Content (When genuine sources are impersonated)
- Manipulated Content (When genuine information or imagery is manipulated to deceive)
- Fabricated Content (News content that is 100% false, designed to deceive and do harm)

3.   Each of these types of mis-/disinformation can be the results of a range of causes and motivations, which Claire Wardle summarized in the following grid table:

	Satire	False Connection	Misleading Content	False Context	Imposter Content	Manipulated Content	Fabricated Content
Poor journalism		X	X	X			
To Parody	X				X		X
To Provoke					X	X	X
Passion				X			
Partisanship			X	X			
Profit		X			X		X

Political Influence			X	X		X	X
Propaganda			X	X	X	X	X

2. What impact has fake news had on public understanding of the world, and also on the public response to traditional journalism? If all views are equally valid, does objectivity and balance lose all value?

4.   Possible impacts of fake news on public understanding, and subsequent responses, fall into two general categories:

   i. Belief in a false claim, potentially resulting in sizable sectors of the population engaging in behaviours that have negative consequences (e.g. MMR scare based on false reporting about research into links between the MMR vaccine and autism).
   ii. A general distrust of all news sources resulting in disengagement/cynicism (e.g. not voting at elections because you can't trust any of 'them'), which reinforces confirmation bias (only believing news that matches previously held beliefs)

5.   Quantitative data on the extent to each of these types of impacts on the population in the UK, or elsewhere, is currently scarce. The 2016 Digital News Report by the University of Oxford Reuters Institute for the Study of Journalism[7] reported the following results from a survey people's trust in the news where people were asked:

"Do you agree that you can trust most news most of the time?" 4% Strongly agree; 46% Tend to agree; 30% Neither; 17% Tend to disagree; 4% Strongly disagree.

It is not clear to what extent any of these results can be linked to awareness or concerns about fake news.

6.   Promotion of a relativistic approach to the truth in which all views are treated equally, by assigning equal 'air-time' them, contributes to the confusion over the trustworthiness of news. Open and balanced debate is necessary, but relevant expertise such as first-hand experience of an issue, must be given its due weight. If a stated claim challenges the accepted views of a whole profession, e.g. MMR Vaccine scare vs. Doctors; Climate change denial vs. climate scientists) the extraordinary nature of the claim must demand an equally extraordinary level of evidence to warrant equal treatment.

7.   The current trend within the media, unfortunately tends towards giving extraordinary claims extra attention (which can easily be confused as a sign of credibility) simply for the sensation it can cause.

3. Is there any difference in the way people of different ages, social backgrounds, genders etc. use and respond to fake news?

8.   To our knowledge there is as yet little evidence to suggest systematic differences in the way different population groups respond to fake news in general. On specific topics we can expect that responses will tend to depend on the level of experience/personal knowledge that exists within the population group.

9.   The 2016 Digital News Report by the Reuters Institute for the Study of Journalism[8] reported a slight difference in trust in the news among under 35s (42% trust) compared to over 35s (52% trust), but again it is not clear to what extent this is related to issues of fake news.

4. Have changes in the selling and placing of advertising encouraged the growth of fake news, for example by making it profitable to use fake news to attract more hits to websites, and thus more income from advertisers?

10. The advertising based business model of most online media platforms has shaped both the development of the platform interfaces and the content.

11. Site interfaces (layout, available tools) are optimized to increase 'engagement', as measured by clicking the 'like' (or equivalent) button and sharing the content with people in the user's network, since this is the basis for revenue from advertisers. To maximizing such 'engagement' it is in the interest of the platforms to make it easy for people to 'like' or share content as quickly as possible, often before the article has been read (other than the headline), and users have not taken the time to assess the validity of the content.

12. These properties of online media platforms tend to disproportionately stimulate the spreading of 'fake news', compared to genuine reporting, since fake news tends to be more sensational. The drive for sensationalism as means to capture audience attention also distorts genuine news reports by encouraging exaggeration and lack of context in order to keep the message short.

5. What responsibilities do search engines and social media platforms have, particularly those which are accessible to young people? Is it viable to use computer-generated algorithms to root out 'fake news' from genuine reporting?

13. A key problem with social media is confusion about the origin of a news story. Social media emphasises the person who shared it, the friend, over the organization/news site that published it. People judge trustworthiness based on the trust they have in their friend, not the reputation of the news source. As soon as one friend has hastily 'liked' an unreliable story the story spreads among their network of friends on the basis of personal trust.

14. Media brand used to be a 'trust mark'. Now the brand is hidden as much as possible. Social media want to users to associate content with their platforms, not go to the website of the media organization, since time spend on the platform translates into (potential) advertising revenues.

15. Due to the use of personalization algorithms to select which content to push to the user, social media platforms are increasingly engaged with editorial behaviours, in the sense of choosing which news stories go on the 'front page'.

16. One concern is the way in which these personalization algorithms amplify and reinforce the natural tendency of platform users to prioritise/engage with content that matches their existing preferences/world-view, the so called 'filter bubble' or 'echo chamber' effect.

17. Research from Facebook[9] has shown that their algorithms do have a significant effect in increasing the filter bubble effect, even if they maintain that the effect of natural human selection bias is even greater. The latter, does not negate the former, nor is it completely clear to which extent the mutually reinforcing feedback loop between the two effects is amplifying both.

18. During discussions with a German consulting company that has been doing interviews with users of social media platforms, we were told that many users are of two minds regarding personalization algorithms. They appreciate the convenience of personalized content regarding product recommendations and commercial content but felt uncomfortable with the use of personalization regarding news content.

19. Another form of algorithmic systems that affect the spread of fake news us the use of bots (software robots) to 'game the system'. There are a number of ways in which bots are used:

20. Bots are used for 'astroturf' campaigns on social media platforms, where they automatically produce a huge volume of activity/posts related to a topic (policy/product) in order to capture attention and make it the 'meme of the day'. Political astroturf are political campaigns made to look like spontaneous "grassroots" activity – to get people to engage who would otherwise mistrust political campaigns[10].

21. Bots can increase the chances that the news recommender systems/personalization algorithms on social media platforms will push a story to users, by repeatedly 'liking' and 'sharing' the fake news content among bot controlled accounts on the social media platform. This method exploits the fact that recommendation algorithms use the number of 'likes' and 'shares' of content as a measure of popularity, and hence 'importance to show to the user'.

22. Bots can use the comment feature on social media platforms to generate supporting comments to 'fake news' stories and critical/doubting comments to real news, that is negative to the desired narrative of the organization controlling the bots.

23. Dr. Ferrara (University of Southern California), among others, are working to try to identify Bots that are feeding social media, however it is an arms race between changes in bot behaviour and bot detection[11].

24. Automated, algorithmic detection of fake news will be very difficult to achieve and probably easy to fool.

25. As indicated in our answer to Q1, fake news comprises of a wide range of types many of which require some form of context understanding, which is still very difficult to achieve through automated systems. Probably the best one could achieve is to detect stories that are based on fabricated content, since these might fail to have corroborating information elsewhere.

26. Such algorithms however may be very easy to fool by using bots to rapidly create 'corroborating' content on websites.

27. A similar problem can be expected for systems such as the 'fake news' reporting button that Facebook has introduced. I would not be surprised if there will rapidly be bots that use these reporting function to flag fact checker websites as fake news.

28. Any method that relies on crowdsourced voting of content reliability will have to find a way to counteract up-/down-voting by bots.

29. One of the simplest, yet possibly most effective, changes to social media platforms that could reduce the spread of fake news is probably to introduce a simple delayed access to 'like'/'share' buttons (at least for news content). The platforms could change their 'like'/'share' buttons so that they only become active after the user has spent enough time seeing the actual full story, such that they can reasonably be expected to have read and assessed the validity of the content. The main obstacle to this approach is the business model of the platforms.

6. How can we educate people in how to assess and use different sources of news?

30. A study by the Stanford History Education Group on "Evaluating Information: the cornerstone of civic online reasoning"[12] tested the ability of US primary, middle and high school children to distinguish between advertising, sponsored content (a form of advertising) and real articles on online media such as 'Slate' magazine. Out of 203 middle school students they found that three-quarters were able to correctly identify traditional advertisement and real news stories correctly, however more than 80% believed that native advertisements, identified by words like "sponsored content" in the box of the story, was a real news story.

31. It should be noted that the above results were in the context of a study where students were explicitly challenged to think about whether not the story was real news or advertising.

32. Often one of the first steps people will take to try and check it a news story is real, is to put key phrases from the story into a search engine to 'google' for sites that might confirm or contradict the story[13]. One of the problem with this approach is that the way in which the issues is phrased will affect the search results. Search engines will naturally prioritise results that use the same language. Websites that support the fake news, often created by the same groups that

created the fake news, will use the same was to phrase the issue. A clear example of this is the case of the Holocaust denial stories. Holocaust denial website tend to use the word 'holocaust' far more frequently than do history websites that provide a thorough explanation of the atrocities of the Nazi regime.

33. Better education on how to critically evaluate online content is clearly necessary. However, since there is no evidence that adults are any better at distinguishing real from fake news, it is Important to develop forms of civic 'critical reading' education suitable for each age group.

7. Are there differences between the UK and other countries in the degree to which people accept 'fake news', given our tradition of public service broadcasting and newspaper readership?

34. There is no evidence that the UK population is more resilient/discerning against fake news than any other people.

35. The tradition of public service broadcasting with limited/no advertising and strong oversight against partisanship means people have less experience /training in dealing with these issues from broadcasting.

36. Within the print media, however, the UK has a tradition of strongly partisan newspapers. In contrast to online content however the partisan bias of the print media brands is generally well known to readers, allowing them to take this into consideration when consuming the news.

March 2017

[1] http://www.horizon.ac.uk

[2] https://www.epsrc.ac.uk/links/councils/research-councils-uk-rcuk/digital-economy-research-rcuk/

[3] http://unbias.wp.horizon.ac.uk

[4] https://standards.ieee.org/develop/project/7003.html

[5] http://isoc-e.org/

[6] https://firstdraftnews.com/fake-news-complicated/

[7] http://www.digitalnewsreport.org/survey/2016/trust-in-the-news-2016/

[8] http://www.digitalnewsreport.org/survey/2016/trust-in-the-news-2016/

[9] https://solomonmessing.wordpress.com/2015/05/24/exposure-to-ideologically-diverse-news-and-opinion-future-research/

[10] https://www.aaai.org/ocs/index.php/ICWSM/ICWSM11/paper/view/2850/3274

[11] https://dl.acm.org/citation.cfm?doid=2963119.2818717

[12] https://sheg.stanford.edu/upload/V3LessonPlans/Executive Summary 11.21.16.pdf

[13] http://www.edweek.org/ew/articles/2016/11/02/why-students-cant-google-their-way-to.html