

2023

Ads Safety Report

Google's work to prevent malicious
use of our ads platforms



Ads Safety Report

A look back at what we did in 2023 to keep our platforms safe for users, advertisers, and publishers.

Billions of people around the world rely on Google products to provide relevant and trustworthy information, including ads. That's why we have thousands of people working around the clock to safeguard the digital advertising ecosystem.

Our continued investment in policy development and enforcement enabled us to block or remove over 5.5 billion ads, restrict over 6.9 billion ads and suspend over 12.7 million advertiser accounts. This represents an account-level increase of nearly double from the previous year. And we blocked or restricted ads from serving more than 2.1 billion publisher pages and across more than 395,000 publisher sites, up from 143,000 in 2022.


The key trend in 2023 was the impact of generative AI. This new technology introduced significant and exciting changes and challenges to the digital advertising industry, and presents a unique opportunity to improve our enforcement efforts significantly. Our teams are embracing this transformative technology, specifically Large Language Models (LLMs), so that we can better keep people safe online.

Advertiser policy enforcement

Our policies are designed to support a safe and positive experience for our users, which is why we prohibit content that we believe to be harmful to users and the overall advertising ecosystem.

5.5 billion bad ads stopped in 2023

Below are the policies that we enforced the most in 2023:



Gen AI bolsters enforcement

Our safety teams have long used AI-driven machine learning systems to enforce our policies at scale. It's how, for years, we've been able to detect and block billions of bad ads before a person ever sees them. But, while still highly sophisticated, these machine learning models have historically needed to be trained extensively — they often rely on hundreds of thousands, if not millions of examples of violative content.

LLMs, on the other hand, are able to rapidly review and interpret content at a high volume, while also capturing important nuances within that content. These advanced reasoning capabilities have already resulted in larger-scale and more precise enforcement decisions on some of our more complex policies. Take, for example, our policy against [Unreliable Financial Claims](#), which includes ads promoting get-rich-quick schemes. The bad actors behind these types of ads have grown more sophisticated. They adjust their tactics and tailor ads around new financial services or products, such as investment advice or digital currencies, to scam users.

To be sure, traditional machine learning models are trained to detect these policy violations. Yet the fast-paced and ever-changing nature of financial trends make it, at times, harder to differentiate between legitimate and fake services, and quickly scale our automated enforcement systems to combat scams. LLMs are more capable of quickly recognizing new trends in financial services, identifying the patterns of bad actors who are abusing those trends, and distinguishing a legitimate business from a get-rich-quick scam. This has helped our teams become even more nimble in confronting emerging threats of all kinds.

To put the impact of AI on this work into perspective: last year more than 90% of our publisher page level enforcement started with the use of machine learning models, including our latest LLMs.

We've only just begun to leverage the power of LLMs for ads safety. Gemini, launched publicly last year, is Google's most capable AI model. We're excited to have started bringing its sophisticated reasoning capabilities into our ads safety and enforcement efforts.



Gemini

Our work to prevent fraud and scams

In 2023, [scams and fraud across all online platforms were on the rise](#). Bad actors are constantly evolving their tactics to manipulate digital advertising in order to scam people and legitimate businesses alike. To counter these ever-shifting threats, we quickly updated policies, deployed rapid-response enforcement teams, and sharpened our detection techniques.

- In November, we launched our [Limited Ads Serving policy](#), which is designed to protect users by limiting the reach of advertisers with whom we are less familiar. Under this policy, we've implemented a "get-to-know-you" period for advertisers who don't yet have an established track record of good behavior, during which impressions for their ads might be limited in certain circumstances — for example, when there is an unclear relationship between the advertiser and a brand they are referencing. Ultimately, Limited Ads Serving, which is still in its early stages, will help ensure well-intentioned advertisers are able to build up trust with users, while limiting the reach of bad actors and reducing the risk of scams and misleading ads.
- A critical part of protecting people from online harm hinges on our ability to respond to new abuse trends quickly. Toward the end of 2023 and into 2024, we faced a targeted campaign of ads featuring the likeness of public figures to scam users, often through the use of deepfakes. When we detected this threat, we created a dedicated team to respond immediately. We pinpointed patterns in the bad actors' behavior, trained our automated enforcement models to detect similar ads, and began removing them at scale. We also updated our [misrepresentation policy](#) to better enable us to rapidly suspend the accounts of bad actors.

Overall, we blocked or removed over

1 billion

advertisements for violating our policy against abusing the ad network, which includes promoting malware

and **206.5 million** advertisements for violating our misrepresentation policy, which includes many scam tactics, and **273.4 million** advertisements for violating our financial services policy.

The fight against scam ads is an ongoing effort, as we see bad actors operating with more sophistication, at a greater scale, and using new tactics such as deepfakes to deceive people. We'll continue to dedicate extensive resources, making significant investments in detection technology and partnering with organizations like the Global Anti-Scam Alliance and Stop Scams UK, to facilitate information sharing and protect consumers worldwide.

Restricted ads

The policies below cover content that is sometimes legally or culturally sensitive. Online advertising can be a powerful way to reach customers, but in sensitive areas, we also work hard to avoid showing ads when and where they might be inappropriate. For that reason, we allow the promotion of the content below, but on a limited basis. These promotions may not show to every user in every location, and advertisers may need to meet additional requirements before their ads are eligible to run.

6.9 billion
restricted ads in 2023

Legal requirements	418.1M
Financial services	282M
Gambling and games	144.7M
Adult content	114.4M
Copyright	87.2M
Healthcare and medicines	85.6M
Alcohol	42.8M

Investing in election integrity

Political ads are an important part of democratic elections. Candidates and parties use ads to raise awareness, share information, and engage potential voters. In a year with several major elections around the world, we want to make sure voters continue to trust the election ads they may see on our platforms. That's why we have long-standing identity verification and transparency requirements for election advertisers, as well as restrictions on how these advertisers can target their election ads. All election ads must also include a "paid for by" disclosure and are compiled in our publicly available [transparency report](#). In 2023, we verified more than 5,000 new election advertisers and removed more than 7.3M election ads that came from advertisers who did not complete verification.

Last year, we were the [first tech company to launch](#) a new disclosure requirement for election ads containing synthetic content. As more advertisers leverage the power and opportunity of AI, we want to make sure we continue to provide people with the greater transparency and the information they need to make informed decisions.

Additionally, we've continued to enforce our [policies](#) against ads that promote demonstrably false election claims that could undermine trust or participation in democratic processes.

Publisher enforcement

We support the open web by helping publishers monetize their content. This content is subject to certain publisher policies and restrictions.

2.1 billion pages taken action against in 2023

Below are the areas that we enforced the most in 2023:



Staying nimble and looking ahead

When it comes to ads safety, a lot can change over the course of a year, from the introduction of new technology such as generative AI to novel abuse trends and global conflicts. And the digital advertising space has to be nimble and ready to react. That's why we are continuously developing new policies, strengthening our enforcement systems, deepening cross-industry collaboration, and offering more control to people, publishers, and advertisers.

In 2023, for example, we launched the [Ads Transparency Center](#), a searchable hub of all ads from verified advertisers, which helps people quickly and easily learn more about the ads they see on Search, YouTube and Display. We also updated our [suitability controls](#) to make it simpler and quicker for advertisers to exclude topics that they wish to avoid across YouTube and Display inventory.

Overall, we made 31 updates to our Ads and Publisher policies.

Though we don't yet know what the rest of 2024 has in store for us, we are confident that our investments in policy, detection and enforcement will prepare us for any challenges ahead.

[Learn more about our policies](#)



[Advertiser](#)



[Publisher](#)