Misinformation in the Digital Age: Comparative Contexts and Policy Lessons Conference

May 5-6, 2022

Titles and Abstracts

Alberto ACERBI, Brunel University London

Misinformation: the long view

Abstract:
In my talk, I will explore how cultural evolution can provide a convenient framework to understand how information is produced, transmitted, and selected in contemporary digital media.
Within this perspective, online misinformation is an aspect of a more general phenomenon, where some cultural traits can be successful because their content taps into widespread cognitive biases. Misinformation, being less constrained by reality than true information, can be manufactured to appeal to these cognitive biases. As such, online misinformation can be characterised not as low-quality information that spreads because of the inefficiency of online communication, but as high-quality information that spreads because of its efficiency. The difference is that `quality` does not denote truthfulness but psychological appeal.
This is one of the reasons why we should expect some, limited, level of misinformation in any communication system. I will discuss how several estimates suggest indeed a relatively low prevalence of misinformation in social media, and how a broad focus on the overall condition of the online information ecosystem could be more advantageous than a narrower one on misinformation.

Bence BAGO, IAST

Why do people disbelieve in climate change?

Abstract:
A widely-held explanation involves politically motivated reasoning: Rather than helping uncover truth, people use their reasoning abilities to protect their partisan identities and reject beliefs that threaten those identities. Despite the popularity of this account, the evidence supporting it (i) does not account for the fact that partisanship is confounded with prior beliefs about the world, and (ii) is entirely correlational with respect to the effect of reasoning. Here, we address these shortcomings by (i) measuring prior beliefs and (ii) experimentally manipulating participants' extent of reasoning using cognitive load and time pressure while they evaluate arguments for or against anthropogenic global warming. The results challenge the politically motivated system 2 reasoning account: engaging in more reasoning led people to have greater coherence between judgments and their prior beliefs about climate change - a process that can be consistent with Bayesian reasoning.

Nadia BRASHIER, Purdue University

Timing Matters When Correcting Fake News

Abstract:
One proposed solution to the misinformation crisis involves flagging misleading content. But people often continue to rely on falsehoods, even after receiving explicit corrections. We tested whether this *continued influence effect* depends on when people receive fact-checks. In two experiments (total *N* = 2,683), participants read true and false headlines taken from social media. In the treatment conditions, "true" and "false" tags appeared before, during, or after participants read each headline. In a control condition, participants received no information about veracity. One week later, all participants rated the same headlines' accuracy. Providing fact-checks after headlines improved subsequent truth discernment more than providing the same information during or before exposure. We recently replicated this finding in a follow-up study where participants (*N* = 1215) completed a more naturalistic task at exposure. Our results inform both the cognitive science of belief revision and social media platform design.

Young Mie KIM, University of Wiconsin-Madison

Digital Disinformation and Election Integrity: Research Problems, Empirical Evidence and Policy Implications

Abstract:
As indicated in the Jan 6 Capitol riot, targeted voter suppression campaigns, election interference and such, the integrity of elections appears to be deeply challenged. Furthermore, recent public debates suggests that social media play a role in such tragic incidents by spreading misleading information and influencing election outcomes. Despite the mounting concerns, relatively little systematic empirical research exists, partly due to some methodological challenges. By utilizing a user-based, real-time online behavior tracking tool and "reverse engineering" techniques, Kim's research empirically evidences digital disinformation on social media, independent from tech platform companies. By combining computational approaches with survey-based approaches, Kim's research investigates the correlates between disinformation exposure, political dispositions, and beliefs of election legitimacy and assesses the effects of digital disinformation on election outcomes such as voter turnout. The talk offers offers insight relevant for regulatory policies and discusses the normative implications for the functioning of democracy.

Stephan LEWANDOWSKY, University of Bristol

Microtargeting: Efficacy, Ethics, Public Opinion, and Boosts

Abstract:
There has been much concern about the "microtargeting" of political messages at individuals on social media based on sometimes sensitive personal characteristics that are inferred by the platforms from mundane data and activities. Evidence suggests that this type of microtargeted advertising, for example based on recipients' personality, can be effective. I review the ethical and political implications of microtargeting and show that public opinion generally opposes the use of sensitive data for targeting. I then present ways in which the public might be protected against microtargeting, for example by "boosting" their ability to detect when they are targeted or by providing machine-learning-based tools that alert users when they are targeted.

Jason REIFLER, University of Exeter

Free speech vs. harmful misinformation: Moral dilemmas in online content moderation
(joint with Anastasia Kozyreva, Stefan M. Herzog, Stephan Lewandowsky, Ralph Hertwig, Philipp Lorenz-Spreen, Mark Leiser, and Jason Reifler)

Abstract:
Content moderators face moral dilemmas between two values — freedom of expression and preventing harm from false information — when they decide whether to remove posts containing dangerous misinformation and whether to suspend accounts that spread such falsehoods. In this study, we used a conjoint survey experiment to explore how people (U.S. respondents,N=2564) approach these moral dilemmas across four misinformation topics: politics (election denial), health (COVID-19 anti-vaccination), history (Holocaust denial), and environment (climate change denial). We found that severity of consequences and the repeated offense had the strongest effects on respondents' decision to remove the posts or to suspend accounts. Type of misinformation shared was also important: Climate change denial was censored the least, whereas Holocaust denial and Election denial elicited the most willingness to censor. In the dilemma between protecting free speech and censoring harmful misinformation, Democrats showed stronger preference for preventing dangerous falsehoods across all 4 scenarios, whereas Republicans sided more with free speech protection and made fewer censorship decisions. Our findings show that features related to the account itself (person behind the account, their partisanship, and number of followers) had no to little effects on censorship decisions. Interestingly, account's partisanship also did not matter along partisan lines. However, Republicans and proponents of freedom of expression over mitigating harmful misinformation were more likely to remove the posts by accounts with lower number of followers, whereas Democrats and proponents of mitigation, penalized accounts with higher number of followers more.

Leah ROSENZWEIG, University of Chicago
(joint with Molly Offer-Westport and Susan Athey)

Optimal Policies to Battle the Coronavirus "Infodemic" among Social Media Users in Sub-Saharan Africa

Abstract:
Alongside the outbreak of the novel coronavirus, an "infodemic" of myths and hoax cures has spread over online media outlets and social media platforms. Building on the literature combating misinformation, we evaluate experimental interventions designed to decrease sharing of false COVID-19 cures. Using Facebook advertisements to recruit social media users in Kenya and Nigeria, we deliver our interventions with a Messenger chatbot, facilitating observation of treatment effects in a realistic setting. We use a contextual adaptive experimental design to target the most effective interventions, and learn and evaluate a contextual policy, improving our understanding of how to tackle harmful misinformation during an ongoing health crisis. This paper brings comparative data to a global problem for which the existing research has largely been focused on the U.S. and Europe.

Alexander STEWART, University of St Andrews

The coercive logic of fake news

Abstract:
The spread of misinformation and "fake news" continues to be a major focus of public concern. A great deal of research has examined who falls for misinformation and why, and what can be done to make people more discerning consumers of news. Comparatively little work, however, has considered misinformation producers, and how their strategies interact with the psychology of news consumers. Here we use game-theoretic models to study the strategic interaction between news publishers and newsreaders. We show that publishers who seek to spread misinformation can generate high engagement with falsehoods by using strategies that mix true and false stories over time, in such a way that they serve more false stories to more loyal readers. These coercive strategies cause false stories to receive higher reader engagement than true stories – even when readers strictly prefer truth over falsehood. In contrast, publishers who seek to promote engagement with accurate information will use strategies that generate more engagement with true stories than with false stories. We confirm these predictions empirically by examining 1,000 headlines from 20 mainstream and 20 fake news sites, comparing Facebook engagement data with 20,000 perceived accuracy ratings collected in a survey experiment. We show that engagement is negatively correlated with perceived accuracy among misinformation sites, but positively correlated with perceived accuracy among mainstream sites. We then use our model to analyze the conditions under which news sites seeking engagement will produce false stories. We show that if a publisher incorrectly assumes that readers prefer falsehoods, their resulting publication strategy can itself manufacture greater engagement with false news – leading to a self-reinforcing cycle of false news promotion. In order to effectively understand and combat misinformation, it is not enough to study how people engage with news in isolation - academics and policy makers must consider the strategic interplay of those who produce and consume (mis)information.

Joshua TUCKER, New York University

Do Your Own Research? How Searching Online Can Increases Belief in Misinformation
(joint with Kevin Aslett, Zeve Sanderson, William Godel, Nathaniel Persily, Jonathan Nagler, and Joshua A. Tucker (presenting))

Abstract:
In order to combat phenomena associated with democratic dysfunction linked to exposure to misinformation, significant attention has been paid to understanding the spread of and belief in online misinformation, with a particular focus on social media platforms. However, the dominant role of search engines in the digital information ecosystem remains under-explored, even though the use of online search to explore the veracity of fake or misleading news when encountered by users is a central component of media literacy interventions encouraged by technology companies, government agencies, and civil society organizations alike. Although conventional wisdom suggests that searching online when evaluating the veracity of false news would reduce belief in false news, there is little empirical evidence with which to evaluate its effect. To address this gap, we test the effect of online search on belief in false news stories. Across five experiments, we present consistent evidence that searching for information online to examine the truthfulness of news articles increases the likelihood of believing misinformation, in some cases by up to 26%. To shed light on this relationship, in our fifth study we combine survey and digital trace data, collected using a custom browser extension, to investigate the cause. We find empirical evidence that the high amount of low-quality information returned by search engines when examining the content of misleading or false news stories increases belief in misinformation, and that this effect may be concentrated in individuals with lower levels of digital literacy.

Sander VAN DER LINDEN, Unviersity of Cambridge

Psychological Inoculation Against Misinformation

Abstract:
In order to combat phenomena associated with democratic dysfunction linked to exposure to Much like a viral contagion, false information can spread rapidly from one individual to another. Moreover, once lodged in memory, misinformation is difficult to correct. Inoculation theory therefore offers a natural basis for developing a psychological 'vaccine' against the spread of fake news and misinformation. Specifically, in a series of randomized lab and field studies, we show that it is possible to "immunize" people against disinformation about a wide range of topics by pre-emptively refuting and exposing them to severely weakened doses of the techniques that underlie its production. This process of psychological inoculation or "prebunking" helps people cultivate cognitive antibodies in a simulated social media setting. During the talk, I'll showcase several award-winning real-world interventions we developed and empirically evaluated in 20 languages—with governments, the WHO, the UN, and social media companies—to help citizens around the world recognize and resist unwanted attempts to influence and mislead.