

# **The Future of Fact-Checking**

by

Abigail Thwaites and Eliot Matthew Watkins

## The Future of Fact-Checking

### Summary

The proliferation of misinformation is a profound social concern in the modern age. This paper explores the prospects for a new and emerging technology which aims to mitigate the spread of misinformation. We call these technologies ‘Automated Fact Checkers’ (AFCs). AFCs are AI-driven softwares whose purpose is to fact-check: to assess (with minimal human intervention) the truth of a claim, and remove or flag claims judged to be false. The technology is still in its infancy, but the hope for AFCs is that they will both minimize research burdens on human fact-checkers and be capable of flagging or removing misinformation almost instantaneously – potentially long before it reaches viral status.

In this paper, we argue that the careful deployment of AFCs is likely to be a profound social good. Our case in favor of AFCs is grounded in attention to the distinctive reasons why misinformation is harmful. So, we begin by asking this question: misinformation seems bad, but why? Contrary to popular belief, we argue that misinformation is not objectionable merely in virtue of its risking *physical* harm. Instead, we argue that the central social concerns raised by misinformation are *epistemic* and *political*.

The epistemic concern: misinformation makes it harder for agents to do good practical reasoning; that is, to make decisions which reflect their values and goals. The political concern: misinformation is bad because it makes it easier for vulnerable people to be exploited by those with power. We argue that these social concerns provide us with robust motivations for developing AFCs.

We consider two objections which suggest that AFCs risk *reinforcing* rather than alleviating these epistemic and political concerns. The first concern is that AFCs risk fostering an environment of ‘epistemic complacency’; too much trust in AFCs may weaken our capacities to engage in critical thinking. The second concern is that AFCs may encode human biases in its judgments of veracity and checkworthiness.

We respond to both of these objections optimistically. We offer realistic proposals for how developers may avoid these pitfalls, and suggest that current developments in the technology of AFCs may already be heading in this direction.

## The Future of Fact-Checking

Word count (excluding references and summary): 3000

### §1. Introduction

Consider the following case:

#### *Viral Tweet*

David, a father of two young children, is scrolling through Twitter. The following tweet is retweeted onto his timeline by a source he trusts:

*A friend who works in seat belt development urged me to tell you this; wearing seat belts is DANGEROUS for children. Seat belts are more likely to hurt your child than save them. If you care about your children, don't let them wear seat belts.*

The claim that seat belts are dangerous for children is in fact false, but David reads this tweet and believes it to be true. David cares a lot about his children, and chooses to act on his new belief. He no longer allows his young children to wear seat belts.

Clearly, something very bad has happened in *Viral Tweet*. David has been exposed to, and subsequently misled by, *misinformation* (a claim about the world which is in fact false, but presented as true). The misinformation which David was exposed to has led David to act in such a way that his children suffer an increased risk of serious harm.

*Viral Tweet* is hypothetical, but not unrealistic. Soon after the outbreak of the Covid-19 pandemic, misinformation concerning alleged treatments for coronavirus were circulated online, including

brushing one's teeth with silver-infused toothpaste<sup>1</sup>, which in fact turns one's skin blue<sup>2</sup>. Part of social media platforms' response to cases like these was to employ human fact-checkers, whose task was to manually flag or remove misinformation.<sup>3</sup> However, the strategy of manual human fact-checking continues to prove limited. For one thing, the sheer quantity of misinformation on the internet vastly outweighs the capabilities of human fact-checkers, and what isn't caught often spreads so rapidly that by the time a human fact-checker has removed or flagged a post it may have already reached many thousands of readers.<sup>4,5</sup>

This paper explores the prospects for new and emerging technologies which aim to mitigate the spread of misinformation. We call these technologies 'Automated Fact Checkers' (AFCs). AFCs are AI-driven softwares whose purpose is to fact-check: to assess (with minimal human intervention) the truth of a claim, and remove or flag claims judged to be false. The claims over which AFCs may have jurisdiction include social media posts, blog posts, and news articles. The technology is still in its infancy, but the hope is that AFCs will both minimize research burdens on human fact-checkers and be capable of flagging or removing misinformation almost instantaneously – potentially long before it reaches viral status.<sup>6</sup>

---

<sup>1</sup> "Alex Jones is Told to Stop Selling Sham Anti-Coronavirus Toothpaste", Luis Ferre-Sadurni and Jesse McKinley, The New York Times, March 13th 2020:

<https://www.nytimes.com/2020/03/13/nyregion/alex-jones-coronavirus-cure.html>

<sup>2</sup> NIH National Center for Complementary and Integrative Health: 'Colloidal Silver':

<https://www.nccih.nih.gov/health/colloidal-silver>

<sup>3</sup> "Facebook doubles down on removing coronavirus conspiracy theories", Shirin Ghaffary and Rebecca Heilweil, Vox, March 4th 2020:

<https://www.vox.com/recode/2020/1/31/21115589/coronavirus-wuhan-china-myths-hoaxes-facebook-social-media-tiktok-twitter-wechat>

<sup>4</sup> "People are bad at spotting fake news. Can computer programs do better?", Maria Temming, Science News, July 26th 2018: <https://www.sciencenews.org/article/can-computer-programs-flag-fake-news>

<sup>5</sup> "The Spread of True and False News Online", Soroush Vosoughi, Deb Roy, and Sinan Aral, *Science*, Vol 359, pp 1146-1151

<sup>6</sup> For a state of the art survey article on AFCs, see "A Survey on Automated Fact-Checking", Zhijiang Guo, Micheal Schlichtkrull, Andreas Vlachos, 2022, *MIT Press Direct*. In *Transactions of the Association for Computational Linguistics vol(10): 178–206*. Available online: [https://doi.org/10.1162/tacl\\_a\\_00454](https://doi.org/10.1162/tacl_a_00454). The task of an AFC has three stages: (i) ascertain which claims are 'check-worthy'; (ii) retrieve evidence relevant to determining the truth-value of these claims; (iii) return flags or recommend deletion for those claims which do not meet the relevant standards for veracity. There is a long way to go before any AFCs operate with the kind of scope and effectiveness that we are imagining in this paper. But AFCs are not a pipe dream, especially given recent advances in natural language processing.

In this paper, we argue that the careful deployment of AFCs is likely to be a profound social good. Our case in favor of AFCs is grounded in attention to the distinctive reasons why misinformation is harmful. So, we begin in §2 by asking this question: misinformation seems bad, but why? Contrary to popular belief, we argue that misinformation is not objectionable merely in virtue of its risking *physical* harm. Instead, we argue that the central social concerns raised by misinformation are *epistemic* and *political*. We argue that these social concerns provide us with distinctive motivations for developing AFCs. In §3 and §4, we consider two objections which respectively suggest that AFCs risk *reinforcing* rather than alleviating the relevant epistemic and political concerns identified in §2. However, we respond to both of these objections optimistically. In particular, we offer realistic proposals for how developers may avoid these pitfalls, and suggest that current developments in the technology of AFCs may already be heading in this direction.

## **§2. Motivating fact-checking**

The motivation for developing and deploying AFCs presumably arises from the basic intuition that misinformation is *bad*. But in what does the badness of misinformation consist? This question is important: understanding *why* misinformation is bad will reveal why we should aspire towards effective AFCs.

One tempting thought is that what makes misinformation bad is merely its potential to cause physical harm. For example, the claim that seat belts are harmful to children is bad because exposure to it led David to put his children at risk of injury. Misinformation about silver as a treatment for Coronavirus was bad because it risked people turning blue. The thought is that our incentive to reduce misinformation derives from its capacity to cause physical harm.

This tempting thought appears to be the received view. As one example of the thought in practice, representatives of *Meta* cited this idea in their justification for extending Facebook's fact-checking

capabilities soon after the outbreak of Covid-19, claiming that Facebook's response was simply another application of their more general commitment to reducing harm:

“We will also start to remove content with false claims or conspiracy theories that have been flagged by leading global health organizations and local health authorities that *could cause harm* to people who believe them. We are doing this as *an extension of our existing policies to remove content that could cause physical harm.*”<sup>7</sup>

What is perhaps so tempting about this idea is that it surely gets something right: misinformation *can* (and does) lead to physical harms which would have been prevented or reduced if individuals had never believed it.

However, this cannot be the whole story. After all, agents' beliefs in *true* claims can equally lead to physical harm which would otherwise be prevented had those agents never believed them. Take the January 6th attacks on the capitol caused by misinformation about the 2020 presidential election. The injuries which individuals suffered there would have been equally as *physically* harmful even if it later emerged that Biden had in fact rigged the vote. Risk of physical harm is not a feature unique to misinformation. To justify unique interventions for misinformation, it must be that the kinds of harms it causes are *distinctively* bad, in a way that harm caused by exposure to the truth is not.

What kinds of distinctive harms would justify intervention in the case of misinformation, yet not justify intervention in the case of harmful truths? We'll suggest two answers to this question.

First, consider David from *Viral Tweet*. It's natural to think that what makes David's actions *distinctively* bad is that his *reasons* for acting are in some sense defective. Not only does he put his

---

<sup>7</sup> 'Keeping People Safe and Informed About the Coronavirus', Kang-Xing Jin (Head of Health at Meta), Meta, January 30th 2020: <https://about.fb.com/news/2020/12/coronavirus/>. Emphasis added.

children at risk of physical harm; what's especially objectionable is that he does so on the basis of a falsehood.

Let's make this more precise. Our claim is that misinformation threatens agents' abilities to do good practical reasoning – in other words, it impairs agents' capacities to make effective decisions about how to act. David has a set of beliefs and a set of goals, and, like any other rational agent, he relies on his beliefs about the world to make plans for how best to realize his goals. In *Viral Tweet*, David's goal is to protect his children (an admirable goal by anyone's lights). However, his attempts to realize his goals are frustrated by the fact that the beliefs on which he acts are false. Misinformation has disrupted David's ability to shape his life in a way that accords with his values. We conceive of this as a kind of 'epistemic' harm; a harm which threatens our ability to reason well:

**Epistemic harm:** Misinformation is bad because it causes agents to act on the basis of defective reasons.

Our second concern with misinformation is political. The spread of misinformation, if unchecked, may facilitate and exacerbate the exploitation of vulnerable individuals:

**Political harm:** Misinformation is bad because it makes it easier for vulnerable people to be exploited by those with power.

Misinformation may exacerbate political harms in a variety of complicated ways, but here we'll point to just one. Typically, the level of exposure which a given claim receives depends on the perceived authority of the source. Unless misinformation is moderated, there is a risk that agents could exploit their own perceived authority for nefarious ends. For example, politicians may freely circulate misinformation which supports their own agenda, to the detriment of vulnerable groups.<sup>8</sup>

---

<sup>8</sup> There is a more subtle concern in this vicinity concerning the extent to which misinformation reinforces 'epistemic injustice' of the kind which Miranda Fricker identifies in *Epistemic Injustice: Power and the Ethics*

Here's one clear example of this concern in practice. Notoriously, Trump spread misinformation about the 2020 presidential election online.<sup>9</sup> Those who trusted Trump were persuaded by this misinformation to donate over \$250 million to 'Stop the Steal' legal funds.<sup>10</sup> It seems to us that the donors were exploited – there was no stolen election, but donors were fed falsehoods by a trusted source which made them believe there was an injustice to be righted.<sup>11,12</sup> Regardless of whether the misinformation leads to *physical* harm in cases like these, something bad has happened here.<sup>13</sup> Individuals in positions of authority should not be capable of manipulating the flow of information for exploitative purposes. A robust and efficient method of fact-checking is one way to temper this.

We have argued so far that misinformation causes unique epistemic and political harms. AFCs, by effectively reducing the spread of misinformation, have the potential to mitigate these harms. It's for this reason that we think AFCs are an exciting technology which researchers should continue to explore and develop.

In what follows, we consider two objections which suggest that AFCs may *exacerbate* rather than mitigate the respective epistemic and political concerns identified here. Without solutions to these challenges, AFCs risk making things worse rather than better. We argue, on the contrary, that AFCs are well equipped to avoid these pitfalls.

---

*of Knowing*, Miranda Fricker, 2007, Oxford University Press. Space constraints prohibit a full exploration of this idea here.

<sup>9</sup> "How Trump's bogus Election Day claims broke through Facebook and Twitter Bans", Mark Scott, November 9 2022, *Politico*:

<https://www.politico.com/news/2022/11/09/trump-bogus-election-claims-facebook-twitter-bans-00065977>

<sup>10</sup>"Trump's 'Stop the Steal' Funds were a 'Rip-Off', Jan. 6th Committee Says", Greg Farrell and Bill Allison, December 23 2023:

<https://www.bloomberg.com/news/articles/2022-12-23/jan-6-report-says-donald-trump-s-stop-the-steal-funds-diverted-in-rip-off?leadSource=verify%20wall>

<sup>11</sup>"Donations under \$8K to Trump 'election defense' instead go to president, RNC", Jarrett Renshaw and Joseph Tanfani, *Reuters*:

<https://www.reuters.com/article/us-usa-election-trump-fundraising-insigh/donations-under-8k-to-trump-election-defense-instead-go-to-president-rnc-idUSKBN27R309>

<sup>12</sup> "Trump's Abuse of Recount Funds Shows the Need for Clear Regulations", Shanna Ports, January 31 2023, *CLC*: <https://campaignlegal.org/update/trumps-abuse-recount-funds-shows-need-clear-regulations>

<sup>13</sup> The intuition that this was exploitative becomes even more pressing when one learns that much of the money went directly to Trump and the RNC, and not to lawsuits or recounts. See source in footnote 11.



### §3. An epistemic challenge for AFCs

The first motivation for developing AFCs was to mitigate the epistemic harm associated with misinformation:

**Epistemic harm:** Misinformation is bad because it may cause agents to act on the basis of defective reasons.

The first challenge for AFC development is that AFCs risk reinforcing, rather than alleviating, epistemic harms. The following objection echoes an argument which John Stuart Mill leveled in defense of free speech.<sup>14</sup>

Consider a world in which AFCs are maximally effective – AFCs can detect misinformation instantaneously and with high levels of accuracy. Scrolling through one’s social media feed, the average user will expect to have minimal exposure to misinformation, and what misinformation they do see will have been flagged. In this environment, it would be rational for users to trust all the unflagged information which they are exposed to. After all (they should think) given that the information has not been flagged by an AFC, it is highly unlikely to be false.

This scenario may initially sound appealing, but the concern here is that reliance on AFCs in this way would foster an environment of epistemic complacency. Suppose, for example, that an agent, Naomi, comes to form a belief in this way. She sees a tweet claiming that the Pfizer vaccine is safe and effective against Covid-19, and she believes that this claim is true on the basis that it would have been flagged if false. Naomi has a true belief. However, she lacks a critical understanding of *why* her belief is true. She simply believes it because that’s what she read, and (almost) everything she reads is true.

---

<sup>14</sup> John Stuart Mill, *On Liberty*, first published 1859, Cambridge University Press

There are various reasons why this kind of epistemic complacency is undesirable. One reason which Mill emphasizes is that only by understanding the reasons in favor of one's belief can one effectively defend that belief in the face of objections.<sup>15</sup> Suppose Naomi is en route to receive a Pfizer vaccine, but is met by a protestor who tells her: "the vaccine contains toxic chemicals which will eventually kill you". If Naomi's belief that the vaccine is safe rests only on the fact that a tweet said so, then Naomi will lack the argumentative resources to respond to the protestor. Under this kind of pressure, the flimsy grounds for her initial belief may buckle, leading Naomi to lose a belief which was in fact true. Some exposure to misinformation is good, according to this objection, because it forces agents to engage in critical thinking. True beliefs acquired through critical thinking will be sturdier in the face of counter-argument.

If we heed the Millian warning, we should worry about fact-checking software being self-undermining; AFCs may reinforce rather than mitigate the original epistemic motivation for implementing them. By *eradicating* misinformation, AFCs would risk replacing one kind of defective reason for another. Rather than allowing agents to act occasionally on the basis of falsehoods, too much trust in AFCs may lead agents to act more generally on the basis of beliefs for which those agents have only flimsy grounds.

The objection is a pressing one. However, rather than speaking against developing AFCs at all, we take the Millian argument to be instructive; it teaches us about what we ought to demand from AFC technology. Importantly, developers should not aspire to 'maximally effective' AFCs which render critical reasoning obsolete. Since the motivation for AFCs is in part to reduce epistemic harm, developers should work towards AFCs which *on balance* promote rather than undermine critical thinking.

We have two concrete suggestions for how AFC developers may secure this commitment. First, an AFC's domain of application should be somewhat limited; that is, the aspiration should never be for

---

<sup>15</sup> *On Liberty*, John Stuart Mill: chapter II

AFCs to flag *all* instances of misinformation or even all instances of viral misinformation. Instead, one natural proposal would be to encode a prioritization procedure into AFCs, such that only categories of information which are deemed ‘high-stakes’ are eligible for flagging or removal. The hope would then be that uncertainty about the AFCs’ domain of application would support individuals’ critical thinking, since users would be unable to infer the truth of a claim directly from its lacking a flag. It is worth noting that AFC developers are already conscious that prioritization of this kind is desirable, and are working to encode prioritization technology.<sup>16</sup> There are of course challenges with this strategy – for example, who decides what counts as ‘high stakes’? This issue returns in the next section. Regardless, we take this response to be at least preliminarily compelling.

Our second recommendation is that AFCs should provide transparent justification for each flagged claim. Where a claim has been flagged by an AFC, access to the evidence on which the AFC judged the claim to be false should be readily available for users’ scrutiny. The hope is that including justifications for flags would encourage readers to critically evaluate reasons in favor or against a particular belief. Again, some developers have already recognised the importance of transparent justification procedures, and are developing technologies which translate the complicated language within which AFCs make veracity judgments, into justifications which are digestible to the average reader.<sup>17,18</sup>

#### **§4. A political challenge for AFCs**

Onto another objection. Recall that the second motivation for developing AFCs was to mitigate the political harms associated with misinformation:

---

<sup>16</sup> “A Survey on Automated Fact-Checking”, Zhijiang Guo, Micheal Schlichtkrull, Andreas Vlachos, 2022, *Transactions of the Association for Computational Linguistics vol(10): 178–206*: section 2.1

<sup>17</sup> “A Survey on Automated Fact-Checking”, Zhijiang Guo, Micheal Schlichtkrull, Andreas Vlachos, 2022, *MIT Press Direct*. In *Transactions of the Association for Computational Linguistics vol(10): 178–206*: section 2.4

<sup>18</sup> “Generating fact checking explanations.”, Pepa Atanasova, Jakob Grue Simonsen, Christina Lioma, and Isabelle Augenstein, 2020, In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7352–7364

**Political harm:** Misinformation is bad because it makes it easier for vulnerable people to be exploited by those with power.

The worry which we consider in this section is that deployment of AFCs may exacerbate the ways in which vulnerable individuals are exploited. Most generally, the worry is that the dataset on which an AFC is trained must to some extent be shaped by human choices. These choices may reflect implicit biases and prejudices of the human developers, in which case the AFC may serve to reinforce rather than alleviate social injustice.

One version of this worry concerns choices about prioritization. Recall, in the last section we suggested that AFCs ought to prioritize flagging ‘high stakes’ claims over those judged ‘lower stakes’. Call claims which are higher stakes in this way ‘check-worthy’. The trouble is this: it’s not immediately clear how an AFC should determine check-worthiness. The risk is that by encoding the AFC with any particular model of check-worthiness, its judgments will thereafter be contaminated with implicit human biases and prejudices. For example, if check-worthiness is defined as a matter of ‘what the general public would be interested in knowing’<sup>19</sup>, an AFC may judge that misinformation about A-list celebrities is more check-worthy than misinformation which risks inflaming hatred towards certain racial groups. At this point, the AFC would be failing to carry through on its promise to mitigate political harms.

One natural response to this challenge is to demand transparency from AFCs and AFC developers. Making their datasets transparent would provide a safety net whereby users, activists and other interested parties could decide for themselves whether the veracity judgments made by the AFC

---

<sup>19</sup> Naeemul Hassan, Chengkai Li, and Mark Tremayne. 2015. Detecting check-worthy factual claims in presidential debates. In *Proceedings of the 24th ACM International Conference on Information and Knowledge Management, CIKM 2015*, pages 1835–1838.

reproduce human bias. Transparency of datasets may also incentivise developers to be more rigorous in weeding out the implicit biases encoded into their own technology.<sup>20</sup>

This response alone is clearly not perfect. Implicit bias is often difficult, if not impossible, to identify even in humans. Once the complicated technical machinery of AI enters the picture, this difficulty only gets compounded. It's perhaps too much (at least at this time) to expect that individuals would be capable of identifying implicit bias in AI any better than they can in humans (for a more detailed argument in this direction, see Zerilli et al (2019)<sup>21</sup>).

Instead, at this time, we think that the best way to address this challenge may be with investment in and attention to further research on implicit bias. Organizations like *The Algorithmic Justice League* founded by MIT researcher Joy Buolamwini, exist to raise awareness of 'algorithmic injustice' and to facilitate research on effective solutions to challenges like these.<sup>22</sup> Alongside this important work, fostering relationships between academics who specialize in social issues like implicit bias, and those who work on AFCs, is a crucial stepping stone. We believe that as this research is developed and applied, the risks of AFCs will be outweighed by the profound social benefits of limiting misinformation which we identified in §2.

## **§5. Conclusion**

We argued that misinformation causes two distinctive social harms: (i) it makes it harder for agents to do good practical reasoning – to make decisions which reflect their values and goals; and (ii) it facilitates the exploitation of the vulnerable by those who have power. We argued that Automated Fact

---

<sup>20</sup> This proposal is clearly connected with the previous section. Explicit, accessible justifications for an AFCs veracity judgments, for example, may help not only with the epistemic concern, but with the political concern too.

<sup>21</sup> Zerilli, J., Knott, A., Maclaurin, J., and Gavaghan, C. (2019) 'Transparency in Algorithmic and Human Decision Making: Is There a Double Standard?', *Philosophy and Technology*, Vol 32, pp 661–683

<sup>22</sup> The Algorithmic Justice League: <https://www.ajl.org/about>

Checkers are well-placed to temper the spread of misinformation, thereby mitigating both of these social ills.

The prospects for AFCs are exciting. Provided they are deployed responsibly, we expect these technologies to prove essential to the healthy functioning of future societies.

## References

### Books

*Epistemic Injustice: Power and the Ethics of Knowing*, Miranda Fricker, 2007, Oxford University Press

*On Liberty*, John Stuart Mill, first published 1859, Cambridge University Press

### Journal Articles

Zerilli, J., Knott, A., Maclaurin, J., and Gavaghan, C. (2019) ‘Transparency in Algorithmic and Human Decision Making: Is There a Double Standard?’, *Philosophy and Technology*, Vol 32, pp 661–683

“The Spread of True and False News Online”, Soroush Vosoughi, Deb Roy, and Sinan Aral, *Science*, Vol 359, pp 1146-1151

“A Survey on Automated Fact-Checking”, Zhijiang Guo, Micheal Schlichtkrull, Andreas Vlachos, 2022, *MIT Press Direct*. In *Transactions of the Association for Computational Linguistics vol(10)*: 178–206

“Generating fact checking explanations.”, Pepa Atanasova, Jakob Grue Simonsen, Christina Lioma, and Isabelle Augenstein, 2020, In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7352–7364

“Detecting check-worthy factual claims in presidential debates”. Hassan Naeemul , Chengkai Li, and Mark Tremayne. 2015. In *Proceedings of the 24th ACM International Conference on Information and Knowledge Management, CIKM 2015, Melbourne, VIC, Australia, October 19–23, 2015*, pages 1835–1838. ACM.

### News Articles

“Alex Jones is Told to Stop Selling Sham Anti-Coronavirus Toothpaste”, Luis Ferre-Sadurni and Jesse McKinley, The New York Times, March 13th 2020:  
<https://www.nytimes.com/2020/03/13/nyregion/alex-jones-coronavirus-cure.html>

“Facebook doubles down on removing coronavirus conspiracy theories”, Shirin Ghaffary and Rebecca Heilweil, Vox, March 4th 2020:  
<https://www.vox.com/recode/2020/1/31/21115589/coronavirus-wuhan-china-myths-hoaxes-facebook-social-media-tiktok-twitter-wechat>

“People are bad at spotting fake news. Can computer programs do better?”, Maria Temming, Science News, July 26th 2018: <https://www.sciencenews.org/article/can-computer-programs-flag-fake-news>

“How Trump’s bogus Election Day claims broke through Facebook and Twitter Bans”, Mark Scott, 9th November 2022, *Politico*:  
<https://www.politico.com/news/2022/11/09/trump-bogus-election-claims-facebook-twitter-bans-00065977>

“Trump’s ‘Stop the Steal’ Funds were a ‘Rip-Off’, Jan. 6th Committee Says”, Greg Farrell and Bill Allison, December 23 2023:  
<https://www.bloomberg.com/news/articles/2022-12-23/jan-6-report-says-donald-trump-s-stop-the-steal-funds-diverted-in-rip-off?leadSource=verify%20wall>

Donations under \$8K to Trump ‘election defense’ instead go to president, RNC”, Jarrett Renshaw and Joseph Tanfani, *Reuters*:  
<https://www.reuters.com/article/us-usa-election-trump-fundraising-insigh/donations-under-8k-to-trump-election-defense-instead-go-to-president-rnc-idUSKBN27R309>

“Trump’s Abuse of Recount Funds Shows the Need for Clear Regulations”, Shanna Ports, January 31 2023, *CLC*:  
<https://campaignlegal.org/update/trumps-abuse-recount-funds-shows-need-clear-regulations>

### **Other Websites**

‘Keeping People Safe and Informed About the Coronavirus’, Kang-Xing Jin (Head of Health at Meta), Meta, January 30th 2020: <https://about.fb.com/news/2020/12/coronavirus/>

The Algorithmic Justice League: <https://www.ajl.org/about>

NIH National Center for Complementary and Integrative Health: ‘Colloidal Silver’:  
<https://www.nccih.nih.gov/health/colloidal-silver>