

February 2019

Google

# How Google Fights Disinformation

# Our work to tackle the intentional spread of misinformation across Google Search, Google News, YouTube, and our advertising platforms.

---

<b>Introduction</b> .....	2
What is disinformation?.....	2
Tackling disinformation across Google.....	3
Teaming up with newsrooms and outside experts.....	5
Getting ahead of future risks.....	7
<b>Google Search, Google News &amp; Disinformation</b> .....	9
Background.....	9
Tackling disinformation on Google Search & Google News.....	10
<b>YouTube &amp; Disinformation</b> .....	17
Background.....	17
Tackling disinformation on YouTube.....	18
<b>Google Advertising Products &amp; Disinformation</b> .....	25
Background.....	25
Tackling disinformation on Google’s advertising platforms.....	26
<b>Conclusion</b> .....	29

---

## Introduction

The open Internet has enabled people to create, connect, and distribute information like never before. It has exposed us to perspectives and experiences that were previously out-of-reach. It has enabled increased access to knowledge for everyone.

Google continues to believe that the Internet is a boon to society – contributing to global education, healthcare, research, and economic development by enabling citizens to become more knowledgeable and involved through access to information at an unprecedented scale.

However, like other communication channels, the open Internet is vulnerable to the organized propagation of false or misleading information. Over the past several years, concerns that we have entered a “post-truth” era have become a controversial subject of political and academic debate.

These concerns directly affect Google and our mission – to organize the world’s information and make it universally accessible and useful. When our services are used to propagate deceptive or misleading information, our mission is undermined.

How companies like Google address these concerns has an impact on society and on the trust users place in our services. We take this responsibility very seriously and believe it begins with providing transparency into our policies, inviting feedback, enabling users, and collaborating with policymakers, civil society, and academics around the world.

This document outlines our perspective on disinformation and misinformation and how we address it throughout Google. It begins with the three strategies that comprise our response across products, and an overview of our efforts beyond the scope of our products. It continues with an in-depth look at how these strategies are applied, and expanded, to Google Search, Google News, YouTube, and our advertising products.

We welcome a dialogue about what works well, what does not, and how we can work with others in academia, civil society, newsrooms, and governments to meet the ever-evolving challenges of disinformation.

## What is disinformation?

As we’ve all experienced over the past few years, the words “misinformation”, “disinformation”, and “fake news” mean different things to different people and can become politically charged when they are used to characterize the propagators of a specific ideology or to undermine political adversaries.

However, there is something objectively problematic and harmful to our users when malicious *actors attempt to deceive them*. It is one thing to be wrong about an issue. It is another to purposefully disseminate information one knows to be inaccurate with the hope that others believe it is true or to create discord in society.

**We refer to these deliberate efforts to deceive and mislead using the speed, scale, and technologies of the open web as “disinformation”.**

The entities that engage in disinformation have a diverse set of goals. Some are financially motivated, engaging in disinformation activities for the purpose of turning a profit. Others are politically motivated, engaging in disinformation to foster specific viewpoints among a population, to exert influence over political processes, or for the sole purpose of polarizing and fracturing societies. Others engage in disinformation for their own entertainment, which often involves bullying, and they are commonly referred to as “trolls”.

Levels of funding and sophistication vary across those entities, ranging from local mom-and-pop operations to well-funded and state-backed campaigns. In addition, propagators of disinformation sometimes end up working together, even unwittingly. For instance, politically motivated actors might emphasize a piece of disinformation that financially motivated groups might latch onto because it is getting enough attention to be a potential revenue source. Sometimes, a successful disinformation narrative is propagated by content creators who are acting in good faith and are unaware of the goals of its originators.

This complexity makes it difficult to gain a full picture of the efforts of actors who engage in disinformation or gauge how effective their efforts may be. Furthermore, because it can be difficult to determine whether a propagator of falsehoods online is acting in good faith, responses to disinformation run the risk of inadvertently harming legitimate expression.

## Tackling disinformation in our products and services

We have an important responsibility to our users and to the societies in which we operate to curb the efforts of those who aim to propagate false information on our platforms. At the same time, we respect our users' fundamental human rights (such as free expression) and we try to be clear and predictable in our efforts, letting users and content creators decide for themselves whether we are operating fairly. Of course, this is a delicate balance, as sharing too much of the granular details of how our algorithms and processes work would make it easier for bad actors to exploit them.

We face complex trade-offs and there is no 'silver bullet' that will resolve the issue of disinformation, because:

- It can be extremely difficult (or even impossible) for humans or technology to determine the veracity of, or intent behind, a given piece of content, especially when it relates to current events.
- Reasonable people can have different perspectives on the right balance between risks of harm to good faith, free expression, and the imperative to tackle disinformation.
- The solutions we build have to apply in ways that are understandable and predictable for users and content creators, and compatible with the kind of automation that is required when operating services on the scale of the web. We cannot create standards that require deep deliberation for every individual decision.
- Disinformation manifests differently on different products and surfaces. Solutions that might be relevant in one context might be irrelevant or counter-productive in others. Our products cannot operate in the exact same way in that regard, and this is why they approach disinformation in their own specific ways.

Our approach to tackling disinformation in our products and services is based around a framework of three strategies: make quality count in our ranking systems, counteract malicious actors, and give users more context. We will outline them in this section, as well as the efforts we undertake beyond the scope of our products and services to team up with newsrooms and outside experts, and to get ahead of future risks. It is worth noting that these strategies are also used to address misinformation more broadly, which pertains to the overall trustworthiness of the information we provide users in our products.

In later sections of this paper, we will detail how these strategies are implemented and expanded for Google Search, Google News, YouTube, and our advertising platforms. We adopt slightly different approaches in how we apply these principles to different products given how each service presents its own unique challenges.

## 1. Make Quality Count

Our products are designed to sort through immense amounts of material and deliver content that best meets our users' needs. This means delivering quality information and trustworthy commercial messages, especially in contexts that are prone to rumors and the propagation of false information (such as breaking news events).

While each product and service implements this differently, they share important principles that ensure our algorithms treat websites and content creators fairly and evenly:

- Information is organized by "ranking algorithms".
- These algorithms are geared toward ensuring the usefulness of our services, as measured by user testing, not fostering the ideological viewpoints of the individuals that build or audit them. When it comes to Google Search, you can find a detailed explanation of how those algorithms operate at [google.com/search/howsearchworks](https://google.com/search/howsearchworks).

## 2. Counteract Malicious Actors

Algorithms cannot determine whether a piece of content on current events is true or false, nor can they assess the intent of its creator just by reading what's on a page. However, there are clear cases of intent to manipulate or deceive users. For instance, a news website that alleges it contains "*Reporting from Bordeaux, France*" but whose account activity indicates that it is operated out of New Jersey in the U.S. is likely not being transparent with users about its operations or what they can trust it to know firsthand.

That's why our policies across Google Search, Google News, YouTube, and our advertising products clearly outline behaviors that are prohibited – such as misrepresentation of one's ownership or primary purpose on Google News and our advertising products, or impersonation of other channels or individuals on YouTube.

Furthermore, since the early days of Google and YouTube, many content creators have tried to deceive our ranking systems to get more visibility – a set of practices we view as a form of 'spam' and that we've invested significant resources to address.

This is relevant to tackling disinformation since many of those who engage in the creation or propagation of content for the purpose to deceive often deploy similar tactics in an effort to achieve more visibility. Over the course of the past two decades, we have invested in systems that can reduce 'spammy' behaviors at scale, and we complement those with human reviews.

### 3. Give Users More Context

Easy access to context and a diverse set of perspectives are key to providing users with the information they need to form their own views. Our products and services expose users to numerous links or videos in response to their searches, which maximizes the chances that users are exposed to diverse perspectives or viewpoints before deciding what to explore in depth.

Google Search, Google News, YouTube, and our advertising products have all developed additional mechanisms to provide more context and agency to users. Those include:

- **“Knowledge” or “Information” Panels** in Google Search and YouTube, providing high-level facts about a person or issue.
- Making it easier to discover the work of **fact-checkers** on Google Search or Google News, by using labels or snippets making it clear to users that a specific piece of content is a fact-checking article.
- A **“Full Coverage”** function in Google News enabling users to access a non-personalized, in-depth view of a news cycle at the tap of a finger.
- **“Breaking News”** and **“Top News”** shelves, and **“Developing News”** information panels on YouTube, making sure that users are exposed to news content from authoritative sources when looking for information about ongoing news events.
- Information panels providing **“Topical Context”** and **“Publisher Context”** on YouTube, providing users with contextual information from trusted sources to help them be more informed consumers of content on the platform. These panels provide authoritative information on well-established historical and scientific topics that have often been subject to misinformation online and on the sources of news content, respectively.
- **“Why this ad”** labels enabling users to understand why they’re presented with a specific ad and how to change their preferences so as to alter the personalization of the ads they are shown, or to opt out of personalized ads altogether.
- **In-ad disclosures and transparency reports on election advertising**, which are rolling out during elections in the US, Europe, and India as a starting point.

We also empower users to let us know when we’re getting it wrong by using **feedback buttons** across Search, YouTube, and our advertising products to flag content that might be violating our policies.

## Teaming up with newsrooms and outside experts

Our work to address disinformation is not limited to the scope of our products and services. Indeed, other organizations play a fundamental role in addressing this societal challenge, such as newsrooms, fact-checkers, civil society organizations, or researchers. While we all address different aspects of this issue, it is only by coming together that we can succeed. That is why we dedicate significant resources to supporting quality journalism, and to weaving together partnerships with many other organizations in this space.

## Supporting quality journalism

People come to Google looking for information they can trust and that information often comes from the reporting of journalists and news organizations around the world.

A thriving news ecosystem matters deeply to Google and directly impacts our efforts to combat disinformation. When quality journalism struggles to reach wide audiences, malicious actors have more room to propagate false information.

Over the years, we've worked closely with the news industry to address these challenges and launched products and programs to help improve the business model of online journalism. These include the Accelerated Mobile Pages Project<sup>1</sup> to improve the mobile web, YouTube Player for Publishers<sup>2</sup> to simplify video distribution and reduce costs, and many more.

In March 2018, we launched the Google News Initiative (GNI)<sup>3</sup> to help journalism thrive in the digital age. With a \$300 million commitment over 3 years, the initiative aims to elevate and strengthen quality journalism, evolve business models to drive sustainable growth, and empower news organizations through technological innovation. \$25M of this broader investment was earmarked as innovation grants for YouTube to support news organizations in building sustainable video operations.

One of the programs supported by the Google News Initiative is Subscribe with Google<sup>4</sup>, a way for people to easily subscribe to various news outlets, helping publishers engage readers across Google and the web. Another is News Consumer Insights, a new dashboard built on top of Google Analytics, which will help news organizations of all sizes understand and segment their audiences with a subscriptions strategy in mind. More details on these projects and others can be found at [g.co/newsinitiative](http://g.co/newsinitiative).

## Partnering with outside experts

Addressing disinformation is not something we can do on our own. The Google News Initiative also houses our products, partnerships, and programs dedicated to supporting **news organizations** in their efforts create quality reporting that displaces disinformation. This includes:

- Helping to launch the **First Draft Coalition** (<https://firstdraftnews.org/>), a nonprofit that convenes news organizations and technology companies to tackle the challenges around combating disinformation online – especially in the run-up to elections.
- Participating in and providing financial support to the **Trust Project** (<http://thetrustproject.org/>), of which Google is a founding member and which explores how journalism can signal its trustworthiness online. The Trust Project has developed eight indicators of trust that publishers can use to better convey why their content should be seen as credible, with promising results for the publishers who have trialed them.
- Partnering with Poynter's **International Fact-Checking Network** (IFCN)<sup>5</sup>, a nonpartisan organization gathering fact-checking organizations from the United States, Germany, Brazil, Argentina, South Africa, India, and more.

In addition, we support the work of **researchers** who explore the issues of disinformation and trust in journalism by funding research at organizations like First Draft, the Oxford University's Reuters News Institute, Michigan University's Quello Center for Telecommunication Management law, and more.

Finally, in March 2018, Google.org (Google's philanthropic arm) launched a \$10 million global initiative to support **media literacy around the world** in the footsteps of programs we have already supported in the UK, Brazil, Canada, Indonesia, and more.

We will continue to explore more ways to partner with others on these issues, whether by building new products that might benefit the work of journalists and fact-checkers, supporting more independent initiatives that help curb disinformation, or developing self-regulatory practices to demonstrate our responsibility.

## Getting ahead of future risks

Creators of disinformation will never stop trying to find new ways to deceive users. It is our responsibility to make sure we stay ahead of the game. Many of the product strategies and external partnerships mentioned earlier help us reach that goal. In addition, we dedicate specific focus to bolstering our defenses in the run-up to elections and invest in research and development efforts to stay ahead of new technologies or tactics that could be used by malicious actors, such as synthetic media (also known as 'deep fakes').

### Protecting elections

Fair elections are critical to the health of democracy and we take our work to protect elections very seriously. Our products can help make sure users have access to accurate information about elections. For example, we often partner with election commissions, or other official sources, to make sure key information like the location of polling booths or the dates of the votes are easily available to users.

We also work to protect elections from attacks and interference, including focusing on combating political influence operations, improving account and website security, and increasing transparency.

To prevent political influence operations, working with our partners at Jigsaw, we have multiple internal teams that identify malicious actors wherever they originate, disables their accounts, and shares threat information with other companies and law enforcement officials. We routinely provide public updates about these operations.<sup>7</sup>

There is more we can do beyond protecting our own platforms. Over the past several years, we have taken steps to help protect accounts, campaigns, candidates, and officials against digital attacks. Our Protect Your Election project<sup>8</sup> offers a suite of extra security to protect against malicious or insecure apps and guards against phishing. To protect election and campaign websites, we also offer Project Shield<sup>9</sup>, which can mitigate the risk of Distributed Denial of Service (DDoS) attacks.

In the run-up to elections, we provide free training to ensure that campaign professionals and political parties are up-to-speed on the means to protect themselves from attack. For instance, in 2018, we trained more than 1,000 campaign professionals and the eight major U.S. Republican and Democratic committees on email and campaign website security.

Furthermore, as a part of our security efforts, for the past eight years, we have displayed warnings to Gmail users who are at risk of phishing by potentially state-sponsored actors (even though, in most cases, the specific phishing attempt never reaches the user's inbox).



Finally, in order to help understand the context for the election-related ads they see online, we require additional verification for advertisers who wish to purchase political ads in the United States, provide transparency about the advertiser to the user, and have established an online transparency report and creative repository on US federal elections.<sup>10</sup>

We look forward to expanding these tools, trainings, and strategies to more elections in 2019, starting with efforts focused on two of the world's largest upcoming elections, which are in Europe<sup>11</sup> and in India.<sup>12</sup>

## Expecting the unexpected

Creators of disinformation are constantly exploring new ways to bypass the defenses set by online services in an effort to spread their messages to a wider audience.

To stay ahead of the curve, we continuously invest resources to stay abreast of the next tools, tactics, or technologies that creators of disinformation may attempt to use. We convene with experts all around the world to understand what concerns them. We also invest in research, product, and policy developments to anticipate threat vectors that we might not be equipped to tackle at this point.

One example is the rise of new forms of AI-generated, photo-realistic, synthetic audio or video content known as "synthetic media" (often referred to as "deep fakes"). While this technology has useful applications (for instance, by opening new possibilities to those affected by speech or reading impairments, or new creative grounds for artists and movie studios around the world), it raises concerns when used in disinformation campaigns and for other malicious purposes.

The field of synthetic media is fast-moving and it is hard to predict what might happen in the near future. To help prepare for this issue, Google and YouTube are investing in research to understand how AI might help detect such synthetic content as it emerges, working with leading experts in this field from around the world.

Finally, because no detector can be perfect, we are engaging with civil society, academia, newsrooms, and governments to share our best understanding of this challenge and work together on what other steps societies can take to improve their preparedness. This includes exploring ways to help others come up with their own detection tools. One example may involve releasing datasets of synthesized content that others can use to train AI-based detectors.<sup>13</sup>

# Google Search, Google News & Disinformation

## Background

Google created its search engine in 1998, with a mission to organize the world's information and make it universally accessible and useful. At the time, the web consisted of just 25 million pages.

Today, we index hundreds of billions of pages – more information than all the libraries in the world could hold – and serve people all over the world. Search is offered in more than 150 languages and over 190 countries.

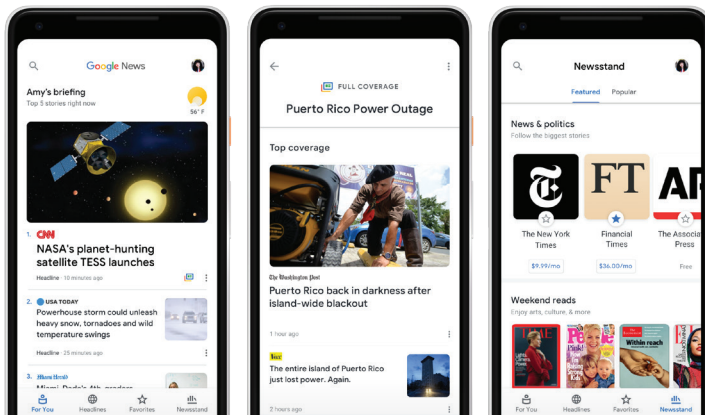
We continue to improve on Search every day. In 2017 alone, Google conducted more than 200,000 experiments that resulted in about 2,400 changes to Search. Each of those changes is tested to make sure it aligns with our publicly available Search Quality Rater Guidelines,<sup>14</sup> which define the goals of our ranking systems and guide the external evaluators who provide ongoing assessments of our algorithms.

Over the past 20 years, we have grappled with the tension between the open access to information and expression that the web enables and the need to ensure trust in authoritative information. Our work on disinformation continues to be informed by these dual goals, as we attempt to strike the right balance in tackling this challenge.

Different types of content may require different approaches to ranking and presentation in order to meet our users' needs. Google News arose from such a realization and was one of the first products Google launched beyond Search. Former Googler Krishna Bharat observed that when people searched for news after the tragic 9/11 attacks in New York, Google responded with old news stories about New York rather than the latest events. He set about to fix that, and on September 22, 2002, Google News was born.

Over time, Google News has improved, including how we present content related to current events in Google Search. In 2018, we launched a reimagined Google News that uses a new set of AI techniques to take

a constant flow of information as it hits the web, analyze it in real-time, and organize it around breaking news events.<sup>15</sup>



Google News App

Through all of this, we've remained grounded in our mission and the importance of providing greater access to information, helping users navigate the open web. We continue to believe that this access is fundamental to helping people make sense of the world around them, exercise their own critical thinking, and make informed decisions as citizens.

## Tackling disinformation on Google Search & Google News

Since Google's early days, malicious actors have attempted to harm or deceive Search users through a wide range of actions, including tricking our systems in order to promote their own content (via a set of practices we refer to as "spam"), propagating malware, and engaging in illegal acts online. The creators and purveyors of disinformation employ many of the same tactics.

Disinformation poses a unique challenge. Google is not in a situation to assess objectively, and at scale, the veracity of a piece of content or the intent of its creators. Further, a considerable percentage of content contains information that cannot be objectively verified as fact. This is because it either lacks necessary context, because it is delivered through an ideological lens others may disagree with, or because it is constructed from contested datapoints.

Disinformation also raises broader concerns of harm. In the worst cases, the impacts of disinformation campaigns can affect an entire society. The stakes of accurately identifying disinformation are higher because disinformation often concerns issues at the core of political society for which the free exchange of ideas and information among genuine voices is of the greatest importance.

To deal with this issue, Google Search and Google News take a pragmatic approach that reinforces the product strategies we have highlighted in the opening section of this paper:

- **Make Quality Count**
  - We use ranking algorithms to elevate authoritative, high-quality information in our products.
  - We take additional steps to improve the quality of our results for contexts and topics that our users expect us to handle with particular care.
- **Counteract Malicious Actors**
  - We look for and take action against attempts to deceive our ranking systems or circumvent our policies.
- **Give Users More Context**
  - We provide users with tools to access the context and diversity of perspectives they need to form their own views.

## Do Google News and Google Search combat disinformation in the same ways?

Google News' focus – coverage of current events – is narrower than that of Google Search. However, their goals are closely related. Both products present users with trustworthy results that meet their information needs about the issues they care about.

For that reason, both products have a lot in common when it comes to the way they operate. For instance, ranking in Google News is built on the basis of Google Search ranking and they share the same defenses against “spam” (attempts at gaming our ranking systems).

In addition, both products share some fundamental principles:

- They use **algorithms, not humans**, to determine the ranking of the content they show to users. No individual at Google ever makes determinations about the position of an individual webpage link on a Google Search or Google News results page.
- Our algorithms are geared toward **ensuring the usefulness of our services, as measured by user testing**, not fostering the ideological viewpoints of the individuals who build or audit them.
- The systems **do not make subjective determinations about the truthfulness of webpages**, but rather focus on **measurable signals** that correlate with how users and other websites value the expertise, trustworthiness, or authoritativeness of a webpage on the topics it covers.

That said, because Google News' purposes are explicitly narrower than those of Google Search and solely focused on coverage of current events, it builds its own *ranking systems* and *content policies* on top of those of Google Search.

When it comes to **ranking**, this means that the systems we use in Google News and in places that are focused on News in Google Search (e.g. our “Top Stories” Carousel or our “News” Tab) make special efforts to understand things like the prominence of a news story in the media landscape of the day, which articles most relate to that story, or which sources are most trusted for specific news topics. It also means that Google News might give additional importance toward factors that indicate a webpage's newsworthiness or journalistic value for users, such as its freshness or (for specific tabs within Google News).

When it comes to **content policies**:

- Google Search aims to make information from the web available to all our users. That's why we do not remove content from results in Google Search, except in very limited circumstances. These include legal removals, violations of our webmaster guidelines, or a request from the webmaster responsible for the page.
- Google Search contains some features that are distinct from its general results, such as Autocomplete. For features where Google specifically promotes or highlights content, we may remove content that violates their specific policies.<sup>16</sup>
- Because Google News does not attempt to be a comprehensive reflection of the web, but instead to focus on journalistic accounts of current events, it has more restrictive content policies than Google Search. Google News explicitly prohibits content that incites, promotes, or glorifies violence, harassment, or dangerous activities. Similarly, Google News does not allow sites or accounts that impersonate any person or organization, that misrepresent or conceal their ownership or primary purpose, or that engage in coordinated activity to mislead users.<sup>17</sup>

With those nuances in mind, it is still safe to think of Google News and Google Search's approaches to disinformation and misinformation are mostly similar, and the content of the following sections apply to both products. Where there is a difference, it will be outlined explicitly in the body of the text or in a dedicated callout box.

## We use ranking algorithms to elevate high-quality information in our products

Ranking algorithms are an important tool in our fight against disinformation. Ranking elevates the relevant information that our algorithms determine is the most authoritative and trustworthy above information that may be less reliable. These assessments may vary for each webpage on a website and are directly related to our users' searches. For instance, a national news outlet's articles might be deemed authoritative in response to searches relating to current events, but less reliable for searches related to gardening.

For most searches that could potentially surface misleading information, there is high-quality information that our ranking algorithms can detect and elevate. When we succeed in surfacing high-quality results, lower quality or outright malicious results (such as disinformation or otherwise deceptive pages) are relegated to less visible positions in Search or News, letting users begin their journey by browsing more reliable sources.

Our ranking system does not identify the intent or factual accuracy of any given piece of content. However, it is specifically designed to identify sites with high indicia of expertise, authority, and trustworthiness.

### How do Google's algorithms assess expertise, authority, and trustworthiness?

- Google's algorithms identify signals about pages that correlate with trustworthiness and authoritativeness. The best known of these signals is PageRank, which uses links on the web to understand authoritativeness.
- We are constantly evolving these algorithms to improve results – not least because the web itself keeps changing. For instance, in 2017 alone, we ran over 200,000 experiments with trained external Search Evaluators and live user tests, resulting in more than 2,400 updates to Google Search algorithms.
- To perform these evaluations, we work with Search Quality Evaluators who help us measure the quality of Search results on an ongoing basis. Evaluators assess whether a website provides users who click on it with the content they were looking for, and they evaluate the quality of results based on the expertise, authoritativeness, and trustworthiness of the content.
- The resulting ratings do not affect the ranking of any individual website, but they do help us benchmark the quality of our results, which in turn allows us to build algorithms that globally recognize results that meet high-quality criteria. To ensure a consistent approach, our evaluators use the Search Quality Rater Guidelines (publicly available online)<sup>18</sup> which provide guidance and examples for appropriate ratings. To ensure the consistency of the rating program, Search Quality evaluators must pass a comprehensive exam and are audited on a regular basis.
- These evaluators also perform evaluations of each improvement to Search we roll out: in side-by-side experiments, we show evaluators two different sets of Search results, one with the proposed change already implemented and one without. We ask them which results they prefer and why. This feedback is central to our launch decisions.

For more information about how our ranking work, please visit:  
[www.google.com/search/howsearchworks](http://www.google.com/search/howsearchworks)

**We take additional steps to improve the trustworthiness of our results for contexts and topics that our users expect us to handle with particular care.**

Our Search Quality Raters Guidelines acknowledge that some types of pages could potentially impact the future happiness, health, financial stability, or safety of users. We call those “Your Money or Your Life” pages or YMYL. We introduced the YMYL category in 2014. They include financial transaction or information pages, medical and legal information pages, as well as news articles, and public and/or official information pages that are important for having an informed citizenry. This last category can comprise anything from information about local, state, or national government processes or policies, news about important topics in a given country, or disaster response services.

General Guidelines		July 20, 2018
<b>General Guidelines Overview</b>		
0.0 Introduction to Search Quality Rating		5
0.1 The Purpose of Search Quality Rating		6
0.2 Raters Must Represent the User		6
0.3 Browser Requirements		6
0.4 Ad Blocking Extensions		6
0.5 Internet Safety Information		6
<b>Part 1: Page Quality Rating Guideline</b>		
1.0 Introduction to Page Quality Rating		7
2.0 Understanding Webpages and Websites		7
2.1 Important Definitions		7
2.2 What is the Purpose of a Webpage?		8
2.3 Your Money or Your Life (YMYL) Pages		9
2.4 Understanding Webpage Content		9
2.4.1 Identifying the Main Content (MC)		9
2.4.2 Identifying the Supplementary Content (SC)		10
2.4.3 Identifying Advertisements/Monetization (Ads)		10
2.4.4 Summary of the Parts of the Page		11
2.5 Understanding the Website		11
2.5.1 Finding the Homepage		11
2.5.2 Finding Who is Responsible for the Website and Who Created the Content on the Page		13
2.5.3 Finding About Us, Contact Information, and Customer Service Information		13
2.6 Reputation of the Website or Creator of the Main Content		14

*Our Search Quality Raters Guidelines*

guidance to our Search Quality Rater Guidelines advising evaluators to give lower quality ratings to informational pages that contain demonstrably inaccurate content or debunked conspiracy theories. While their ratings don't determine individual page rankings, they are used to help us gather data on the quality of our results and identify areas where we need to improve. This data from Search Evaluators also plays a significant role in determining which changes we roll out to our ranking systems.

Beyond specific types of content that are more sensitive to users, we realize that some contexts are more prone to the propagation of disinformation than others. For instance, breaking news events, and the heightened level of interest that they elicit, are magnets for bad behavior by malicious players. Speculation can outrun facts as legitimate news outlets on the ground are still investigating. At the same time, malicious actors are publishing content on forums and social media with the intent to mislead and capture people's attention as they rush to find trusted information. To reduce the visibility of this type of content, we have designed our systems to prefer authority over factors like recency or exact word matches while a crisis is developing.

In addition, we are particularly attentive to the integrity of our systems in the run-up to significant societal moments in the countries where we operate, such as elections.<sup>19</sup>


**We actively look for and take action against attempts to deceive our ranking systems or circumvent our policies.**

Google is designed to help users easily discover and access the webpages that contain the information they are looking for. Our goals are aligned with those of site owners who publish high-quality content online because they want it to be discovered by users who might be interested. That's why we provide extensive tools and tips

For these “YMYL” pages, we assume that users expect us to operate with our strictest standards of trustworthiness and safety. As such, where our algorithms detect that a user's query relates to a “YMYL” topic, we will give more weight in our ranking systems to factors like our understanding of the authoritativeness, expertise, or trustworthiness of the pages we present in response.

Similarly, we direct our Google Search evaluators to be more demanding in their assessment of the quality and trustworthiness of these page than they would otherwise. Specifically, in 2016, we added additional

to help webmasters and developers manage their Search presence and succeed in having their content, sites, and apps found. We provide interactive websites, videos, starter guides, frequent blog posts, user forums, and live expert support to inform webmasters. Our publicly available webmaster guidelines complement these resources by outlining some of the tips and behaviors that we recommend webmasters adopt to make it easiest for our systems to crawl and index their websites.<sup>20</sup>



Google Webmaster Guidelines

Not all site owners act in good faith. Since the early days of Google, many have attempted to manipulate their way to the top of Search results through deceptive or manipulative behavior, using any insights into the functioning of our systems they can get to try to circumvent them. The earliest example of such attempts dates back to 1999, when Google's founders published a seminal paper on PageRank, a key innovation in Google's algorithm.<sup>21</sup> The paper described how our algorithms use links between websites as an indicator of authority. Once that paper was published, spammers tried to game Google by paying each other for links.

These manipulative behaviors aim to elevate websites to users not because they are the best response to a query, but because a webmaster has deceived our systems. As such, they are considered "spam" and run afoul of our core mission. Our webmaster guidelines clearly spell out actions that are prohibited and state that we will take action against websites engaging in such behaviors.

While not all spammers engage in disinformation, many of the malicious actors who try to distribute disinformation (at all levels of sophistication or funding) engage in some form of spam. The tactics they use are similar to those of other spammers. Therefore, our work against spam goes hand-in-hand with our work against disinformation.

Our algorithms can detect the majority of spam and demote or remove it automatically. The remaining spam is tackled manually by our spam removal team, which reviews pages (often based on user feedback) and flags them if they violate the Webmaster Guidelines. In 2017, we took action on 90,000 user reports of search spam and algorithmically detected many more times that number.

As our tactics improve and evolve, so does spam. One of the trends we observed in 2017 was an increase in website hacking, both for spamming search ranking and for spreading malware. We focused on reducing this threat and were able to detect and remove from Search results more than 80 percent of these sites over the following year.

We continue to be vigilant regarding techniques used by spammers and remain conscientious of what we share about the ways our ranking systems work so as not to create vulnerabilities they can exploit.

## Google News policies against deceptive content


In addition to other efforts to fight spam, Google News’ content policies prohibit:

- Sites or accounts that impersonate any person or organization;
- Sites that misrepresent or conceal their ownership or primary purpose;
- Sites or accounts that engage in coordinated activity to mislead users – including, but not limited to, sites or accounts that misrepresent or conceal their country of origin or that direct content at users in another country under false pretenses.


In addition to algorithmic signals that might indicate such behavior, where there is an indication that a publisher may be violating our policies, such as through a user report or suspicious account activity, our Trust and Safety team will investigate and, where appropriate, take action against that site and related sites that can be confirmed to be operating in concert.

## We provide users with the context and diversity of perspectives they need to form their own views.

From the very beginning of Google Search, the very nature of Google Search results pages has ensured that when looking for information on news or public interest topics they care about, users are presented with links to multiple websites and perspectives.



Fact-Check on Google Search



Local Coverage & Timeline in Google News

This remains true today. When users search for news on Google, they are always presented with multiple links. In many cases, they are also presented with additional elements that help them get more context about their search. For instance, “Knowledge Panels” might show in Search results to provide context and basic information about people, places, or things that Google knows about. Fact-check tags or snippets might show below links in Google Search and Google News, outlining that a specific piece of content purports to fact-check a claim made by a third party.<sup>22</sup> Or we might call out related searches or questions that users tend to ask about the topic of a search query.

In Google News, additional cues may help users pick up on points of context that are particularly relevant to News stories, such as “Opinion” or “User-Generated Content” tags under articles that news publishers want to signal as such; or algorithmically generated story timelines that let users explore at-a-glance the milestones of a news story over the weeks or months that led to the day’s events.

## Does Google personalize the content that shows in Google Search and Google News so that users only see news consistent with their views, sometimes known as “filter bubbles”?

We try to make sure that our users continue to have access to a diversity of websites and perspectives. Google Search and Google News take different approaches toward that goal.



**Google Search:** contrary to popular belief, there is very little personalization in Search based on users' inferred interests or Search history before their current session. It doesn't take place often and generally doesn't significantly change Search results from one person to another. Most differences that users see between their Search results and those of another user typing the same Search query are better explained by other factors such as a user's location, the language used in the search, the distribution of Search index updates throughout our data centers, and more.<sup>23</sup> Furthermore, the Top Stories carousel that often shows in Search results in response to news-seeking searches is never personalized.

**Google News:** To meet the needs of users who seek information on topics they care about, Google News aims to strike a balance between providing access to the same content and perspectives as other users and presenting content that relates to news topics one cares about. To do this, Google News offers three interconnected ways to discover information:

- *'Headlines' and 'Top Stories':* To help users stay on top of the top trending news in their country, the "Headlines" and "Top Stories" tabs shows the major stories and issues that news sources are covering at any point in time and shows them to everyone in a non-personalized manner.
- *'For You':* To help users stay on top of the news that matter to them, the "For You" tab lets them specify the topics, publications, and locations they are interested in so they can see the news that relates to those selections. Additionally, and depending on their permission settings, the "For You" tab may show them news they may be interested in in light of their past activity on Google products.
- *'Full Coverage':* To help users access context and diverse perspectives about the news stories they read, the "Full Coverage" feature in Google News lets users explore articles and videos from a variety of publishers related to an article or news story of their choice. The "Full Coverage" feature is not personalized and is accessible in one click or tap from most articles in the "For You" and "Headlines" tabs.

Importantly, for both services, **we never personalize content based on signals relating to point of view on issues and/or political leanings** – our systems do not collect such signals, nor do they have an understanding of political ideologies.

## **We constantly improve our algorithms, policies, and partnerships, and are open about issues we have yet to address.**

Because the malicious actors who propagate disinformation have the incentive to keep doing so, they continue to probe for new ways to game our systems, and it is incumbent on us to stay ahead of this technological arms race. A compounding factor in that challenge is that our systems are constantly confronted with searches they have never seen before. Every day, 15% of the queries that our users type in the Google Search bar are new.

For these reasons, we regularly evolve our ranking algorithms, our content policies, and the partnerships we enter into as part of our efforts to curb disinformation.

We are aware that many issues remain unsolved at this point. For example, a known strategy of propagators of disinformation is to publish a lot of content targeted on "data voids", a term popularized by the U.S. based think tank Data and Society to describe Search queries where little high-quality content exists on the web for Google to display due to the fact that few trustworthy organizations cover them.<sup>24</sup> This often applies, for instance, to niche conspiracy theories, which most serious newsrooms or academic organizations won't make the effort to debunk. As a result, when users enter Search terms that specifically refer to these theories, ranking algorithms can only elevate links to the content that is actually available on the open web – potentially including disinformation.

We are actively exploring ways to address this issue, and others, and welcome the thoughts and feedback of researchers, policymakers, civil society, and journalists around the world.

# YouTube & Disinformation

## Background

YouTube started in 2005 as a video-sharing website and quickly evolved into one of the world's most vibrant online communities. Thousands, then millions, then billions of people connected through content that educated, excited, or inspired them. YouTube is one of the world's largest exporters of online cultural and learning content and is a significant driver of economic activity, providing many of its creators with the ability to make a livelihood by using its services.

Disinformation is not unique to YouTube. It is a global problem afflicting many platforms and publishers. When a platform fosters openness, as we do at YouTube, there's a risk that unreliable information will be presented. While disinformation has been a problem as long as there's been news to report, the Internet has made it possible for disinformation to spread further and faster than ever before. We take our responsibility to combat disinformation in this domain seriously. To be effective at our size, we invest in a combination of technological solutions with a large and growing base of human talent. Technology provides scale and speed, while human talent provides the contextual knowledge needed to fine-tune and improve our systems every step of the way.

YouTube has developed a comprehensive approach to tackling controversial content on our platform. This approach is guided by three principles:

1. Keep content on the platform unless it is in violation of our Community Guidelines
2. Set a high bar for recommendations
3. Monetization is a privilege

From these principles, we create robust systems to responsibly manage all types of controversial content, including disinformation.

Given how broad the spectrum of disinformation is, we implement the three product strategies mentioned in our opening section in ways that are relevant to YouTube's specific products, community, and challenges:

- **Make Quality Count**

- We deploy effective product and ranking systems that demote low-quality disinformation and elevate more authoritative content

- **Counteract Malicious Actors**

- We rigorously develop and enforce our content policies
- We protect the integrity of information tied to elections through effective ranking algorithms, and tough policies against users that misrepresent themselves or who engage in other deceptive practices
- We remove monetary incentives through heightened standards for accounts that seek to utilize any of YouTube's monetization products

- **Give Users Context**

- We provide context to users via information panels on YouTube

We take our responsibilities as a platform seriously. We believe a responsible YouTube will continue to embrace the democratization of access to information while providing a reliable and trustworthy service to our users.

## Tackling disinformation on YouTube

Given the spectrum of content and intent, it is necessary to have a nuanced approach that strikes the right balance between managing our users' expectations to express themselves freely on the platform with the need to preserve the health of the broader community of the creator, user, and advertiser ecosystem.

Let's take a closer look at the three guiding principles on which we base our approach for YouTube:

### 1. Keep content on the platform unless it is in violation of our Community Guidelines

YouTube's Community Guidelines<sup>25</sup> prohibit certain categories of material, including sexually explicit content, spam, hate speech, harassment and incitement to violence. We aim to balance free expression with preventing harmful content in order to maintain a vibrant community. Striking this balance is never easy, especially for a global platform. YouTube has always had Community Guidelines, but we revise them as user behavior changes and as the world evolves.

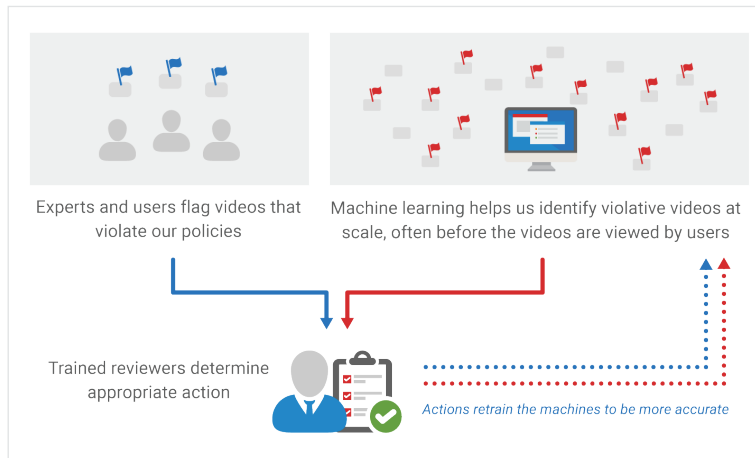
YouTube also maintains a more detailed and living set of enforcement guidelines that provide internal guidance on the enforcement of the public Community Guidelines. These enforcement guidelines are extensive and dynamic to ensure that the policies apply to changing trends and new patterns of controversial content online. YouTube does not typically disclose these updates to the public because doing so would make it easier for unscrupulous users to evade detection.

To help formulate rules that are consistent, unbiased, well-informed, and broad enough to apply to a wide scope of content, YouTube often relies on external subject matter experts and NGOs to consult on various issues. YouTube has also worked with independent experts as a member of the Global Network Initiative (GNI),<sup>26</sup> to establish key principles to guide content review efforts and systems, including notifying users if a video is removed and allowing for appeals. To honor YouTube's commitment to human rights, we also make exceptions to the Community Guidelines for material that is educational, documentary, scientific, and/or artistic.

#### Consistent enforcement

With hundreds of hours of new content uploaded to YouTube every minute, clear policies and enforcement guidelines are only part of what matters. To maintain a site where abuse is minimized, the systems used to curtail abuse must scale. YouTube has always relied on a mix of humans and technology to enforce its guidelines and will continue to do so.

YouTube has thousands of reviewers who operate 24/7 to address content that may violate our policies and the team is constantly expanding to meet evolving enforcement needs. Our review teams are diverse and global. Linguistic and cultural knowledge is needed to interpret the context of a flagged video and decide whether it violates our guidelines. Reviewers go through a comprehensive training program to ensure that they have a full understanding of YouTube's Community Guidelines. We use frequent tests as part of the training process to ensure quality and knowledge retention. Human reviewers are essential to evaluating context and to ensuring that educational, documentary, scientific, and artistic content is protected.



YouTube flagging and review decisions continuously improves our systems

We strive to be as transparent as possible when it comes to actions we take on content on our platform. That is why we release a Community Guidelines Enforcement Report<sup>27</sup> where we give insight into the scale and nature of our extensive policy enforcement efforts. It shows that YouTube's 'crime rate' is low – only a fraction of YouTube's total views are on videos that violate company policies.

**Application to disinformation**

There are several policies in the Community Guidelines that are directly applicable in some form to disinformation. These include policies against spam, deceptive practices, scams,<sup>28</sup> impersonation,<sup>29</sup> hate,<sup>30</sup> and harassment.<sup>31</sup>

The policy against spam, deceptive practices, and scams prohibits posting large amounts of untargeted, unwanted, or repetitive content in videos, comments, private messages, especially if the main purpose of the content is to drive people off to another site. Similarly, activity that seeks to artificially increase the number of views, likes, dislikes, comments, or other metrics either through the use of automated systems or by serving up videos to unsuspecting viewers is against our terms. Additionally, content that exists solely to incentivize viewers for engagement (views, likes, comments, etc.), or by coordinating at scale with other users to drive up views for the primary purpose of interfering with our systems, is prohibited.

One of the abuses this policy covers is content that deliberately seeks to spread disinformation that could suppress voting or otherwise interfere with democratic or civic processes. For example, demonstrably false content that claims one demographic votes on one day while another votes on a separate day would be in violation of our policies.

Another applicable policy regards impersonation. Accounts seeking to spread disinformation by misrepresenting who they are via impersonation are clearly against our policies and the account will be removed. For example, if a user copies a channel's profile, background, or text, and writes comments to make it look like somebody else's channel posted the comments, we remove the channel. Impersonation can also occur if a user creates a channel or video using another individual's real name, image, or other personal information to deceive people into thinking they are someone else on YouTube.

YouTube has clear policies against hate and harassment. Hate speech refers to content that promotes violence against, or has the primary purpose of inciting hatred against, individuals or groups based on certain attributes, such as race or ethnic origin, religion, disability, gender, age, veteran status, or sexual orientation/ gender identity. Harassment may include abusive videos, comments, messages, revealing someone's personal information, unwanted sexualization, or incitement to harass other users or creators. Users that spread disinformation that runs afoul of either our hate or harassment policies will be removed and further appropriate action taken.

## 2. Set a high bar for recommendations

Our primary objective when it comes to our search and discovery systems is to help people find content that they will enjoy watching, whether through their Homepage, Watch Next, or Search results. We aim to provide content that lets users dive into topics they care about, broaden their perspective, and connect them to the current zeitgeist. When a user is seeking content with high intent – subscribing to a channel or searching for the video – it is our responsibility to help users find and watch the video. On the other hand, in the absence of strong or specific intent for a particular video, we believe it is our responsibility to not proactively recommend content that may be deemed low quality.

### **How our approach to recommendations has evolved**

When YouTube's recommendation systems first launched, it sought to optimize for content that got users to click. We noticed that this system incentivized creators to publish misleading and sensationalist clickbait, and users would click on the video but very quickly realize the content was not something they liked. The system was failing to meet our user-centric goals.

To provide a better service for our users, we began to look at the amount of time a video was watched, and whether it was watched to completion, rather than just whether it was clicked. Additionally, we began to demote clickbait. We realized that watchtime was a better signal to determine whether the content we were surfacing to users was connecting them to engaging content they'd enjoy watching. But we learned that just because a user might be watching content longer does not mean that they are having a positive experience. So we introduced surveys to ask users if they were satisfied with particular recommendations. With this direct feedback, we started fine-tuning and improving these systems based on this high-fidelity notion of satisfaction.

The efforts to improve YouTube's recommendation systems did not end there. We set out to prevent our systems from serving up content that could misinform users in a harmful way, particularly in domains that rely on veracity, such as science, medicine, news, or historical events.

To that end, we introduced a higher bar for videos that are promoted through the YouTube homepage or that are surfaced to users through the "watch next" recommendations. Just because content is available on the site, it does not mean that it will display as prominently throughout the recommendation engine.

As has been mentioned previously, our business depends on the trust users place in our services to provide reliable, high-quality information. The primary goal of our recommendation systems today is to create a trusted and positive experience for our users. Ensuring these recommendation systems less frequently provide fringe or low-quality disinformation content is a top priority for the company. The YouTube company-wide goal is framed not just as "Growth", but as "Responsible Growth".

Beyond removal of content that violates our community guidelines, our teams have three explicit tactics to support responsible content consumption. They are:

- **Where possible and relevant, elevate authoritative content from trusted sources.** In areas such as music or entertainment, relevance, newness, or popularity might be better signals to tilt our systems to achieve the user's desired intent and connect them to quality content they'd enjoy. But as we describe in our Search section, in verticals where veracity and credibility are key, including news, politics, medical, and scientific domains, we work hard to ensure our search and recommendation systems provide content from more authoritative sources.
- **Provide users with more context (often text-based information) to make them more informed users on the content they consume.** On certain types of content, including content produced by organizations that receive state or public funding or topical content that tends to be accompanied by disinformation online, we have started to provide information panels that contain additional contextual information and links to authoritative third-party sites so that our users can make educated decisions about the content they watch on our platform.
- **Reduce recommendations of low-quality content.** We aim to design a system that recommends quality content while less frequently recommending content that may be close to the line created by our Community Guidelines, content that could misinform users in harmful ways, or low-quality content that may result in a poor experience for our users, like clickbait. For example, content that claims that the Earth is flat or promises a "miracle cure" for a serious disease might not necessarily violate our Community Guidelines, but we don't want to proactively recommend it to users. Our incentives are to get this right for our users, so we use everyday people as evaluators to provide input on what constitutes disinformation or borderline content under our policies, which in turn informs our ranking systems.

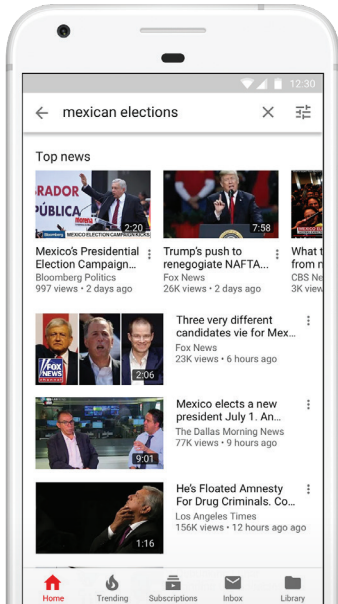
## Case Study: Applying Our Principles to YouTube News & Politics

Disinformation in news and politics is a priority given its importance to society and the outsized impact disinformation can have during fast-moving news events. Although news content only generates a small fraction of overall YouTube watchtime, it's a specific use case that is especially important to us.

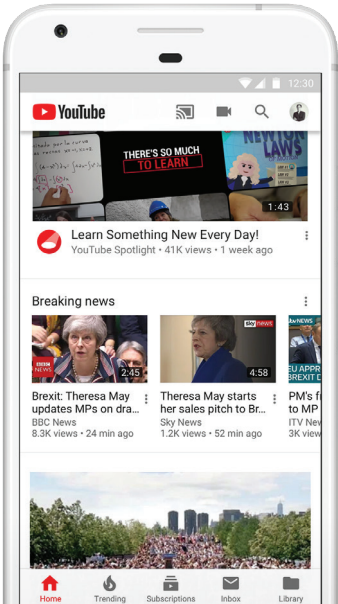
In July 2018, YouTube announced product and partnership tactics that directly apply our guiding content management principles to news.<sup>32</sup>

The first solution included **making authoritative sources readily available**. To help meet this goal, we created a system to elevate authoritative sources for people coming to YouTube for news and politics. For example, if a user is watching content from a trusted news source, the "watch next" panel should similarly display content from other trusted news sources. Assumed within this principle is the demotion of disinformation content that we outlined earlier.

The team has also built and launched two cornerstone products – the Top News shelf and the Breaking News shelf – to prominently display authoritative political news information. The Top News shelf triggers in response to search queries that have political news-seeking intent, and provides content from verified news channels. These systems rely on a variety of signals that are derived from Google News and from our internal systems when a user might be seeking information on a given topic.



Top News Shelf on YouTube Search



Breaking News Shelf on YouTube Home

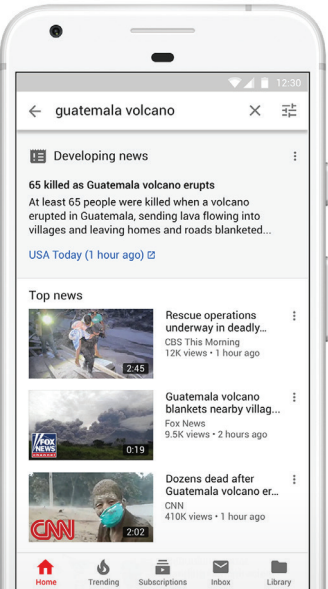
The Breaking News shelf triggers on the YouTube homepage automatically when there is a significant news event happening in a specific country.

Similar to the Top News shelf, only content from authoritative sources is eligible to be displayed in the Breaking News shelf.

More recently, YouTube has been developing products that directly address a core vulnerability involving the spread of disinformation in the immediate aftermath of a breaking news event. After a breaking news

event, it takes some time for journalists to create authoritative fact-based video content and upload it to YouTube, while unscrupulous uploaders can more quickly upload bizarre conspiracy theories.

In these events, YouTube's systems historically delivered the most relevant content that matched the typed query, and without appropriate guardrails would display content from these users seeking to exploit this vulnerability.




Information Panel Providing Breaking News Context

The first step toward a resolution involved creating systems that determine when a breaking news event might be happening, which tilts results tied to that event toward authority and away from strict relevance, popularity, or recency. This helped display content from credible sources.


Furthermore, while authoritative video content takes time, credible text-based reporting is much quicker. As a result, YouTube launched a product that displays an information panel providing text-based breaking news content from an authoritative news source while a significant news event is developing. The information panel also links directly to the article's website so that viewers can easily access and read the full article about the news event.

Once a critical mass of authoritative news videos has been published on the topic, Breaking News and Top News shelves begin to take over as the primary news consumption experiences on the platform.

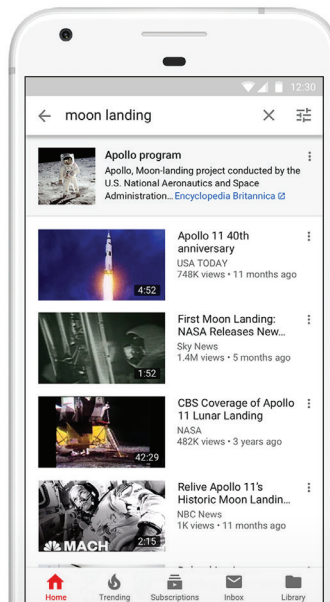


Information Panel Providing Publisher Context

Another information panel product that was launched in 2018 aims to provide additional factual information from outside sources around topics that tend to be accompanied by disinformation online, particularly on YouTube. Users may see panels alongside videos on topics such as “moon landing hoax” that link to information from credible sources, including Encyclopaedia Britannica and Wikipedia.



Information Panel Providing Topical Context on YouTube Search and Watch Pages



An information panel providing topical context may appear in Search results or on the watch page of videos. It will include basic, independent information about a given topic and will link to a third-party website to allow viewers to learn more about the topic. This information panel appears alongside all videos related to the topic, regardless of the opinions or perspectives expressed in the videos. Information panels do not affect any video features or monetization eligibility.

The third solution involves **supporting journalism with technology that allows news to thrive.**

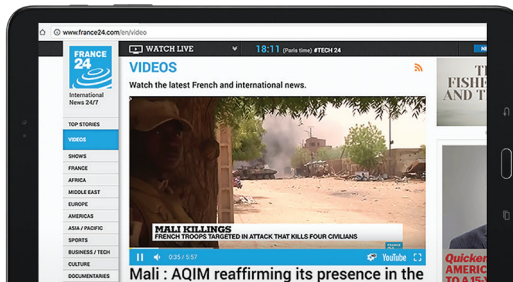
These important societal issues go beyond any single platform and involve shared values across our society. The fight against disinformation is only as good as the quality of news information available in the ecosystem, and are intent on doing our part to support the industry if we are to truly address the broader problem of disinformation. We believe quality journalism requires sustainable revenue streams and that we have a responsibility to support innovation in products and funding for news.

Several years ago, YouTube developed a solution in partnership with key publishers to help newsrooms improve and maximize their video capabilities. The program is called Player for Publishers, and it allows publishers to use YouTube to power the videos on their sites and applications. The program is free



and 100% of the advertising revenue for ads sold by the publisher and served on their own properties goes to the publisher.

In addition, last year, YouTube committed \$25M in funding, as part of a broader \$300M investment by the Google News Initiative,<sup>33</sup> to support news organizations in building sustainable video operations. YouTube announced<sup>34</sup> the winners of our first ever innovation funding program. These partners hail from 23 countries across the Americas, Europe, Africa, and Asia-Pacific, representing a diverse mix of broadcasters, traditional publishers, digital publishers, news agencies, local media, and creators. Best practices gained from this program will be shared publicly via case studies, providing all newsrooms the opportunity to learn and apply insights as we work together to support the development of long-term, sustainable news video businesses.



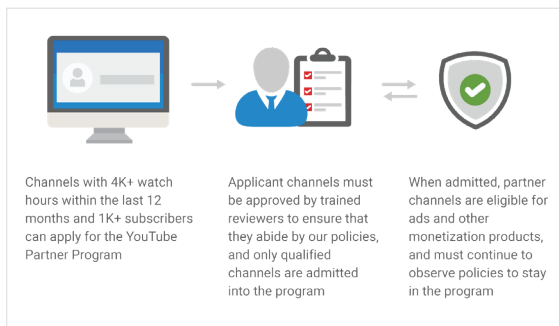
Player for Publishers

In conjunction with this investment, YouTube created a news working group – a quarterly convening of news industry leaders with whom we are collaborating to shape the future of news on YouTube. The news working group consists of top broadcasters, publishers, creators, and academics around the world and they have been providing feedback on items such as how to better quantify authoritativeness, what additional types of information might be useful to provide to users in our information panel projects, and what more we can do to support online video operations in newsrooms. Given how complex these issues are, we know we can't work in a silo. We must partner with industry and civil society to come together around solutions that work.

### 3. We view monetization on our platform as a privilege

Many people use YouTube simply to share their content with the world. Creators who meet the eligibility criteria can apply to join the YouTube Partner Program, which makes their videos eligible to run advertising and earn money through Google's advertising products. Monetizing creators must comply with advertiser-friendly content guidelines. Advertising will be disabled from running on videos that violate these policies.

Over the last few years, YouTube and Google's advertising products have taken steps to strengthen the requirements for monetization so that spammers, impersonators, and other bad actors can't hurt the ecosystem or take advantage of good creators. To apply for membership in the YouTube Partner Program, the thresholds were increased for channels to be deemed eligible. Channels must have generated 4,000 watch hours in the previous 12 months and have more than 1,000 subscribers.



Process for Admittance to YouTube Partner Program

Following application, YouTube's review team ensures the channel has not run afoul of monetization, content, and copyright policies prior to admitting them to the program. Only creators with sufficient history and demonstrated advertiser safety will receive access to ads and other monetization products. In changing these thresholds, YouTube has significantly improved the protections in place against impersonating creators.

More information on how Google protects its monetization services from abuse is discussed in the next section, "Google Ads & Disinformation".

# Google Advertising Products & Disinformation

## Background

Google provides multiple products to help content creators – website builders, video makers, and app developers – make money from doing what they love. Our advertising products enable creators to place ads within the content they make and manage the process by which they sell space in that content. In addition, we have products for advertisers to purchase that inventory across various content creators.

Google Ads and DV 360 help businesses of all sizes get their messages to the audiences they need to grow. These services are “front doors” for advertisers of all sizes to buy ads across Google’s monetization products and platforms – connecting them with billions of people finding answers on Search, watching videos on YouTube, exploring new places on Google Maps, discovering apps on Google Play, and more.

AdSense, AdMob, and Ad Manager, support content creators’ and publishers’ efforts to create and distribute their creations. We launched AdSense in 2003 to help publishers fund their content by placing relevant ads on their website. Over time, it has become a core part of our advertising products, serving more than 2 million website owners around the world.

Our ads and monetization products enable businesses of all sizes from around the world to promote a wide variety of products, services, applications, and websites on Google and across our partner sites and apps, making it possible for Internet users to discover more content they care about.

We also understand that the content of both ads and publisher sites needs to be safe and provide a positive experience for users. We aim to protect users and ensure a positive ad experience across our partner sites and apps as well as owned and operated properties like Maps and Gmail, by creating clear policies that govern what content can and cannot be monetized. When we create these policies, we think about our values and culture as a company, as well as operational, technical, and business considerations. We regularly review changes in online trends and practices, industry norms, and regulations to keep our policies up-to-date. And, we listen to our users’ feedback and concerns about the types of ads they see.

As we create new policies and update existing ones, we strive to ensure a safe and positive experience for our users. We also consider the impact that certain kinds of content will have on our advertisers and publishers. For example, some advertisers do not want their ads shown next to particular types of publisher content, and vice versa.

At the same time, we are mindful that the advertisers and publishers who use our services represent a broad range of experiences and viewpoints and we don’t want to be in the position of limiting those viewpoints or their ability to reach new audiences.

Oftentimes, these goals are in tension with one another. We aim for a balanced approach that prevents harm to our users by putting limits on the types of content we allow to be monetized without being overly restrictive, while also creating clear, enforceable, and predictable policies for advertisers and publishers.

We have a responsibility to balance the imperatives of making sure we leave room for a variety of opinions to be expressed, while preventing harmful or misrepresentative content on our advertising platforms.

## Tackling disinformation on Google’s advertising products

The considerations described above influence the policies we create for advertisers and publishers, and those policies are the primary way by which our advertising platforms implement the strategies to counter disinformation that we mention in the opening section of this paper, including:

- **Counteract Malicious Actors**
  - We look for, and take action against, attempts to circumvent our policies.
- **Give Users More Context**
  - “Why this ad” labels enabling users to understand why they’re presented with a specific ad and how to change their preferences so as to alter the personalization of the ads they are shown, or to opt out of personalized ads altogether.
  - In-ad disclosures and transparency reports on election advertising, which are rolling out during elections in the U.S., Europe, and India as our starting point.

Google’s policies to tackle disinformation on our advertising platforms favor an approach that focuses on misrepresentative or harmful behavior by advertisers or publishers while avoiding judgments on the veracity of statements made about politics or current events. To that end, we have developed a number of policies designed to catch bad behaviors, including many that can be associated with disinformation campaigns.

While we do not classify content specifically as “disinformation”, we do have a number of long-standing content policies aimed at preventing deceptive or low-quality content on our platforms. These policies complement, and build on, the strategies we outline in the opening section of this paper.

Each of these policies reflect a behavior-based approach to fighting deceptive content. Rather than make a judgment on specific claims, we enforce policies against advertiser and publisher behavior that is associated with misrepresentative or harmful content.

The policies described in this document are current as of the publication of this paper but are also subject to continual refinement and improvement to take into account emerging trends and threats to ensure the integrity of our platforms and the information we are providing to partners and users.

### Managing “scraped” or unoriginal content

In order to ensure a good experience for users and advertisers, we have policies for publishers that limit or disable ad serving on pages with little to no value and/or excessive advertising.<sup>35</sup> This results in a significant number of policy violations. In 2017, we blocked over 12,000 websites for “scraping,” duplicating and copying content from other sites, up from 10,000 in 2016.<sup>36</sup>

Also, Google Ads does not allow advertisements that point users to landing pages with insufficient original content. This includes content that is replicated from another source without providing any additional value in the form of added content or functionality. For example, a site that consists of news articles that are scraped from other sources without adding additional commentary or value to the user would not be allowed to advertise with us.<sup>37</sup>

## Misrepresentation

We have long prevented ads that intend to deceive users by excluding relevant information or giving misleading information about products, services, or businesses. This includes making false statements about the advertiser's identity or qualifications, or making false claims that entice a user with an improbable result.

Our policies on misrepresentation were extended to content that is available via our monetization products (AdSense, AdMob, and Ad Manager) in 2016, and are publicly available online.<sup>38</sup>

We made an additional update to our Google Ads and AdSense policies in 2018 to specifically state that it's not acceptable to direct content about politics, social issues, or matters of public concern to users in a country other than your own, if you misrepresent or conceal your country of origin or other material details about yourself or your organization.

## Inappropriate content

We also have long-standing policies to disallow monetization of shocking, dangerous, or inappropriate content on our advertising platforms, the details of which are publicly available online.<sup>39</sup> This includes derogatory content, shocking or violent content, or ads that lack reasonable sensitivity toward a tragic event.

## Political influence operations

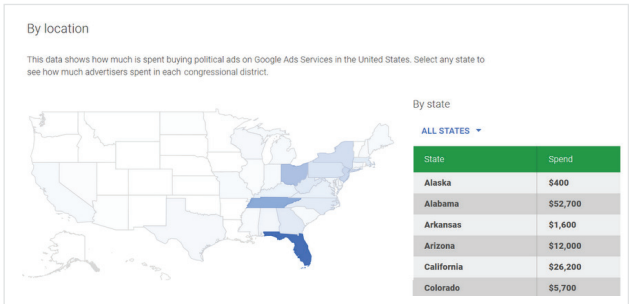
As discussed in a blog post in August 2018, we have also conducted investigations into foreign influence operations on our advertising platforms.<sup>40</sup> To complement the work of our internal teams, we engage independent cybersecurity experts and top security consultants to provide us with intelligence on these operations. Actors engaged in these types of influence operations violate our policies and we swiftly remove such content from our services and terminate these actors' accounts.

## Election integrity

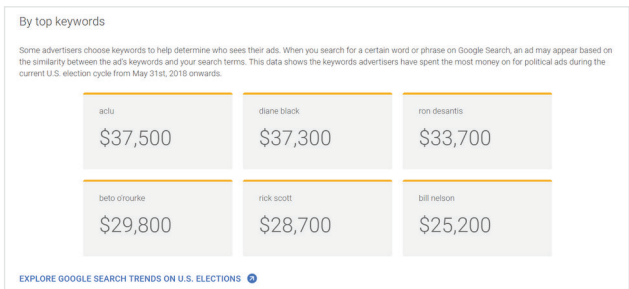
When it comes to elections, we recognize that it is critical to support democratic processes by helping users get important voting information, including insights into who is responsible for the political advertising content they see on our platforms.

Beginning with the 2018 U.S. Congressional midterm election, we require additional verification for anyone who wants to purchase an election ad on Google in the U.S., and require that advertisers confirm they are a U.S. citizen or lawful permanent resident.<sup>41</sup> In an effort to provide transparency around who is paying for an election ad, we also require that ad creatives incorporate a clear disclosure of who is paying for it. Additionally, we released a Transparency Report specifically focused on election ads.<sup>42</sup> This Report describes who is buying election-related ads on our platforms and how much money is being spent. We have also built a searchable library for election ads where anyone can find election ads purchased on Google and who paid for them.<sup>43</sup> In parallel, we updated our personalized ads policies to require verification for all advertisers who use our limited political affiliation options to target ads to users or to promote advertisers' products and services in the United States.<sup>44</sup>

As we look ahead to 2019, we are also planning to extend these election integrity efforts to other elections around the globe. Similar to our approach for U.S. federal elections, we will be requiring verification and disclosure for election ads in the European Union Parliamentary elections<sup>45</sup> and the Indian Lok Sabha election.<sup>46</sup> Ads that mention a political party, candidate, or current office holder will be verified by Google and be required to disclose to voters who is paying for the ad. We will also be introducing respective Political Ads Transparency Reports and searchable ad libraries for each of these elections that will provide more information about who is purchasing election ads, who is being targeted, and how much money is being spent.



The new political advertising section in our Transparency Report shows how much money is spent across states and congressional districts for U.S. federal elections



The political advertising section in our U.S. Transparency Report also shows which ads had the highest views, the latest election ads running on our platform, and explores specific advertisers' campaigns

In addition to these specific efforts, we're thinking hard about elections and how we continue to support democratic processes around the world, including by bringing more transparency to political advertising online, by helping connect people to useful and relevant election-related information, and by working to protect election information online. We will continue to invest in initiatives that build further on our commitment to election transparency.

## Consistent Enforcement

Our enforcement teams use a variety of robust methods to ensure content on our advertising platforms adheres to our policies, including machine learning, human review, and other technological methods. This approach is very similar to the one used by YouTube and described earlier in this paper. We have always relied on a combination of humans and technology to enforce our policies and will continue to do so.

When we find policy violations we take action to enforce our policies. Depending on the policy violation, this can include blocking a particular ad from appearing and removing ads from a publisher page or site. In cases of repeated or egregious violations, we may disable an account altogether.<sup>47/48</sup> Every year we publish a report on our efforts to remove bad actors from our advertising ecosystem.<sup>49</sup>

We also know that some content, even if it complies with our policies, may not be something that all advertisers want to be associated with. That's why, in addition to these policies, we provide advertisers with additional controls, and help them exclude certain types of content that, while in compliance with our policies, may not fit their brand or business. These controls let advertisers exclude certain types of content or terms from their video, display and search advertising campaigns. Advertisers can exclude whole categories of content such as politics, news, sports, beauty, fashion, and many others. Similarly, publishers can also review and block certain ads from showing on their pages, including by specific advertiser URL, general ad category like "apparel" or "vehicles", and sensitive ad category like "religion" or "politics".

## Conclusion

Tackling the propagation of false or misleading information is core to Google's mission and to ensuring our products remain useful to the billions of users and partners who utilize our services every day. While we have always fought against malicious actors' efforts to manipulate our systems and deceive our users, it's never been more important to thwart them and to ensure we provide our users with information worthy of the trust they have in our services.

As we have outlined in this paper, this is not a straightforward endeavor. Disinformation and misinformation can take many shapes, manifest differently in different products, and raise significant challenges when it comes to balancing risks of harm to good faith, free expression, with the imperative to serve users with information they can trust.

We believe that we're at our best when we improve our products so they continue to make quality count, to counteract malicious actors, and to give users context, as well as working beyond our products to support a healthy journalistic ecosystem, partner with civil society and researchers, and stay one step ahead of future risks.

We are constantly striving to make progress on these issues. This is by no means a solved problem, and we know that we have room to make progress. We welcome a constructive dialogue with governments, civil society, academia, and newsrooms on what more can be done to address the challenges of misinformation and disinformation and hope that this paper will be useful in sparking these conversations.

## References

- 1 <https://europe.googleblog.com/2015/10/introducing-accelerated-mobile-pages.html>
- 2 <https://www.blog.google/outreach-initiatives/google-news-initiative/digital-news-initiative-introducing/>
- 3 <https://g.co/newsinitiative>
- 4 <https://blog.google/topics/google-news-initiative/introducing-subscribe-google/>
- 5 <https://www.poynter.org/channels/fact-checking>
- 6 <https://jigsaw.google.com/>
- 7 <https://www.blog.google/technology/safety-security/update-state-sponsored-activity/>
- 8 <https://protectyourelection.withgoogle.com/>
- 9 <https://projectshield.withgoogle.com/>
- 10 <https://transparencyreport.google.com/political-ad>
- 11 <https://www.blog.google/around-the-globe/google-europe/supporting-european-union-parliamentary-elections/>
- 12 <https://india.googleblog.com/2019/01/bringing-more-transparency-to-indian.html>
- 13 On January 31st 2019, we have made a dataset of synthetic speech available to all participants in the third-party and independent 2019 ASVspoof challenge, which invites researchers from all around the world to test countermeasures against fake (or 'spoofed') speech. Blog post: <https://www.blog.google/outreach-initiatives/google-news-initiative/advancing-research-fake-audio-detection/>
- 14 More on the guidelines: <https://www.google.com/search/howsearchworks/mission/web-users/>
- 15 <https://www.blog.google/products/news/new-google-news-ai-meets-human-intelligence/>
- 16 E.g. for Autocomplete: <https://support.google.com/websearch/answer/7368877>
- 17 Link to Google News content policies: <https://support.google.com/news/producer/answer/6204050>
- 18 Link to our Search Quality Raters Guidelines: <https://static.googleusercontent.com/media/www.google.com/en//insidesearch/howsearchworks/assets/searchqualityevaluatorguidelines.pdf>
- 19 More information on our work to protect elections can be found in our opening section, [p.7-8](#)
- 20 <https://support.google.com/webmasters/answer/35769>
- 21 <http://infolab.stanford.edu/~backrub/google.html>
- 22 Webmasters can signal fact-check content to Google Search and Google News using dedicated HTML Code – more information about the technical and content criteria applying to these fact-checks here: <https://developers.google.com/search/docs/data-types/factcheck>
- 23 For more information: <https://twitter.com/searchliaison/status/1070027261376491520>
- 24 <https://datasociety.net/output/data-voids-where-missing-data-can-easily-be-exploited/>
- 25 In addition to these Community Guidelines, we have guidelines related to copyright, privacy, and impersonation that are not discussed in this paper. See our list of Community Guidelines here: <https://www.youtube.com/yt/about/policies/#community-guidelines>
- 26 <https://www.globalnetworkinitiative.org/>
- 27 <https://transparencyreport.google.com/youtube-policy/overview>
- 28 <https://support.google.com/youtube/answer/2801973?hl=en>
- 29 [https://support.google.com/youtube/answer/2801947?hl=en&ref\\_topic=2803176](https://support.google.com/youtube/answer/2801947?hl=en&ref_topic=2803176)
- 30 [https://support.google.com/youtube/answer/2801939?hl=en&ref\\_topic=2803176](https://support.google.com/youtube/answer/2801939?hl=en&ref_topic=2803176)
- 31 [https://support.google.com/youtube/answer/2802268?hl=en&ref\\_topic=2803176](https://support.google.com/youtube/answer/2802268?hl=en&ref_topic=2803176)
- 32 <https://youtube.googleblog.com/2018/07/building-better-news-experience-on.html>
- 33 <https://newsinitiative.withgoogle.com/>
- 34 <https://youtube.googleblog.com/2018/07/building-better-news-experience-on.html>
- 35 Read more: <https://support.google.com/adsense/answer/1346295?hl=en>
- 36 Read more: [https://support.google.com/adsense/answer/1348737?hl=en&ref\\_topic=1261918](https://support.google.com/adsense/answer/1348737?hl=en&ref_topic=1261918)
- 37 Read more: [https://support.google.com/adspolicy/answer/6368661?hl=en&ref\\_topic=1626336](https://support.google.com/adspolicy/answer/6368661?hl=en&ref_topic=1626336)
- 38 Google Ads Misrepresentation Policy: <https://support.google.com/adspolicy/answer/6020955> - AdSense misrepresentative content policy: <https://support.google.com/adspolicy/answer/1348688>
- 39 Google Ads Dangerous or Derogatory Policy: <https://support.google.com/adspolicy/answer/6015406>; AdSense Dangerous or Derogatory Policy: [https://support.google.com/adspolicy/answer/1348688?hl=en&rd=1#Dangerous\\_or\\_derogatory\\_content](https://support.google.com/adspolicy/answer/1348688?hl=en&rd=1#Dangerous_or_derogatory_content); AdSense Shocking Content policy: [https://support.google.com/adspolicy/answer/1348688?hl=en&rd=1#Shocking\\_content](https://support.google.com/adspolicy/answer/1348688?hl=en&rd=1#Shocking_content)
- 40 Read the post: <https://www.blog.google/technology/safety-security/update-state-sponsored-activity/>
- 41 Read more: <https://support.google.com/adspolicy/answer/9002729>
- 42 Read more: <https://transparencyreport.google.com/political-ads/overview>
- 43 Search the library: <https://transparencyreport.google.com/political-ads/library>
- 44 Read more: <https://support.google.com/adspolicy/answer/143465?#533>
- 45 Read more: <https://www.blog.google/around-the-globe/google-europe/update-our-work-prevent-abuse-ahead-eu-elections/>
- 46 <https://india.googleblog.com/2019/01/bringing-more-transparency-to-indian.html>
- 47 More information regarding Google Ads enforcement: [https://support.google.com/adspolicy/answer/7187501?hl=en&ref\\_topic=1308266](https://support.google.com/adspolicy/answer/7187501?hl=en&ref_topic=1308266)
- 48 More information regarding AdSense enforcement: <https://support.google.com/adspolicy/answer/7003627?hl=en>
- 49 Our most recent report from March 2018: <https://blog.google/technology/ads/advertising-ecosystem-works-everyone>

Google