OXFORD JOURNAL
OF STUDENT
SCHOLARSHIP

The Age of Misinformation: How Exposure, Belief, and Sharing Shape Susceptibility Across Age Groups

Oscar Li
oscarsmily@gmail.com

ABSTRACT

This paper explores three distinct yet interconnected interactions with misinformation on digital platforms in order to analyze the factors that may stand behind age-based correlation. It draws from research from 2016 to 2025, using a structured literature review to argue that susceptibility cannot be reduced to just cognitive decline, which is what has been assumed in the past. By organizing the literature around these three forms of interaction, this paper argues that older adults are often better than their younger counterparts at detecting misinformation and that susceptibility to misinformation stems primarily from repeated exposure, motivated reasoning, and the engagement-optimizing algorithms used by social media platforms. Three case studies (the 2016 election, the COVID-19 pandemic, and the rise of generative AI) show the growing impact misinformation has in our digitized world and potential harms in the near future. The paper concludes that misinformation persists not solely due to individuals, but also in part because it may serve the interests of political actors and technological platforms.

INTRODUCTION

"Pope Francis Endorses Donald Trump for President." This fabricated headline from WTOE 5 News, a self-proclaimed fantasy website, was shared almost a million times on Facebook before election day in 2016 (Silverman). In fact, this false story generated more engagement on Facebook than top election stories from major news outlets including the New York Times, Washington Post, and NBC News. This event in 2016 can now be seen as an ominous foreshadowing of how social media would completely shift the nature of misinformation, with significant impacts on global politics.

Fake news has become a pervasive feature of online discourse, whether it fabricates headlines to attract clicks or portrays events in a biased manner to fit a certain argument. During the pandemic, COVID and vaccine misinformation generated 3.8 billion views through a 12 month period (Denniss and Lindberg). Yet the extent to which misinformation shapes political discourse, voter behavior, and public trust remains a neglected and understudied aspect of social and political life.

While most people are familiar with the expression "fake news," defining it concretely poses an immediate issue. Many scholars disagree over the true definition of "fake news", and there is debate around whether it is a discrete category of content or a spectrum of reliability. Studying various existing definitions provides insight on the common threads that run through the scholarly literature.

On the one hand, Allcott and Gentzkow define fake news as "articles that are intentionally and verifiably false, and could mislead readers" (213). This definition emphasizes deliberate intent to deceive and demonstrable inaccuracy. Baptista and Gradim expand on this definition by including motive, describing fake news as "misleading and/or false statements… intentionally created to manipulate…typically in pursuit of ideological or financial gain" (2). Other scholars approach fake news through typology rather than definition. Vasu identifies five overlapping forms: state-sponsored disinformation distributed to undermine national security; politically motivated misinformation spread by domestic groups; non-political misinformation that becomes viral; satirical or parodic falsehoods; and falsehoods distributed for financial gain (5).

Across these definitions, fake news is generally characterized by falsity, intent, and news-like form. This is especially true with reference to intent. To clarify that dimension, this paper uses the information disorder framework proposed by Wardle and Derakhshan, who define misinformation as false information shared without intent to harm and disinformation as false information deliberately created to deceive (5). Fake news, despite its colloquial use, will appear infrequently due to the contested status of its definition. While Wardle and Derakhshan distinguish these terms by intent, this paper uses misinformation as a broader umbrella term encompassing both, as is consistent with common usage in the field. Disinformation will be specified when intent is relevant.

This paper narrows a focus on fake news to a specific question about the intersection of age and use: How does age impact exposure to, belief in, and sharing of online misinformation, and what cognitive, social, and structural factors drive those effects? Drawing on a body of research spanning 2016 to 2025, it argues that while age is one of the strongest demographic predictors of engagement with misinformation, susceptibility to believing it cannot be reduced only to cognitive decline. Instead, susceptibility is driven by motivated reasoning, the engagement-driven design of social media platforms, and social factors that shape how people encounter and share information regardless of their ability to detect falsehood. In order to elucidate this argument, the present paper proceeds in five stages. After offering a brief historical overview of the place of social media within the broader history of misinformation, the paper reviews empirical literature around the mechanisms of exposure, belief, and sharing. Thereafter, the fourth section of the paper looks into the overarching impacts of fake news and the fifth delves into who actually benefits from the increase in misinformation. After reviewing these findings, the paper suggests that misinformation persists not because of individual failures of judgement, but in part because it may serve the interests of political actors and technology companies whose incentives reward engagement over truth.

HISTORICAL BACKGROUND

The use of mass media to misinform and deceive the public has existed for centuries. However, the conditions under which misinformation spreads have shifted drastically in recent years, as editorial gatekeeping has been eroded through the rise of the internet and attention-driven business models. These

shifts matter for understanding age as each generation formed their own expectations about news and trust within a different media environment. What distinguishes modern day media from previous moments is not the essence of misinformation, but the digital environment in which it can now spread.

For much of the twentieth century, news was concentrated in a small number of institutional outlets, each with its own set of editors and operating under professional norms focusing on accuracy. This did not eliminate error entirely, but it did implement institutional checkpoints between a claim and its audience. The repeal of the Fairness Doctrine in 1987 removed the requirement that broadcasters must present controversial issues with viewpoints on either side, allowing openly partisan content to flourish. This gatekept environment shaped how today's older generations came to understand news as they learned to trust that content carrying the form and authority of an established outlet had passed through editorial verification, a habit that would prove costly when applied to online news.

By the early 2000s, the rise of social media represented a more fundamental shift in how information could be disseminated and accessed. As Allcott and Gentzkow observe, social media platforms are structurally different from older technologies, allowing users to publish claims without third-party filtering, fact-checking, or editorial judgment (211). It also removes all barriers to mass distribution, allowing an individual user with no journalistic background to reach audiences comparable to those of major news outlets. Crucially, this is the environment in which younger generations came of age. Unlike older adults, who formed their habits when news passed through editorial gatekeepers, younger users grew up native to a system with no such filters, learning from the start to treat online content with skepticism and to verify claims themselves.

In addition to the collapse of gatekeepers, digital media operated under a different set of incentives than traditional media. Traditional media like newspapers depended on subscriptions and advertising, and readers were unlikely to pay for sources they didn't trust. On the other hand, social media generates revenue through sensationalism, and their algorithms are designed to drive clicks, shares, and comments. Research by Vosoughi, Roy, and Aral found that false news on Twitter spreads "farther, faster, deeper, and more broadly than the truth in all categories of information", suggesting that there is tension between the incentives of social media platforms and the goals of accurate journalism (1146) Taken together, these shifts provided an environment fundamentally different from the one in which older generations received their news. As such, older adults often extend that same trust they had in vetted institutions to online content that often has no supporting sources, setting the stage for the age-based patterns examined in the sections that follow.

METHODOLOGY

The following analysis draws on existing empirical research rather than producing original data. Sources were gathered primarily through Google Scholar and JSTOR, though citation chaining was used to locate additional sources. Searches were conducted between May and June of 2026 using the following terms:

"misinformation and age," "exposure to misinformation," "misinformation belief," and "sharing of misinformation" with results limited to journal articles and book chapters. Only peer-reviewed work was calculated into the analysis of this paper, including journal articles and a Pew Research Center report. These studies were selected by their relevance to the intersection of age and online misinformation. Additionally, the publication window of these sources spanned from 2016-2025, beginning with the surge of scholarly attention prior to the 2016 election. Studies were excluded if they did not address age, were purely theoretical, or fell outside this window. A small number of non-peer-reviewed sources, including news reports and institutional publications, were used to document recent events that had not been covered in the academic literature. Each study was then classified according to the primary interaction it examined: exposure, belief, sharing, or real-world impact.

The specific contribution of the paper addresses a gap in the existing research. Most peer-reviewed publications I found only focus on one form of interaction with misinformation in isolation. Yet they paint a fragmented picture. While it is well established that age correlates with engagement with misinformation, the extent to which it affects specific interactions is not compared. Furthermore, there are inconsistent definitions of "older adults" across studies. Thus, this paper contributes through the comparison and synthesis it offers, building a coherent framework built around exposure, belief, sharing, and impact. In doing so, it demonstrates that age operates differently at each stage of interaction and traces the field's growing emphasis on motivation and platform design alongside cognitive explanations.

THE UNIQUE MECHANISMS BEHIND EXPOSURE

The core of this essay traces a process of interaction and engagement with social media that moves from exposure, to belief, and finally sharing. Exposure is the first of the three interactions examined within this paper. Exposure differs from belief and sharing in one crucial way: it does not require any judgment on the part of the reader. Encountering misinformation is the initial step before someone can believe or share that false story, and exposure is shaped by different factors than the other two mechanisms. However, it is still important to understand how exposure is decided, whether through algorithms or content creators. The studies below span from 2016 to 2020, examining how different age groups are exposed to misinformation.

Gottfried and Shearer provide the baseline on which many later exposure studies build upon (2016). They utilize a large representative sample of 4,654 U.S. adults and map how these participants encounter news across 9 different social networking sites. They found that 62% of Americans got news on social media, which was an increase from 49% just 4 years prior. Out of these sites, Facebook was the largest platform, reaching 44% of U.S. adults. However, the main core of this study comes from the age breakdown. 58% of adults who get their news from Instagram are between 18 and 29 in contrast to just 2% being 65 and older. The study also made a distinction between passive and active consumption, finding that users often encountered news incidentally rather than actively seeking it out. This suggests that exposure is shaped in

part by how platforms distribute content, rather than what users deliberately choose. It is important to note the limitations of this study, however; exposure to any online content on these social media sites are not evenly distributed by age due to the fact that different age groups use specific platforms more often than others, complicating the claim that a specific age group sees more misinformation in absolute terms.

Three years later, Grinberg et al. directly measured misinformation exposure on Twitter during the 2016 election season, one in which misinformation circulated extensively (2019). They linked 16,442 Twitter accounts to public voter registration records to tie the results to real votes. They also classified 300 websites as fake news sources based on existing fact-checker lists and a manual annotation process, then tracked which of these sites appeared in each panel member's feed. The study's central finding was that just 1% of panel members accounted for 80% of all exposure to fake news sources, and these "superconsumers" had, on average, nearly 4,700 daily exposures to political URLs compared with only 49 for the typical panel member. Setting aside this outlier group, the strongest predictors of how much misinformation an individual's feed contained were age and overall political engagement, with age positively associated with increased exposure across every point on the political spectrum. Still, the study found that even the heaviest misinformation consumers received the bulk of their political content from mainstream outlets, with fake sources making up on average just 1.18% of political exposures. Grinberg and colleagues' contribution is showing that age-driven exposure is real but concentrated in a small population of highly engaged users, rather than reflecting the typical older user's feed, and that misinformation exposure exists as a persistent but minor thread within a media diet still dominated by mainstream sources.

Loos and Nijenhuis tested exposure directly by building their own misinformation website and spreading 14 fake articles as Facebook advertisements between February and June 2018, spending 700 euros to promote them across various age brackets (2020). They found that older age groups were far more likely to see these advertisements, with the 55-64 age group receiving 24.7% while 13-17 just received 4.4%. This furthers the idea that age skew in misinformation exposure is not just a byproduct of who uses the platform, but may also reflect the algorithmic delivery systems.

In conclusion, these three studies complicate a simple story about age and exposure. If platform demographics were the whole picture, younger users should see more misinformation simply by virtue of spending more time on social media, as Gottfried and Shearer's numbers suggest (2016). Instead, the opposite holds specifically for misinformation, and Grinberg and colleagues locate that reversal in a narrow slice of highly active older users rather than the platform's older population as a whole (2019). Loos and Nijenhuis push the explanation one step further, suggesting that the pattern is not merely a byproduct of who uses the platform: when identical fake advertisements were distributed across multiple age groups, the platform's own delivery algorithm disproportionately showed them to older users. Though the study does not provide specific reasoning, it indicates that the algorithmic distribution of content, rather than user choice alone, contributes to age-based differences in exposure (2020). The exposure literature therefore suggests a mechanism distinct from the belief or sharing dynamics explored in the sections that follow, one where age-based patterns emerge before a person ever exercises judgment.

Exposure - Comparison of Studies

Study	Sample & Method	Age Definition	Key Finding
Gottfried & Shearer (2016)	4,654 U.S. adults; Pew survey across 9 platforms	18–29 vs. 65+	Younger users cluster on social platforms (58% of Instagram news users are 18–29 vs. 2% for 65+); news is often encountered incidentally rather than sought out
Grinberg et al. (2019)	16,442 Twitter accounts linked to voter records	Age as a continuous predictor	1% of users ("superconsumers") accounted for 80% of fake-news exposure; age positively predicts exposure, but concentrated in a small, highly engaged group
Loos & Nijenhuis (2020)	Field study; 14 fake articles run as Facebook ads	Brackets from 13–17 to 65+	Older users saw far more fake ads (55–64: 24.7% vs. 13–17: 4.4%), suggesting algorithmic delivery skews exposure by age

BELIEF AND SOCIAL MEDIA

Belief refers to the internal acceptance of a piece of misinformation as true. Where exposure is concerned about whether or not people encounter false content, belief looks at whether or not they accept it as accurate. To understand susceptibility to misinformation, one must understand why people accept false information as true. Early explanations often blamed specific deficiencies, whether in analytical ability,

memory, or digital literacy. Yet more recent works have complicated this view, introducing a range of factors that all shape belief regardless of cognitive ability. By tracing academic literature on this topic chronologically, this review reveals how the field has increasingly emphasized motivated reasoning alongside earlier cognitive explanations, rather than replacing one with the other. The studies below span from 2019 to 2025, exploring the field from a variety of perspectives, including neurophysiological measurement, cognitive synthesis, surveys, and behavioral analysis.

Research on belief in misinformation has consistently pointed to political identity as a powerful distorting force. Moravec and colleagues provided neurophysiological evidence for this, asking 83 undergraduate students, aged 18 to 34, to evaluate the believability of news headlines while their brain activity was recorded through electroencephalography (2019). Some headlines aligned with participants' political beliefs and others contradicted them, and some were flagged as false by a fact-checking indicator. The study produced three notable findings. First, confirmation bias strongly shaped credibility judgments: participants were significantly more likely to rate a headline as credible when it aligned with their existing political beliefs, regardless of its actual accuracy. Second, the "false" flags had no measurable effect on whether participants believed headlines that aligned with their views, suggesting that corrective labels do little to overcome motivated reasoning. Third, the EEG data revealed heightened activity in the frontal cortices when participants encountered belief-consistent headlines that had been flagged as false, a pattern the authors interpret as cognitive dissonance, indicating that participants expended additional mental effort to reconcile the conflict between their beliefs and the warning. The study's primary contribution to the question of belief is its demonstration that political identity, rather than evidence, often governs what people accept as true, and that interventions like warning labels may fail precisely because they collide with motivated reasoning. Its limitations include a small, young, undergraduate sample, which constrains generalization across the age spectrum.

Where Moravec and colleagues focused on a young sample, Brashier and Schacter turned specifically to the question of why older adults appear especially susceptible to believing misinformation (2020). Rather than conducting a single experiment, the authors synthesized decades of cognitive psychology research to evaluate three competing explanations: cognitive decline, social changes, and digital illiteracy. They argue, counterintuitively, that cognitive decline alone cannot account for older adults' engagement with misinformation. While source memory does deteriorate with age, other relevant capacities remain intact or even improve. Processing fluency is preserved, and accumulated general knowledge often allows older adults to reject false claims that contradict what they know. The authors note that older adults are actually more accurate than younger adults at distinguishing true from false headlines at first glance. Susceptibility, they argue, emerges instead through repeated exposure: when a false claim is encountered multiple times, source memory failures combine with the illusory truth effect, whereby repetition increases the felt sense that a claim is true. The authors also emphasize social changes in late adulthood such as increased interpersonal trust and a greater emphasis on relational rather than informational goals, and gaps in digital literacy, since older adults are relative newcomers to the conventions of social media. The study's central contribution is its argument that belief in misinformation among older adults is not a simple product of cognitive failure but the result of multiple interacting mechanisms, a framing that

reshaped how the field approaches age and misinformation.

The same year, Rampersad and Althiyabi examined whether demographic factors such as age generalize beyond Western contexts, investigating how demographics and culture shape the acceptance of misinformation in Saudi Arabia (2020). They distributed a quantitative online survey via WhatsApp to Saudi students abroad. This survey was taken by 107 respondents, and analyzed using structural equation modeling to test the effects of age, gender, education, and culture on misinformation acceptance. Among demographic factors, age had the biggest influence on the acceptance of misinformation, supporting the claim that older individuals were more accepting of misinformation. Education on the other hand, showed a negative effect, implying that more educated respondents were less likely to accept false content. However, when considering culture, it exerted the strongest overall influence on acceptance as it shaped how comprehensible a news item appeared, which in turn drove acceptance. This study is important because it showed how the effects of age were not just a Western phenomenon. Its limitations are its small, non-representative sample of students abroad, limiting how far the findings can generalize to the broader Saudi Arabian population.

In the following year, Lyons and colleagues shifted attention towards a specific psychological trait rather than demographics: overconfidence (2021). Utilizing data from two YouGov surveys totaling over 8000 participants representative of the US population, the researchers tested whether those who are overconfident are more vulnerable to misinformation. They had respondents rate the accuracy of 12 headlines framed as Facebook posts and separately estimate how their ability compares to the average American. Overconfidence was represented as the gap between actual and perceived ability. The findings were striking. Across both surveys, ~73% of respondents rated their ability above their actual ability by an average of 22 percentiles. Those who performed worst at detecting misinformation were also the least aware of their deficiency. The bottom quartile of performers rated themselves on average at the 63rd percentile. Furthermore, overconfidence led to a greater willingness to share misinformation. Overall, this study looks at a metacognitive failure as the driver of susceptibility, rather than demographic factors studied previously. The large representative samples used strengthens the generalizability of its conclusions.

More recently, Lutz and colleagues studied how the mechanism of belief works, examining how linguistic features of an article can affect cognitive (careful, conscious thinking) and affective (automatic, unconscious evaluation) thinking, and how that leads to acceptance of misinformation (2024). Their methodology was a study where 42 college students read 40 articles, half false, and half true. The researchers tracked cognitive processing through eye fixations and affective processing through heart rate variability, where lower variability signals discomfort, and computed thirty-three linguistic cues per article. Participants correctly judged veracity only 73.5% of the time, underscoring the difficulty of the task even for educated readers. On the cognitive side, people thought more carefully about longer fake articles but less about simply written, adverb-heavy ones, implying that plainly written misinformation tends to slip past scrutiny. The more novel finding concerned affective processing: significant effects appeared only for misinformation, and articles written with more analytic-sounding language produced

greater discomfort. The authors interpret this to mean that analytic language may make a fake story appear more authoritative, triggering an unconscious response that could lead readers to accept it without realizing it. Notably, an exploratory analysis found that older participants engaged more in cognitive processing while younger participants were more prone to the affective response, hinting that age shapes not just whether people believe misinformation but how they process it. Its limitations include a small undergraduate sample and an artificial reading environment.

The most recent research has begun to challenge the cognitive framing that dominated the previous decade. In a study reported by *CU Boulder Today*, Ramos and Van Boven surveyed roughly 2,500 adults aged 18 to 80 across the United States and Brazil, presenting them with true and false political headlines and asking both whether they would share each one and whether they believed it (Marshall, 2025). The study found no evidence that older adults are less analytically capable or less able to distinguish true from false news. Instead, older adults were significantly more partisan, and that partisanship shaped how critically they evaluated content: they applied stricter standards of evidence to claims that reflected poorly on their own political side and more lenient standards to claims that flattered it. Because this pattern held across both the United States, a two-party system, and Brazil, which has many competing parties, the authors argue that the effect reflects age-related partisanship rather than the structure of any particular political system. This study represents a significant reframing of the question of belief: where earlier research located older adults' susceptibility in cognitive decline, source-memory failure, or digital illiteracy, Ramos and Van Boven argue it is in motivated reasoning. The finding directly complicates interventions built on the assumption that older adults simply need better skills.

Together, these works trace the evolution in how researchers have explained belief in misinformation. The earliest work claimed that individual characteristics were to blame. Subsequent research identified more specific mechanisms like metacognitive overconfidence and the unconscious response to linguistic features. Yet one thing is clear across these studies: purely cognitive explanations are insufficient on their own. Brashier and Schacter found that older adults are actually more proficient at first-glance detection (2020). This does not mean that cognitive factors are unimportant, but rather that they interact with social ones. This emphasis on motivated reasoning frames a central tension in the paper: if belief in misinformation is driven partly by what people want to believe, then common interventions aimed at building skills may not address the root of the misinformation crisis.

Belief - Comparison of Studies

Study	Sample & Method	Age Definition	Key Finding
Moravec et al. (2019)	83 undergraduates; EEG experiment	18–34	Confirmation bias governs belief; "false" flags have no effect on belief-consistent headlines
Brashier & Schacter (2020)	Review of cognitive-psychology research	"Older adults" (65+ focus)	Cognitive decline alone cannot explain susceptibility; older adults detect better at first glance; repeated exposure and the illusory-truth effect drive belief
Rampersad & Althiyabi (2020)	107 Saudi students; survey + structural equation modeling	Age as a demographic factor	Age was the strongest demographic predictor of acceptance, though culture had the strongest overall effect
Lyons et al. (2021)	8,000+ U.S. adults; two YouGov surveys	Representative adult sample	Overconfidence, a metacognitive failure, predicts susceptibility and greater willingness to share

Lutz et al. (2024)	42 college students; lab experiment (eye-tracking, HRV)	Mean age 26	Analytic-sounding language may make misinformation feel more authoritative; older participants processed more cognitively, younger more affectively
Ramos & Van Boven (2025)	~2,500 adults, U.S. & Brazil; survey	18–80	No age deficit in detection; partisanship, not cognition, best explains older adults' heightened susceptibility

THE RELATIONSHIP BETWEEN BELIEF AND SHARING ON SOCIAL MEDIA

Sharing refers to the act of passing misinformation along to others, whether through posting, reposting, or sending privately. This section therefore focuses less on the act of sharing itself than on its underlying motivations, examining why people choose to spread misinformation. Researchers have identified a range of drivers, including social identity, emotional resonance, and a sense of a story's relevance to one's community, that operate independently of whether the content is accurate. The studies below, from 2018 to 2020, trace these motivations across different platforms and age groups.

In 2018, Alice Marwick, a media scholar, delved into the reasoning behind the sharing of misinformation (2018). She wanted to understand if there was a deeper reason, investigating whether or not political factors were involved and the impact of “deep stories.” A “deep story,” drawn from sociologist Arlie Hochschild, is an emotional narrative about how the world works that feels true to a particular group of people. Marwick built a model from two samples: the top 100 BuzzFeed-identified misinformation stories on Facebook and the Snopes “Hot 50.” Each story was coded as political or not political, and she tested two hypotheses: H1 was misinformation mostly right-leaning and H2 was right-leaning misinformation tracks conservative “deep stories.” From her analysis, the majority of problematic content was non-political, proving H1 false. However, it is important to note that among the political stories, most were right-leaning, with some notably discussing conservative talking points like Obama’s birth certificate and sensational claims around Hillary Clinton’s criminality. H2 on the other hand, was

supported as right-wing misinformation continued to map onto conservative "deep stories." The relevance of Marwick's study is the conclusion that people spread misinformation due to it resonating with their group's deep story. This conclusion shapes the act of sharing as one of belonging, rather than a claim about factual accuracy. However, Marwick's study analyzes specific stories rather than individuals. It does not distinguish between age demographics, a question that later studies would take up.

A year later, Guess, Nagler, and Tucker investigated how prevalent misinformation sharing was on Facebook and what characteristics predicted it (2019). They used a representative survey of around 3,500 Americans linked directly to their Facebook data. Misinformation was defined through a curated list of false domains from BuzzFeed and sharing from all these participants were tracked. They found that sharing was extremely rare overall, with over 90% of participants sharing no misinformation at all. Yet a striking age gap emerged: users over 65 shared nearly seven times as many misinformation articles as the youngest group (18-29), even after controlling for education, ideology, and partisanship. Earlier, a study by Brashier and Schacter was featured, which concluded that older adults are better than their younger counterparts at first-time recognition, highlighting how sharing is inherently different from belief.

At the other end of the age spectrum, Herrero-Diz examined how and why adolescents share misinformation, with a specific focus on WhatsApp (2020). The researchers surveyed 480 teenagers (14-18) drawn through random sampling from secondary schools in Seville, Spain, during the 2018-2019 school year. They had each participant rate several cases of real content and misinformation on the likelihood that they would share the content. The study found that teens were equally willing to share true items and fabricated ones as they largely could not distinguish content by its truthfulness. However, clickbait was shared less often than other content. When asked why they shared these articles, the participants said informational reasons the most and popularity the least. These findings suggest that sharing reflects whether or not the content seems relevant to their network rather through verification. This study extends the conclusions to the youngest users of these platforms, showing that the gap between sharing and belief is not unique to older adults, but rather applies all across the age spectrum

Together, these studies reveal that sharing misinformation is affected by a range of social and identity-based factors that persist throughout the age spectrum. Marwick demonstrated the impact of a group's "deep story," one that feels true regardless of the facts (2018). Guess and colleagues found that older adults share far more misinformation than younger users, despite being better at detecting misinformation than their younger counterparts (2019). Herrero-Diz showed that adolescents share news driven by whether or not content is useful to them (2020). Overall, these findings all indicate a gap between what people share and what they believe, that across ages, the decision to share content does not come from verification. This distinction matters for resolving this issue as it suggests that reducing the spread of misinformation must also target the social motivations driving it.

Sharing - Comparison of Studies

Study	Sample & Method	Sharing Behavior	Age Definition	Key Finding
Marwick (2018)	Content analysis of top fake-news stories (BuzzFeed, Snopes)	Broad circulation of stories	N/A (studies stories, not people)	People spread misinformation that resonates with their group's "deep story"; sharing is an act of belonging, not verification
Guess et al. (2019)	~3,500 U.S. adults linked to Facebook data	Public posting to Facebook feed	18–29 vs. 65+	Sharing is rare (90%+ share none), but users over 65 share ~7x more than the youngest group
Herrero-Diz et al. (2020)	480 teens; survey	Private forwarding on WhatsApp	14–18	Teens share true and false content equally, driven by perceived relevance rather than verification

THE IMPACT UPON THE MOST VULNERABLE

Together, the previous sections examined how misinformation is transmitted and whether people believe it. This section seeks to answer the question of whether or not misinformation has a real-world impact. It will look into three cases chronologically: the 2016 election, COVID-19, and AI-generated content. The section will also seek to discover whether or not age is a factor in how people respond to misinformation.

During the 2016 election, misinformation reached an enormous scale on social media, penetrating political discourse and raising serious concerns over democratic discourse online. Bovet and Makse studied the prevalence of misinformation prior to the election with a dataset of 171 million tweets in the 5 months leading up to the election (2019). 30 million of those tweets linked to a news outlet, and 25% of those pointed at sources spreading fake or extremely biased news, giving insight into how deeply misinformation had penetrated political conversation on the platform. This study also revealed how most misinformation were spread by unverified or later-deleted accounts, suggesting that distribution is often an anonymous effort.

Allcott and Gentzkow provide insight on whether or not misinformation was persuasive during the election (2017). They agree that it was pervasive since the average adult in their study saw and remembered at least one misinformation story in the months around the election. Yet they explicitly stated that even if each instance of misinformation was as effective as a TV ad campaign, the impact on voters would be on a much smaller scale than Trump's margins of victory in swing states (2017, 232). As such, Allcott and Gentzkow conclude that misinformation was widespread but its actual impact on the outcome was likely insignificant.

While the impact of misinformation on the 2016 election is hard to quantify, the COVID-19 pandemic offers a case where the consequences are far more observable. Health misinformation during the pandemic produced real harm, because it affected people's actions.

Rocha and colleagues conducted a systematic review of 14 studies across 14 countries, examining how misinformation about COVID-19 impacted people's lives (2021). They found that in 36% of the studies, exposure to infodemic content generated fear, panic, depression, and anxiety. Furthermore, a popularized myth that drinking alcohol would eliminate the virus killed around 800 people in Iran, and hospitalized over 5000 others. Their results conclude that there are measurable impacts of health misinformation, leading to psychological and physical harm.

While Rocha and colleagues documented the human toll of misinformation, Moscadelli focuses on the scale (2020). They utilized the analytics tool BuzzSumo, an analytics tool that measures how widely content is shared across social media platforms to review 2102 of the most-shared pandemic-related articles circulating in Italy between December 2019 and April 2020, classifying them as either fake or

science-based. They found that misinformation links were shared over 2.3 million times, or 23.1% of all shares. Furthermore, misinformation was far more likely to be shared than verified news when standardized for the same number of articles. For example, on the topic of garlic as a cure, 18% of fake articles accounted for nearly 70% of the shares. Taken together, these two sources show that COVID-19 misinformation was just as pervasive as during the 2016 election, but that it was directly harmful as well.

Finally, the rise of AI in recent years has led to a change in how misinformation is transmitted and the basic state of misinformation. With generative AI increasing in capacity year after year, it's become increasingly difficult to distinguish between AI-generated content and real-world contexts. Loth, Kappes, and Pahl study how generative AI affects misinformation in its creation and detection (2024). They found, through a structured literature review of 59 articles, that AI can automate the creation of high-quality, individually targeted misinformation for almost no cost. This changes the environment instantly from the previous two scenarios. During the 2016 election and COVID, misinformation was human-made, and limited by labor. However, AI removes that ceiling. The authors point to deepfakes, synthetic content impersonating real people, as an emerging frontier and the location of an "arms race" between the creation of this content and the technology to detect it (2024).

The rise of AI has already led to real impacts with human detection and belief. Bashardoust and colleagues did an experimental study with 988 human subjects, who each were faced with 20 different misinformation items that were either human generated or AI-generated in random order (2024). For each, users were asked to assess the accuracy of each item and the willingness to share the item. They found that 98.9% of users failed to correctly identify at least one AI-generated news. However, they also found that users were more easily fooled by human-generated misinformation over AI-generated counterparts. This implies that AI-generated misinformation, while having a measurable impact, has not yet surpassed human-written misinformation in its ability to deceive.

Across all three cases, the reach of misinformation is undeniable, while its precise effect is harder to establish. During the 2016 election, misinformation accounted for a quarter of news-linking tweets, yet there is no definite answer on whether this led to a change in the outcome of the election. On the other hand, the wide pervasive nature of misinformation across the world led to hundreds of preventable deaths and widespread anxiety. Now with the rise in generative-AI, the world must be prepared for an influx of cheap yet high quality false content into social media. These cases demonstrate that the consequences of misinformation are real and growing, even when the exact magnitude cannot be measured numerically.

WHO BENEFITS

A lingering question that all of this raises about the impact of misinformation is who are the beneficiaries from the drastic increase in misinformation. If the results of the studies surveyed above are overwhelmingly negative, why does social media continue to propagate such misinformation? One recalls that misinformation has always been a part of political gaming, but the modern increase in scale points

towards two likely beneficiaries: politicians whose partisan aims are served through division and tech companies whose profit structures reward engagement over veracity.

In the past decade political actors have repeatedly exploited misinformation that serves their aims by dividing and mobilizing the electorate. This dynamic appears across all three cases this paper examines. During the 2016 election, the users who shared the most misinformation were self-reported as conservative or very conservative and the content ran heavily in one political direction (Guess et al. 2019). Grinberg and colleagues found the same result on X, formerly known as Twitter, where fewer than 5% of users on the left or center ever shared misinformation compared to 11% and 21% on the right and extreme right respectively (2019, 376). This partisan skew reveals that misinformation is not just neutral noise, but a tool that is used to reinforce existing political divisions and energize the electorate. Similarly, politicians used COVID as a cover to serve political interests. In April 2020, President Trump halted US funding to the WHO, accusing the organization of being "China-centric" ("China Wins"). Yet a Trump campaign confidant, speaking anonymously to TIME, remarked that "China has now become weaponized to our advantage," highlighting the potential political logic behind the withdrawal. AI marks a shift in this pattern. While previously, politicians would frame existing content in a way to either discredit their opponents or energize their supporters, generative AI allows political operatives to cheaply create and distribute their own content. This became reality in January 2024, where voters in New Hampshire received a robocall using the voice of President Joe Biden urging them to not vote in the coming election ("Political Consultant"). The operative behind this faced criminal charges and federal fines, but this example showcases how generative AI can lead to election manipulation, lowering the cost of deception.

Platforms also profit heavily from the engagement that arises with misinformation. Even without explicit intent, these profit structures contribute towards its spread. Misinformation, which is often provocative, is especially effective at commanding that attention. Vosoughi, Roy, and Aral found that false news spreads significantly faster on Twitter than true news, which they attribute to the strong emotional responses it provokes in readers (2018, 1146). This is also supported by the findings of exposure studies. Loos and Nijenhuis showed that a platform's algorithm can route fabricated content towards specific demographics (2020). Grinberg and colleagues also discovered that misinformation concentrated among the most engaged accounts (2019). This implies a system fundamentally opposed to accuracy because platforms are rewarded for attention. These platforms have little incentive in curbing misinformation, even if they publicly oppose it.

Taken together, the interests of political actors and technology platforms may help explain why misinformation persists despite its overwhelmingly negative effects. misinformation endures not because it serves the public, but because its existence appears to benefit those who create misinformation online. Former Facebook product manager Frances Haugen, who disclosed thousands of internal documents in 2021, said the company found that if they changed its algorithm to be safer, "people will spend less time on the site, they'll click on less ads, they'll make less money" ("Facebook Whistleblower"). The costs disproportionately fall upon users who are most vulnerable to misinformation. Understanding who

benefits reframes the problem as the product of a system whose incentives reward division and engagement over truth.

CONCLUSION

This paper set out to answer a question: how does age impact exposure to, belief in, and sharing of online misinformation, and what cognitive, social, and structural factors drive those effects? According to the studies analyzed here, age emerges as one of the strongest demographic predictors of engagement with misinformation, whether it was through sharing, encountering, or believing. Yet cognitive decline with age is not the full picture. Older adults have been found to be better than younger adults at recognizing misinformation at first glance, and recent research points towards partisanship as a more convincing reason for their heightened susceptibility. Yet the evidence points towards a combination of factors including repeated exposure, motivated reasoning, overconfidence, and gaps in digital literacy, each operating differently across the age spectrum.

Examining these mechanisms reveals how age matters at every stage. Exposure is shaped by algorithms that route content differently to various age demographics. Political identity drives belief, and sharing follows social motives that operate independently of whether a person believes what they share. The case studies this paper looked into show how the growth in the spread of misinformation has disproportionately affected the stakes in real life.

The conclusions described here, however, should be read in light of several limitations. First, the studies analyzed define "older adults" inconsistently, using age cutoffs that range from the mid-fifties to sixty-five and, in some cases, treating age as a continuous variable, which complicates direct comparison across findings. Second, much of the evidence is correlational rather than causal; these studies establish that age is associated with engagement with misinformation, but they cannot prove that age itself is the cause. Third, several of the key studies rely on small or non-representative samples, such as undergraduate participants, which limits how far their findings generalize. Finally, the literature remains heavily concentrated in Western contexts, and this review drew on a small number of non-peer-reviewed sources to document the most recent events, which are less rigorous than peer-reviewed research. The field of social media analysis is also still young, and more work is needed for a comprehensive understanding of online misinformation.

At the same time, one implication is clear: if susceptibility is driven less by people's actions and more by the systems that deliver information to them, then interventions built solely on teaching digital literacy skills will address only part of the problem. Confronting misinformation will require attention not only to individual judgment, but to the motivations and structures that allow falsehoods to spread. One wonders how social media companies and platforms may take decisive action in advance of the 2028 election, given what has happened in the past two election cycles.

REFERENCES

Allcott, Hunt, and Matthew Gentzkow. "Social Media and Fake News in the 2016 Election." *Journal of Economic Perspectives*, vol. 31, no. 2, 2017, pp. 211–236.

Baptista, João Pedro and Anabela Gradim. "Understanding Fake News Consumption: A Review." *Social Sciences*, vol. 9, no. 10, 2020, article 185.

Bashardoust, Amirsiavosh, et al. "Comparing the Willingness to Share for Human-Generated vs. AI-Generated Fake News." *Proceedings of the ACM on Human-Computer Interaction*, vol. 8, no. CSCW2, 2024, article 489.

Bovet, Alexandre and Hernán A. Makse. "Influence of Fake News in Twitter during the 2016 US Presidential Election." *Nature Communications*, vol. 10, 2019, article 7.

Brashier, Nadia M. and Daniel L. Schacter. "Aging in an Era of Fake News." *Current Directions in Psychological Science*, vol. 29, no. 3, 2020, pp. 316–323.

"China Wins: Why Trump's WHO Funding Cut Is a Gift to Beijing." *TIME*, 15 Apr. 2020, time.com/5822078/china-wins-trumps-who-funding-cut/.

Denniss, Emily and Rebecca Lindberg. "Social Media and the Spread of Misinformation: Infectious and a Threat to Public Health." *Health Promotion International*, vol. 40, no. 2, 2025, daaf023.

"Facebook Whistleblower Frances Haugen Details Company's Misleading Efforts on 60 Minutes." *CBS News*, 5 Oct. 2021, www.cbsnews.com/news/facebook-whistleblower-frances-haugen-misinformation-public-60-minutes-2021-10-03/.

Gottfried, Jeffrey and Elisa Shearer. "News Use Across Social Media Platforms 2016." *Pew Research Center*, 26 May 2016, www.pewresearch.org/journalism/2016/05/26/news-use-across-social-media-platforms-2016.

Grinberg, Nir, et al. "Fake News on Twitter during the 2016 U.S. Presidential Election." *Science*, vol. 363, no. 6425, 2019, pp. 374–378.

Guess, Andrew, et al. "Less than You Think: Prevalence and Predictors of Fake News

Dissemination on Facebook." *Science Advances*, vol. 5, no. 1, 2019, article eaau4586.

Herrero-Diz, Paula, et al. "Teens' Motivations to Spread Fake News on WhatsApp." *Social Media + Society*, vol. 6, no. 3, 2020, pp. 1–14.

Loos, Eugène and Jordy Nijenhuis. "Consuming Fake News: A Matter of Age? The Perception of Political Fake News Stories in Facebook Ads." *Human Aspects of IT for the Aged Population*, edited by Qin Gao and Jia Zhou, Springer, 2020.

Loth, Alexander, et al. "Blessing or Curse? A Survey on the Impact of Generative AI on Fake News." *arXiv*, 2024, arXiv:2404.03021.

Lutz, Bernhard, et al. "Which Linguistic Cues Make People Fall for Fake News? A Comparison of Cognitive and Affective Processing." *Proceedings of the ACM on Human-Computer Interaction*, vol. 8, no. CSCW1, 2024, article 191.

Lyons, Benjamin A., et al. "Overconfidence in News Judgments Is Associated with False News Susceptibility." *Proceedings of the National Academy of Sciences*, vol. 118, no. 23, 2021, pp. 1–10.

Marshall, Lisa. "Older Adults Share More Political Misinformation. Here's Why." CU Boulder Today, University of Colorado Boulder, 5 Nov. 2025, www.colorado.edu/today/2025/11/05/older-adults-share-more-political-misinformation-heres-why.

Marwick, Alice E. "Why Do People Share Fake News? A Sociotechnical Model of Media Effects." *Georgetown Law Technology Review*, vol. 2, no. 2, 2018, pp. 474–512.

Moravec, Patricia L., et al. "Fake News on Social Media: People Believe What They Want to Believe When It Makes No Sense At All." *MIS Quarterly*, vol. 43, no. 4, 2019, pp. 1343–1360.

Moscadelli, Andrea, et al. "Fake News and COVID-19 in Italy: Results of a Quantitative Observational Study." *International Journal of Environmental Research and Public Health*, vol. 17, no. 16, 2020, article 5850.

"A Political Consultant Faces Charges and Fines for Biden Deepfake Robocalls." *NPR*, 23 May 2024, www.npr.org/2024/05/23/nx-s1-4977582/.

Rampersad, Giselle, and Turki Althiyabi. "Fake News: Acceptance by Demographics and

Culture on Social Media." *Journal of Information Technology & Politics*, vol. 17, no. 1, 2020, pp. 1–11.

Rocha, Yasmim Mendes, et al. "The Impact of Fake News on Social Media and Its Influence on Health during the COVID-19 Pandemic: A Systematic Review." *Journal of Public Health*, 2021.

Silverman, Craig. "This Analysis Shows How Viral Fake Election News Stories Outperformed Real News on Facebook." *BuzzFeed News*, 16 Nov. 2016, www.buzzfeednews.com/article/craigsilverman/viral-fake-election-news-outperformed-real-news-on-facebook.

Vasu, Norman, et al. "Unpacking Fake News." *FAKE NEWS: NATIONAL SECURITY IN THE POST-TRUTH ERA*, S. Rajaratnam School of International Studies, 2018, pp. 5–9. *JSTOR*, http://www.jstor.org/stable/resrep17648.5. Accessed 18 July 2026.

Vosoughi, Soroush, et al. "The Spread of True and False News Online." *Science*, vol. 359, no. 6380, 2018, pp. 1146–1151.

Wardle, Claire, and Hossein Derakhshan. Information Disorder: Toward an Interdisciplinary Framework for Research and Policy Making. Council of Europe, 2017.