

Report of Systemic Risk Assessments

Table of Contents

Table of Contents	1
1. Introduction	6
Our Philosophy	6
About this Report	7
Scope and Purpose	7
Findings	7
Next Steps	7
Structure of the Report	8
Our Risk Assessment Methodology	9
Assessing Risk for each VLOSE and VLOP	9
2. Background	11
Maintaining User Trust and Safety	11
Our Commitments	11
Investing in Systemic Risk Prevention	12
Promoting Trustworthy Content and User Safety	15
One: Protecting Users from Harm	16
Preventing Harm with Safety by Design	16
Preparing for the Unexpected	17
Designing Appropriate Content Policies	17
Reviewing Content Policies and Practices	18
Counterbalancing Risk	19
Detecting and Responding to Harmful Content at Scale	20
Handling Government Removal Requests	20
Evaluating Content Across Languages	22
Two: Delivering Reliable Information	24
Surfacing Quality Information	24
Using Recommender Systems	24
Fighting Misinformation	25
Addressing the Risks and Opportunities of Artificial Intelligence	27
Equipping Users	30
Three: Partnering to Create a Safer Internet	31
Partnering for Information Quality	31
Consulting with Experts	31
Sharing Tools and Technology	32
Collaborating with Companies and Stakeholders	33
Developing Best Practices	34
Setting High Standards for Advertising	35

<u>3. Methodology</u>	<u>38</u>
<u>Introduction</u>	<u>38</u>
<u>Step One: Classifying Risk</u>	<u>39</u>
<u>Step Two: Identifying Risk</u>	<u>40</u>
<u>Engaging Stakeholders</u>	<u>41</u>
<u>Step Three: Assessing Inherent Risks</u>	<u>43</u>
<u>Step Four: Assessing Preparedness</u>	<u>44</u>
<u>Taking a Human Rights-Based Approach</u>	<u>45</u>
<u>Step Five: Identifying Additional Mitigations</u>	<u>45</u>
<u>Step Six: Reporting the Results</u>	<u>46</u>
<u>4. Results of the Assessment</u>	<u>47</u>
<u>Search</u>	<u>48</u>
<u>Description of Service and Associated Risk Profile</u>	<u>48</u>
<u>Systemic Risk Assessment Results and Associated Observations</u>	<u>51</u>
<u>Removing Content</u>	<u>52</u>
<u>Removing Illegal Content</u>	<u>52</u>
<u>Addressing Violations of Intellectual Property Rights</u>	<u>52</u>
<u>Detecting, Removing, and Reporting CSAM</u>	<u>53</u>
<u>Removing NCEI and ISPI</u>	<u>54</u>
<u>Investing in Search Information Quality</u>	<u>54</u>
<u>Addressing Sensitive, Harmful, and Policy Violative Content</u>	<u>55</u>
<u>Informing Users</u>	<u>55</u>
<u>Providing SafeSearch</u>	<u>56</u>
<u>Tailoring our Content Policies</u>	<u>56</u>
<u>Addressing Civics Misinformation</u>	<u>57</u>
<u>Respecting Freedom of Opinion, Expression, Media Pluralism, and Civic Discourse</u>	<u>58</u>
<u>Addressing Disinformation</u>	<u>58</u>
<u>Service Design</u>	<u>59</u>
<u>Addressing Unfair Commercial Practices and Fraudulent Content about a Business</u>	<u>59</u>
<u>Respecting Privacy</u>	<u>60</u>
<u>Protecting Children’s Rights</u>	<u>61</u>
<u>Obtaining Age Assurance</u>	<u>61</u>
<u>Enabling Parental Control</u>	<u>62</u>
<u>Providing Ads Protections</u>	<u>62</u>
<u>Providing SafeSearch on by Default</u>	<u>62</u>
<u>Maps</u>	<u>63</u>
<u>Description of Service and Associated Risk Profile</u>	<u>63</u>
<u>Systemic Risk Assessment Results and Associated Observations</u>	<u>65</u>
<u>Content Moderation</u>	<u>66</u>
<u>Removing Illegal Content</u>	<u>66</u>

Addressing Content that Violates our Policies	66
Developing Content Policy	66
Enforcing Content Policy	67
Undertaking Automated Detection and Removal	67
Undertaking Human Review	68
Undertaking Enforcement Proactively	68
Posting Restrictions for Repeat Violators	69
Assessment Results for Specific Content Risks	69
Protecting Civic Discourse	69
Protecting Consumers and the Freedom to Conduct a Business	69
Respecting Freedom of Opinion, Expression, and Media Pluralism	70
Service Design	71
Respecting Privacy	71
Protecting Users of Maps	71
Protecting Contributors to Maps	71
Addressing Risks Relating to Images on Maps	71
Protecting Children’s Rights	72
Google Play	73
Description of Service and Associated Risk Profile	73
Systemic Risk Assessment Results and Associated Observations	75
Content Moderation	75
Removing Illegal Content	75
Addressing Content that Violates our Policies	76
Maintaining Developer Policies	77
Developing Policy	77
Enforcing Policy	78
Addressing Specific Content-Related Risks	78
Preventing Review Bombing and Ensuring Rating and Review Integrity	79
Protecting Civic Discourse	79
Platform Design	80
Protecting Privacy	80
Protecting Children’s Rights	81
Maintaining Additional Policies for Minors	81
Providing a Teacher-Approved Program	82
Obtaining Age Assurance	82
Enforcing Content Ratings and Content Restrictions	82
Shopping	84
Description of Service and Associated Risk Profile	84
Systemic Risk Assessment Results and Associated Observations	85
Content Moderation	85

<u>Removing Illegal Content</u>	85
<u>Identifying and Blocking Illegal Products and Services</u>	85
<u>Prohibiting and Detecting Violations of Intellectual Property Rights</u>	86
<u>Addressing Content that Violates our Policies</u>	86
<u>Maintaining Shopping Policies</u>	86
<u>Maintaining Guardrails for User-Contributed Content</u>	87
<u>Preventing Unfair Commercial Practices</u>	87
<u>Preventing Fraudulent Business Information</u>	88
<u>Service Design</u>	88
<u>Respecting Privacy</u>	88
<u>Vetting Merchants</u>	89
<u>Monitoring Merchants and Listings</u>	89
<u>Protecting Children's Rights</u>	89
<u>YouTube</u>	91
<u>Description of Service and Associated Risk Profile</u>	91
<u>Systemic Risk Assessment Results</u> <u>and Associated Observations</u>	93
<u>Content Moderation</u>	94
<u>Removing Illegal Content</u>	94
<u>Two Examples: Terrorist or Violent Extremist Content and Child Sexual Abuse</u> <u>Material (CSAM)</u>	94
<u>Identifying and Removing Violent Extremist Content</u>	94
<u>Detecting, Removing, and Reporting CSAM</u>	95
<u>Prohibiting and Detecting Infringement of Intellectual Property Rights</u>	96
<u>Addressing Content that Violates our Policies</u>	96
<u>Developing Policy</u>	97
<u>Providing EDSA Exceptions</u>	97
<u>Enforcing Policy</u>	98
<u>Undertaking Automated Detection and Removal</u>	98
<u>Maintaining a Priority Flagger Program</u>	98
<u>Enforcing a Three-Strike System for Repeat Violators</u>	99
<u>Measuring Success: Violative View Rate</u>	99
<u>Elevating Authoritative Sources</u>	101
<u>Providing Information Panels with Topical Context</u>	101
<u>Addressing Specific Content Risks</u>	101
<u>Addressing Misinformation and Disinformation</u>	101
<u>Addressing Public Health Related Violative Content</u>	102
<u>Addressing Civic-Discourse-Related Violative Content</u>	103
<u>Detecting and Removing Harassment and Bullying in YouTube Comments</u>	104
<u>Prohibiting and Removing Hate Speech</u>	105
<u>Service Design</u>	106
<u>Respecting Privacy</u>	106

Protecting Children's Rights	106
Maintaining Guardrails for Children's Access to Content	107
Addressing Potentially Addictive Behaviour in Children	108
Protecting Children's Data	108
Protecting Children's Safety in YouTube Comments	108
Promoting Equity	110
5. Conclusions	111
Annex A: Full List of Risk Statements	113
Illegal Content	113
Fundamental Rights	114
Freedom of Expression and Information	114
Pluralism in the Media	114
Privacy and Data Protection	114
Human Dignity	115
Consumer Protection	115
Child Rights	115
Equality and Non-Discrimination	115
Freedom to Conduct a Business	116
Civic Discourse	117
Civic Discourse and Elections	117
Public Security	117
Public Health	118
Public Health	118
Gender-based Violence	118
Physical and Mental Wellbeing	118
Annex B: List of Mitigations	119
Background	119
Article 35 Mitigation Types	119
Mitigations Applicable to Multiple Services	121
Google Maps	122
Google Play	122
Google Search	123
Shopping	124
YouTube	125
Annex C: List of Consultations	127
Prior Stakeholder Engagement	127
Stakeholder Engagement to Inform the Systemic Risk Assessment	129

1. Introduction

Our Philosophy

People use Google services to access, share, and contribute to a vast and ever-expanding universe of information on the web. They turn to Google in moments that matter, including when looking for information about natural disasters or breaking news. They also use our services to consume digital content: watching videos, playing games, shopping for products, listening to music, and reading books. We do the work of organising and serving the information users are seeking in the most usable format, because information is only as useful as it is accessible. As we have long said, our mission is to organise the world's information and make it universally accessible and useful.

We're proud that our services expand knowledge, power businesses, and provide opportunities for expression and connection. The internet amplifies and makes available to the world the benefits of technology and humanity's collective knowledge, but like any forum, it can also reflect prejudice, hate, and greed if left unchecked. The range of potential effects, and the scale of their impacts, demand that we provide access to relevant and trustworthy information and content and opportunities for free expression to users across our services, while minimising the inherent risk of abuse and harm.

We have felt that responsibility since the beginning. In their [first letter to shareholders](#), our founders described Google's goal to "develop services that significantly improve the lives of as many people as possible." To that end, we have long designed services and policies, built teams, and developed technologies with the wellbeing of users in mind. That commitment has become increasingly important as more users have come to trust and depend on our largest services.

Maintaining a diverse, high-quality, and thriving digital ecosystem is also a business imperative. The ability to access information on our services as well as the quality and safety of our services are directly linked to our ability to attract users, which in turn is critical to our continued success as a business.

User and societal safety is a dynamic challenge without simple answers. It requires a collective effort. So we welcome input from experts, civil society, the people who use our services, and governments. We view the systemic risk assessment under the Digital Services Act (DSA) as the beginning of close engagement with the European Commission and other relevant stakeholders on these important issues. As this report will show, we have historically been highly attuned to the systemic risks identified in the DSA, but advancements in the technologies and techniques that bad actors use mean that there is an ongoing need to identify and mitigate emerging risks.

This report describes the methodology and results of the first systemic risk assessments we have undertaken for our designated very large services to meet the requirements of Article 34 of the DSA, and the mitigation measures satisfying Article 35 of the DSA. We welcome the opportunity to present these results and demonstrate how our approach helps keep users safe online and furthers the EU's aspirations of an information society in which the rights of all users of digital services are protected. As our founders put it in 2004, “[w]e believe a well functioning society should have abundant, free and unbiased access to high quality information. Google therefore has a responsibility to the world.”

About this Report

Scope and Purpose

This report is issued by Google Ireland Limited. The report and the appendices meet the requirement under Article 42(4) of the DSA that the providers of very large online search engines (“VLOSEs”) or very large online platforms (“VLOPs”) make available to the Digital Services Coordinator of establishment and the European Commission a report setting out: (a) the results of the systemic risk assessment undertaken to meet the requirements of Article 34 of the DSA; (b) the mitigation measures put in place pursuant to Article 35(1) of the DSA; and (c) information about the consultations conducted in support of the risk assessments and design of the risk mitigation measures.

Article 34 of the DSA requires VLOSEs and VLOPs to identify, analyse, and assess enumerated systemic risks in the EU stemming from the design or functioning of their services and related systems, while Article 35 requires providers of VLOSEs and VLOPs to put in place reasonable, proportionate, and effective mitigation measures to address systemic risks identified in the Article 34 risk assessment. In scope for this report are Google’s designated VLOSE (Google Search) and VLOPs (Google Maps, Google Play, Shopping, and YouTube).

Findings

We concluded that our mitigation measures generally address the highest inherent risks and are well tailored to the purposes of the Google services we assessed. However, we also concluded that there are several areas where we can enhance our mitigations, such as new or enhanced user reporting and appeals channels, improved translation and content moderation across languages, and more robust efforts to address disinformation and misinformation. We found that risks associated with highly motivated bad actors seeking to misuse our services remain a cause of concern, and addressing them requires that our mitigations keep pace with the evolving social context and the changing nature of technology and risks online.

Next Steps

We have already begun planning for future annual systemic risk assessments. We expect to further embed the systemic risk assessment process into our broader risk management frameworks and systems, address new technologies, and further test our approaches with users, independent experts, civil society organisations, and other stakeholders.

Structure of the Report

This report has five sections:

- **Background:** We describe how Google uses service design and content moderation to create and maintain services that balance maximising the benefits they provide with minimising potential negative externalities.
- **Methodology:** We describe the six-step methodology used to conduct the systemic risk assessments.
- **Results:** We share the results of the systemic risk assessments conducted for each of our VLOPs (i.e., Google Maps, Google Play, Shopping, and YouTube) and our VLOSE, (i.e., Google Search). Each section includes:
 - Discussion of the identification and assessment of the most important inherent and residual risks.
 - Description and assessment of our long standing content policies, safety- and private-by-design practices, and other measures designed to mitigate systemic risk.
 - Mitigation enhancements that represent additional commitments by Google to further address systemic risk in the EU and, in many cases, globally. Taken in combination with our existing measures, these enhancements help ensure that our mitigations are reasonable, proportionate, and effective.

Throughout each VLOP and VLOSE section, we describe how the internal and external factors articulated in Article 34(2) of the DSA and regional or linguistic considerations had an impact on the assessment of risks or mitigations.

- **Conclusion:** We provide observations on the future of systemic risk assessments at Google, in the EU, and beyond.
- **Appendices:** We outline more details about the systemic risk assessment.
 - A complete list of risk statements for each VLOSE and VLOP.
 - A list of the mitigations being adopted pursuant to Article 35 of the DSA.
 - A list of consultations used in support of the risk assessment and the design of the risk mitigation measures.

Our Risk Assessment Methodology

Article 34 of the DSA requires VLOSEs and VLOPs to identify, analyse, and assess systemic risks in the EU stemming from the design or functioning of their services and their related systems or from the use of their services. We developed our systemic risk assessment methodology by combining the systemic risk assessment requirements of the DSA with proven risk assessment methodologies, such as those used to assess enterprise risk, human rights risk, compliance risk, and systemic risk assessments in other sectors.

Assessing Risk for each VLOSE and VLOP

Step One: Classification. We established 40 “risk statements” across the four categories of systemic risk in the DSA. The risk statements are plain language articulations of the potential adverse impacts for each risk category and provide the focus for each systemic risk assessment.

Step Two: Identification. We identified the risk drivers and exposure scenarios that may lead to inherent risk for each risk statement and pinpointed the quantitative and qualitative insights needed to assess systemic risk.

Step Three: Assessment of Inherent Risks. We assessed each risk statement according to the potential severity of the negative impacts that could arise from that risk, and the probability or frequency of the risk’s occurrence. Combined, these elements produce an estimate of the inherent risk—the risk absent our risk reduction efforts. That estimate was in the later steps then used as the foundation to review how well we address each risk. In practice, the inherent risk does not reflect actual risk on the service because all services are launched with risk mitigations.

Step Four: Assessment of Preparedness. We reviewed the mitigations (e.g., policies, controls, enforcement practices, and other measures) we have in place to address each risk and assessed the level of our preparedness, resulting in an estimate of residual risk (i.e., the risk after our mitigation efforts) for each risk statement. We considered the extent to which the combination of mitigations prevents or significantly addresses adverse impacts of the risk.

Step Five: Additional Mitigations. We used the results of the risk assessment to identify where additional mitigations are appropriate. We identified these additional measures to ensure that there are reasonable, proportionate, and effective mitigations in place to address the specific systemic risks we identified, consistent with Article 35 of the DSA.

Step Six: Reporting. We disclose the results of the systemic risk assessments in this report, including a discussion of the most important inherent and residual risks and our efforts to address them. We will publish this report (subject to removal of confidential information) in due course, consistent with the requirements of Articles 35 and 42 of the DSA.

We discuss this methodology in more detail in Section 3.

2. Background

Maintaining User Trust and Safety

While the DSA's formal systemic risk assessment paradigm is new, our commitment to examining and addressing the impacts our services can have on societal risks is not. We have built teams, service protections, tools, and partnerships to address risks arising from the increasing use of the internet by society, and risks that may result from the use of our services.

We begin with an overview of key teams at Google that work to promote user safety and combat potential harm, then detail our approach to preventing risk at scale.

Our Commitments

Our approach to maintaining user trust and safety on a global scale stems from our overall philosophy that we have a responsibility for the impacts of our services on people and societies. That perspective is reflected in a number of policy frameworks:

- Our [Human Rights Policy and White Paper](#), which set out our commitment to respecting human rights and upholding the standards established in the United Nations Guiding Principles on Business and Human Rights (UNGPs).
- Our [Responsible AI Principles](#), which describe our commitment to developing technology responsibly and work to establish specific application areas we will not pursue.
- Our [Information Quality and Content Report](#), which outlines the key considerations that guide our product, policy, and enforcement decisions.
- Our [Transparency Center](#), which outlines the content policies that help keep users safe from harm and abuse, as well as information about how we develop and enforce those policies.
- Our [Privacy and Terms](#) Center, which sets out our Privacy Policy, Terms of Service, Privacy and Security Principles, and other relevant guides and resources.

We provide regular updates on [The Keyword](#), our official blog for product and technology announcements, news, and stories.

Investing in Systemic Risk Prevention

Each of our services seeks to help users while keeping them safe from potential harms. Within each VLOP and VLOSE are well-developed functions that refine and enforce product policies, and design and maintain features aimed at avoiding and/or mitigating risks to our users. But in addition to these important service-level efforts, we have made significant investments in another layer of systemic risk mitigation by building central, cross-service teams. These organisations lead our efforts to mitigate specific types of negative impacts potentially caused by our services. These teams of subject matter experts match the systemic risks identified in Article 34 of the DSA because we have long strived to address these risks. The descriptions below of a sample of those teams are evidence of our specific risk-focused efforts:

- **Trust and Safety:** Together, our Trust and Safety teams are located worldwide, are fluent in multiple languages, and are able to carefully evaluate flagged content 24 hours a day, seven days a week. Our teams also monitor emerging trends to address new harm vectors before they can become a larger issue.
 - **Google Trust and Safety:** We pioneered the now industry-wide practice of investing in Trust and Safety specialists who are trained to analyse bad actors, abusive practices, content issues, and the effectiveness of existing policies. Today, our Trust and Safety teams consist of tens of thousands of experts, specialists, and engineers working to keep people safe online by using the latest technology to enforce our policies and moderate content. These teams partner with external experts and across Google to carry out our mission to keep people safe online and protect our services and products from abuse.
 - **YouTube Trust and Safety:** YouTube has built its own Trust and Safety team, with expertise in addressing the unique content challenges that arise on an open video-first service. Like Google's company-wide Trust and Safety organisation, YouTube Trust and Safety partners with members of our legal, operations, public policy, product management, and engineering teams to develop innovative ways to combat harmful content. Hundreds of hours of new content are uploaded to YouTube every minute, and we use a combination of people and automated systems to detect problematic content at scale.
- **Human Rights:** The Human Rights program is a central function responsible for ensuring that we are meeting our [human rights commitments](#) across all functions, products, and services. The program advances company-wide strategy on civil and human rights, advises product teams on potential civil and human rights impacts, conducts human rights due diligence, and engages external experts and stakeholders. The program also partners with our Responsible Innovation team within the Office of Compliance and Integrity to undertake human rights due diligence of our advanced technologies to help meet our commitment to the [AI Principles](#).
- **Privacy, Safety, and Security:** The Privacy, Safety, and Security (PSS) organisation combats digital threats to users and is committed to keeping the internet as a whole protected. We do this because we are an internet company, and our fate is tied to the fate of the internet. So we do not just design solutions to protect our users, we eliminate entire classes of threats from being effective on our services and products and across the internet.

PSS comprises industry-leading experts focused on protecting users and data, improving governance and assurance practices related to security, and increasing our technical and operational capabilities. PSS develops and implements automatic protections from bad actors in the data and security space across our services. Our PSS efforts include the following focus points:

- **Privacy:** Our Privacy program teams drive strategy for and provide leadership on Google's privacy priorities. The central Privacy program teams are responsible for administering privacy policies, training, and documentation that ensure that our products and services protect the privacy of our users. We have also embedded privacy teams and specialists in product areas to ensure that privacy goals are part of product work, and to ensure that we maintain a consistent and high standard of privacy protection and support across the company. Central and product privacy specialists coordinate across the company in working groups that focus on privacy issues that are relevant to particular products or sectors and track best practices and developments relating to particular policy topics. Our privacy subject matter experts also oversee privacy review processes to verify that our services and products vigilantly protect the privacy of our users.
- **User Protection:** Within our User Protection framework, our Threat Analysis Group (TAG) is responsible for countering threats from government-backed attackers, coordinated information operations, and serious cybercrime networks. TAG actively monitors threat actors and studies the evolution of their tactics and techniques, using research to continuously improve the safety and security of our products, improve Google's defences, and protect users.

TAG shares intelligence with our industry peers and publicly releases information about the operations it disrupts via public bulletins highlighting the group's work. For example, TAG has been closely tracking and disrupting campaigns targeting individuals and organisations in Ukraine, and frequently [publishes reports on Russian threat actors](#). The group also works closely with product teams to detect and remove malicious ads, videos, or channels that may be spreading disinformation, malware, or other types of cyber threats (see [examples](#)).

Complementing the work of TAG, our Account and Device Integrity (ADI) team within User Protection keeps users safe by ensuring products interact with legitimate users and devices. ADI's technology works to ensure that accounts and devices have access to Google products and services in ways that are proportional to their demonstrated integrity. In addition, ADI's offering limits opportunities for accounts to be created, compromised, or operated at scale to abuse our products or violate the privacy and security of people who use our services.

- **Civics:** Our Civics team works across our services, addressing threats to democratic participation in partnership with Trust and Safety specialists. The Civics team oversees products, initiatives, and promotional efforts that aim to safeguard the integrity of elections-related information and provide users with candidate information from authoritative sources. These teams also provide 24/7 support to triage emergent issues during elections.
- **Health:** People come to Google Search daily with health-related questions. Because of the ties between these queries and our users' health and wellbeing, we have prioritised building products to empower people with accurate, actionable health information. In 2019, we established improving health-related information as a key goal for the company. To implement this goal, we have recruited experts with decades of experience in health care, public health, and life sciences who help us translate clinical knowledge into product impact. Many of our product areas have policies prohibiting content

that contradicts well-established medical consensus, and our Clinical Team helps enforcement teams calibrate medical claims and ensure we are not exposing users to harmful medical misinformation.

During the COVID-19 pandemic, our users needed to quickly access high-quality information on topics relating to a rapidly evolving public health crisis. We supported public health officials through projects such as Exposure Notifications and Community Mobility Reports. We also lift up accurate and timely information on COVID-19 vaccines, fight misinformation, and support vaccine equity.

- **Kids and Families:** Our Kids and Families program includes a Kids and Family Steering Committee, which brings together executives and leaders from relevant services. It also includes a central team tasked with managing minors' accounts, creating age-appropriate experiences across our services, and advancing child safety protections. The program and its staff have built on years of input from experts and research insights to build tools and features that empower kids and teens while also giving families the ability to exercise choice over their children's relationship with technology. The results are products, features, and policies such as [YouTube Kids](#), [Assistant for Families](#), [Family Link](#), and [Google Play Families Policies](#).

These teams and experts are some of the key groups that partner with other teams across Google to assess and mitigate systemic risks. Their work helps us make good on our commitments to protect users from harm, deliver reliable information, and partner to create a safer internet.

Promoting Trustworthy Content and User Safety

Three core concepts guide our approach to providing access to trustworthy information and content while keeping users protected.

- **Protect users from harm.** We keep users and society safe through built-in protections utilising the latest technology that enable us to prevent, detect, and respond to illegal and harmful content.
- **Deliver reliable information.** We enable confidence by delivering reliable information and best-in-class tools that give additional context and put users in charge of evaluating content.
- **Partner to create a safer internet.** We scale our industry-leading practices to help keep users safe online through proactive partnership with experts and organisations to both inform and share our resources and technologies.

While we pursue these principles in all of our endeavours, we also recognise that working towards user trust and safety requires constant adaptation to changing social context, evolving threats, and new techniques employed by bad actors. We can never bring the threat of systemic risks to zero, but these principles guide our efforts to constantly increase trust and safety across all of our services.

One: Protecting Users from Harm

We work hard to keep users and society safe through built-in protections that enable us to prevent, detect, and respond to illegal and harmful content.

Preventing Harm with Safety by Design

Our first line of defence is the set of safety features we build into our products to protect user data and prevent abuse.

We present risk assessment results for our four VLOPs and our VLOSE as one report because many of the most effective protections we offer to users are implemented at the Google account level, and these protections are effective across our different service offerings. These account design features protect users whether they are browsing Google Search or downloading books on Google Play. And because we scale privacy and security solutions across all our services, we are able to minimise the number of times our services collect user data and the number of places we store that data.

Clear account settings options, robust account verification, and a secure sign-in process are fundamental to user safety and data security. Strong protections around these processes help guard user data from bad actors, empowering users and their family members to interact with our services the way that they wish. We invest in protecting these processes because they are the primary entry points for many risks. That's why we have developed features like Google's 2-Step verification, which requires a second layer of verification after a user enters a password, and helps guard against compromised passwords.

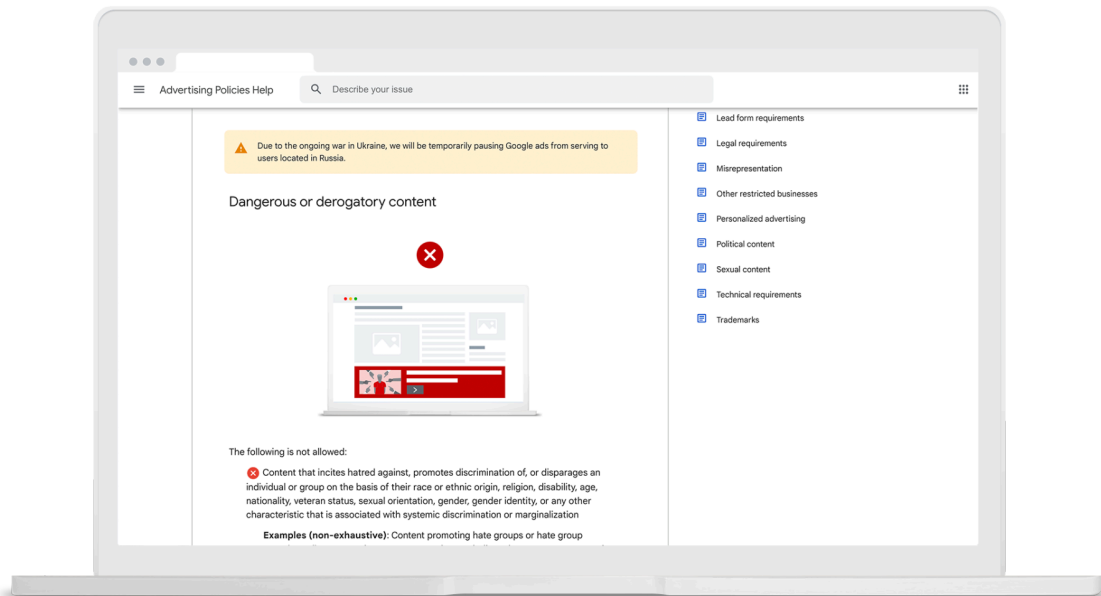
We recently began [rolling out passkeys](#) across Google Accounts as an easier and more secure way to sign in to apps and websites, and a major step towards a "passwordless" future. Passkeys let users sign in to apps and sites the same way they unlock their devices: with a fingerprint, a face scan, or a screen lock PIN. Unlike passwords, passkeys are resistant to online attacks like phishing, making them more secure than solutions like SMS one-time codes.

We also apply protections for signed-in and signed-out users who we believe are minors, and have engineered easy management of ads preferences and privacy settings through the My Ads Center. These protections, and many others, are designed as an integral part of our services, making it simple and quick for our users to benefit from advanced security infrastructure.

We also build services that consider safety at the outset and incorporate safety considerations into service design. For example, on Google Play, we reduce the risk of "review bombing" and sham ratings by using a percentage of the most recent reviews, not the average of all the ratings, to determine the overall rating for an app.

Preparing for the Unexpected

Other product policies and protections are focused on less likely, but potentially serious events. These policies, like many others described in this report, permit us to more nimbly respond to unexpected events. For example, Google Ads' sensitive events framework is designed to prevent ads that potentially profit from or exploit a sensitive event, such as a natural disaster, public health emergency, act of terrorism, conflict, or act of mass violence. We disallow ads that seek to profit from a tragic event with no discernible benefit to users, engage in price gouging that restricts access to vital supplies, or use keywords related to a sensitive event to drive traffic. We enforced our sensitive event framework in response to the war in Ukraine, [prohibiting ads that dismiss or condone the war](#).



Ads prohibited under our Sensitive Events policy

Designing Appropriate Content Policies

We design content policies across our services to protect users from harm and improve the intended use and function of each service for the benefit of our users. Our content policies, which are publicly available, articulate the purpose and intended use of each service to which they apply. They explain what types of content and conduct are not allowed, and the process by which a piece of content, or the user responsible for it, may be removed from the service. Our content policies have long been online and accessible to our users, and we regularly update them as our services evolve and new threats arise. You can find them [here](#) for [Google Search](#), [Google Maps](#), [Google Play](#), [Shopping](#), and [YouTube](#). Additionally, ads may be presented on these services, which have distinct [Google Ads](#) content policies.

We carefully tailor the rules about allowable content on each service according to the core purpose of that service and available levers to enforce the rules. For example, Search is intended to facilitate the exploration of a broad range of information from a wide variety of sources on the open web. Search's objective is to [maximise access to information](#), and we remove web results from Search in only very limited and clearly-defined circumstances. When listings and other information are presented as Search features (like featured snippets), however, users may interpret the information as having greater quality or credibility, and we apply more restrictive policies.

By contrast, our advertising services have policies that restrict certain types of harmful content because we do not believe the digital advertising ecosystem should profit from the sale of harmful or illegal content or experiences. Similarly, because Maps is designed to be a source of reliable information about places and experiences, its policies place a greater emphasis on accuracy, authenticity, and relevance.

YouTube's policies support the interests of creators generating expressive content, viewers who come to YouTube to watch user-generated content, and society at large. YouTube's policies give creators the freedom to share a broad range of experiences and perspectives through video, but because it also hosts and serves user-generated content, YouTube has different content policies than Search does.

Reviewing Content Policies and Practices

Promoting high-quality content and responding to harmful content is a dynamic challenge that requires constant adaptation. To help us identify emerging harms and gaps in our existing policies, we consider expert input, user feedback, and regulatory guidance. This is part of a continuous risk review and feedback process that each VLOP and VLOSE engages in with respect to all of its policies. We conduct research into the evolving tactics deployed by bad actors, safety trends observed across different services, and emerging cultural issues that require further observation. When we identify significant trends, we review existing policies and amend them to provide better tailored protections to users.

For example, [YouTube regularly reviews its policies](#) to make sure that they are effective at preventing real-world harm, and to ensure they properly address changes occurring both on and off our service.

YouTube works directly with civil society organisations, academics, and relevant experts with varying viewpoints and from different countries to inform this policy review. Much of YouTube's work on content policies, which we call the YouTube Community Guidelines, focuses on analysing, assessing, and addressing emerging issues before they reach, or become widespread on, YouTube. Similarly, as risks change and evolve, so do our content policies for Maps (e.g., fake engagement, misrepresentation, and misinformation policies), Play (e.g., user-generated content policies), Search (e.g., highly personal information), and Shopping (e.g., vehicle ads).

Counterbalancing Risk

Fundamental rights are interdependent. The fulfilment of one right (e.g., freedom of expression) may facilitate the fulfilment of other rights (e.g., civic participation and democracy) or come at the expense of others (e.g., freedom from discrimination).

As a result, fundamental rights are sometimes in tension with each other.¹ For example, the pursuit of child safety may limit adult users' rights and present risks to different rights held by children, such as their rights to participation, privacy, and freedom of expression and information. We address these tensions through various means, such as providing parents or guardians with controls that allow them to supervise minors' access to content, and giving users extensive controls over their privacy settings.

When efforts to protect or advance one right may result in the limitation of another right, our approach is to identify and implement sensible mitigation measures to address potential adverse impacts. This balancing involves considering appropriate and proportionate mitigation techniques, such as protecting freedom of expression via appeals mechanisms or raising authoritative content to address lower quality content that may appear on the service, rather than removing low-quality content altogether unless it is unequivocally harmful.

Throughout the VLOSE- and VLOP-specific sections of this report, we explain why one risk may take precedence over another in certain circumstances, describe how the nature and purpose of the service being assessed inform these choices, and set out the reasonable and proportionate mitigations we believe strike the right balance.

¹ This is recognised in the DSA. Recital 153 of the DSA states that “in situations where the relevant fundamental rights conflict, a fair balance between the rights concerned, in accordance with the principle of proportionality” should be achieved.

Detecting and Responding to Harmful Content at Scale

In every country in which we operate, different laws govern what is considered permissible expression. To address these nuances, we have teams and systematic processes to develop and deploy localised policies and enforcement practices. When users [report content they believe violates the law](#) on our services, we carefully review whether to block, limit, or remove access to it.

Handling Government Removal Requests

Courts and government agencies around the world regularly request that we remove user-generated content from our services. We were the first company to publish (in 2010) a formal transparency report about such requests, and you can read more about our process and the volume of requests we receive in the [Government Requests to Remove Content](#) segment of our latest Transparency Report.

As illegal content may also violate product policies (like YouTube's Community Guidelines), it is worth noting most content that is otherwise illegal is first detected by our automated classifiers for violations of our policies and quickly removed. In those cases, though the content may also be illegal, we treat these as violations of our policies, since we may not be in a position to make conclusive determinations about the legality of content.

We maintain a robust process to receive, evaluate, and act on government removal requests. We review these requests closely to confirm that they are supported by local laws and international norms of human rights and to determine whether we should remove content as a matter of law or policy. Consistent with our commitment to the [Global Network Initiative Principles](#), we assess the legitimacy and completeness of government requests, which must be in writing, as specific as possible about the content to be removed, and clear in their explanation of how the content is illegal. We do not honour requests that have not been made through appropriate channels. If we receive an oral request, we ask for it in writing.

In some narrow cases, to protect the rights of users, we do not act on orders that appear illegitimate or inapplicable. For example, we examine the legitimacy of every document we receive, and if we determine that a court order is forged, we won't comply with it. In other cases, we do not need to take action because the content has already been removed by the uploader.

The enforcement of content policy is a joint effort between people and an array of technologies, including automated systems employing machine learning technology, which work together to achieve consistently high levels of accuracy when reviewing content. We design models and train classifiers² to identify potentially violative content, use machine learning to constantly improve those classifiers, take automated actions when we have a high degree of confidence that the content violates our policies, and enqueue content for review by specialist teams when we have lower confidence in fully automated techniques. These human content moderators help confirm whether machine-identified content should be removed, and we use the results of the human review to further train our classifiers and improve their ability to detect evolving violative content.

This collaborative approach helps improve the accuracy of our models over time, as models continuously learn and adapt based on human feedback. And it also means our enforcement systems can manage the scale of content that's available on our services, while still rendering nuanced decisions on whether a piece of content violates our policies. Examples of automated systems and humans working in combination are provided in each of the VLOSE and VLOP sections that follow.

We heavily invest in the training of machine learning classifiers and human content reviewers to increase accuracy. Sometimes we enforce policies broadly to err on the side of user safety, which may result in removal of some content from our services that does not actually violate our policies. In many cases, appeals channels are an appropriate way to fulfil our commitment to freedom of expression and to the UN Guiding Principles on Business and Human Rights by providing a check against incorrect removal and ensuring that content creators have redress.

Whether an appeal is meritorious requires a case-by-case determination, but on our VLOP services, with limited exceptions, we allow users to appeal enforcement actions they believe may have occurred in error. In some other cases, consideration of other equities counsels against providing appeals, such as those involving repeat or abusive violators, ancillary content, or egregious conduct.

We seek to ensure that these mechanisms are accessible and work to learn from appeals outcomes, including to modify our content efforts to help them become more accurate. Insights gained from these appeals processes also inform policy changes to prevent future adverse impacts. Sections of this report specific to each VLOP will describe our appeals mechanisms and where we are expanding them to mitigate risks to freedom of expression.

In addition to our own review and legal removal requests and user flags of illegal content (described above), we offer a variety of mechanisms for users to report and request removal of policy violating content. For example, Google Maps users can [flag content](#) that violates our policies or [profiles of users](#) who are contributing false information, uploading offensive content, or taking other abusive actions. On YouTube, users can [flag videos](#) that may violate our policies. Trained content moderators then review credible flags and take appropriate action, which may result in content being removed, age restricted, geo-restricted, or left up.

Our approach to flagging also involves partnering with other organisations. One example of that is the [YouTube Priority Flagger program](#). This program provides robust tools to government agencies and

² A classifier is an algorithm that identifies and sorts content into types of content. For example, a classifier may proactively identify content with a high likelihood of violating a specific Google policy.

non-governmental organisations (NGOs) to improve our content by notifying YouTube of content that might violate [Community Guidelines](#) (i.e., YouTube's content policies).

Our Public Interest Framework guides policy and enforcement decision-making by safeguarding against content actions that could potentially contribute to or exacerbate adverse impacts due to allowing or removing content following government removal requests, escalations, or other policy enforcement contexts. The framework helps us consider how (1) content, if not removed, could adversely impact the rights of an individual, community, or society as a whole, or (2) whether allowing the content is in the public interest because it furthers the understanding of social, political, cultural, civic, and economic affairs, and so should remain.

Evaluating Content Across Languages

Automated systems such as algorithms and classifiers detect violating content and behaviour at scale. But human operators are often required to review, validate, and train these automated systems because humans can evaluate content or other signals in ways that might be difficult for current automated systems, such as understanding nuance, context, and slang.

Taken together, Google services maintain Trust and Safety coverage, including human content moderators, across nearly all official EU languages, as well as many other languages commonly spoken in the EU.³ However, given the important role played by automated systems, we assessed the risk that algorithms may be less well trained in some languages, dialects, and vernaculars than others. This is an industry-wide challenge not unique to Google, and can be especially important for languages, dialects, and vernaculars that are less commonly spoken and that therefore lack sufficient training data.

One important element of this assessment is the review of how significant advances in machine translation assist with review of content at scale. We have found that English-speaking reviewers relying on machine translation perform nearly as well as native language speakers for the vast majority of official EU languages.⁴ Given the operational challenge of having content moderators available 24/7 for even less widely used languages, the use of these tools enables us to undertake moderation of content at scale more rapidly, consistently, and effectively.

Performance across languages is also important when we use classifiers (such as derogatory or offensive speech classifiers) to identify potentially violative content. We use both human content moderators and machine learning to constantly train these classifiers and improve their ability to accurately detect such content across different languages, dialects, and vernaculars. In addition to using human content moderators native or fluent in many languages, we also consult specialists for cultural or language nuances and use these insights to inform continuous learning and performance improvement.

³ Exceptions are Irish, Maltese, and Slovak.

⁴ Exceptions are Irish, Maltese, and Hungarian.

Our systemic risk assessments found some residual risk remaining for the performance of automated systems across languages, dialects, and vernaculars for all VLOP and VLOSE services reviewed. This is a risk where Google-wide advances (such as continuous improvement in machine translation) can support the work of different Google services as each develops custom models, thresholds, and confidence levels tailored to their own policy enforcement needs.

Going forward we will continue to test the performance of classifiers to identify differences in performance across languages and pursue continuous improvement, including a rolling program to identify priority languages for investment in enhancing translation and content moderation quality.

We will also continue to keep pace with developments in local contexts—including how language and terminology may evolve with potential for higher-risk events, such as upcoming elections—and use human content moderators and native speakers to improve the quality of automated systems, including classifiers.

We are constantly improving our automated systems' ability to operate equally well in many different languages, and recently took action to address the ongoing challenge of gender bias in machine translation.

Our researchers [developed a new dataset for studying and preventing gender bias in machine learning](#) by exploring gender translation between English and Spanish and English and German. Grammatical differences can pose a challenge for machine translation systems, especially when translating from a language without subject pronouns (such as Spanish) to one that requires gendered subject pronouns (such as English). The researchers built a new “context-aware” model that incorporated context from surrounding sentences or passages to improve gender accuracy when translating personal pronouns. This dataset provided useful performance measurements for the new context-aware models; using the dataset, researchers determined that context-aware models made 67% fewer gender translation errors than previous models that translated sentence by sentence.

These improvements enhance our ability to address risks where an accurate understanding of gender identity may be essential to content moderation efforts, such as when assessing content related to risks in the areas of gender-based violence, discrimination, bullying, and harassment.

Two: Delivering Reliable Information

Providing access to high-quality information to all users is core to our mission. We also provide users with best-in-class tools that give additional context that help them evaluate content.

Surfacing Quality Information

The world wide web holds an unprecedented, and growing, volume of information that is not ordered or easily navigable. But when automated systems operating at scale sort, organise, and deliver relevant information, users can find the needles in humanity's largest haystack.

Algorithms power our services by prioritising relevant information in Search results, making app recommendations on Google Play, and providing relevant product listings in Shopping. Our algorithms sort through hundreds of billions of pieces of content to find the most relevant and useful results.

Algorithms enable us to advance quality and relevance while reducing systemic risk to users and society. Search uses signals such as meaning, relevance, quality, usability, and context to help [determine which results are returned](#) and prioritised for each query. Our systems use these and hundreds of other signals to prioritise the results that seem most helpful, in particular content that seems to demonstrate expertise, experience, authoritativeness, and trustworthiness. These signals are especially important for what we call “Your Money or Your Life” (YMYL) topics, defined as those that may significantly impact or affect the health, financial stability, or safety of individual people, or the welfare of society.

To help us [test and improve](#) our Search algorithms we put all possible changes to Search through a rigorous evaluation process to analyse metrics and decide whether to implement a proposed change. We work with external [Search Quality Raters](#) to evaluate if our search systems are generating helpful results that demonstrate experience, expertise, authoritativeness, and trustworthiness.

This overall approach is summarised in [How Search Works](#).

Using Recommender Systems

Some risk factors under Article 34(2), such as recommender systems, may increase or decrease risk. Poorly designed or controlled recommender systems may increase the risk that harmful content goes viral. But properly functioning recommender systems should decrease risk by increasing the visibility of high-quality and trustworthy content and by promoting a diversity of topics and sources for users to explore.

Recommender systems are an essential tool as we navigate the inherent tensions that come with respecting countervailing fundamental rights while fulfilling our mission to organise the world's information and make it universally accessible and useful. We aim to make all of our recommendations useful, inclusive, and empowering.

Using recommender systems to order the presentation of content, including by elevating high-quality and trustworthy content, is often a more proportionate approach to addressing harmful content risk than removing content altogether, which can present risks to freedom of expression and information.

Fighting Misinformation

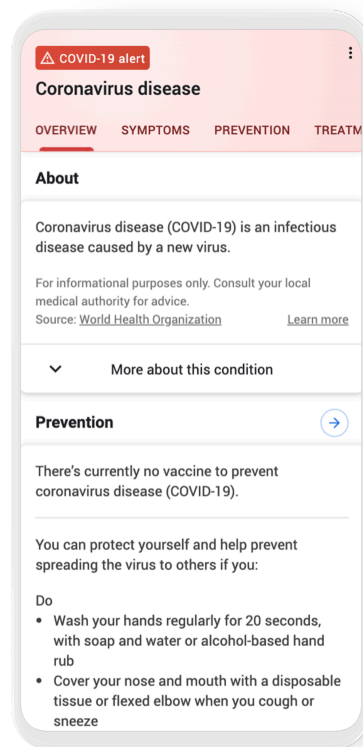
Google heavily invests in elevating authoritative sources and countering misinformation, particularly as it relates to people's finances, health, livelihood, or civic participation and to sensitive events. Misinformation can manifest itself in different ways on different services across the open web, such as misleading pages attempting to monetise their content with our Ads services, health misinformation videos on YouTube, or websites spreading misinformation appearing in Search results. Other examples of misinformation practices include fraud, deceptive behaviour (such as the use of deep fakes), impersonation, misrepresentation of ownership, and medical misinformation. We take action to prevent the spread of this type of content at scale.

For example, during a breaking news cycle, speculation and misinformation can outrun facts while legitimate news outlets are still investigating. Bad actors may publish content with the intent to mislead, or to attract attention and traffic on the basis of unverified information. To defend against these risks, YouTube and Search have automated systems designed to promote authoritative content.

We have long recognised the importance of multi-stakeholder approaches to misinformation, including the EU's 2018 [Code of Practice on Disinformation](#) and a [Strengthened Code](#) that Google signed in June 2022. As part of the Strengthened Code, we have committed to providing the European Commission with [reports](#) detailing how we have implemented our commitments. Our commitment to the Strengthened Code applies to Search, YouTube, and Google Ads, and you can read more about it in the Search and YouTube sections of this report.

Our commitment to fighting misinformation guided our reaction to the COVID-19 crisis. Our products have long-standing policies in place to ban harmful or misleading medical or health-related content, but while the policies are longstanding, they continuously evolve to meet changing global medical consensus. Accordingly, the COVID-19 crisis required new policy and enforcement work to address misinformation and prevent a range of new abuses.

As the pandemic unfolded, there was a surge in online searches for health-related issues, like testing, vaccinations, and masks. The public's interest in health-related information created opportunities for malicious actors, and we faced a wide range of new abuses, including phishing attempts, malware, dangerous conspiracy theories, and fraudulent schemes. Our teams tackled these issues across all our services, enacting policy revisions and stepping up enforcement. We raised the visibility of authoritative content through features such as Health Knowledge Panels and structured search results designed to make trusted information easy to access.



Structured results for searches about COVID-19 on Google Search help make high-quality information easier to access

Ukraine and the broader Central and Eastern European region is facing a disinformation crisis, and our commitment to fighting misinformation has again guided our approach. We're [monitoring the threat landscape in Eastern Europe](#) and disrupting coordinated influence operations from Russian threat actors. By early 2023, we had removed more than 85,000 videos and 9,000 channels on YouTube related to the war for violating our Community Guidelines and Terms of Service, and early in the war we blocked YouTube channels associated with Russian state-funded news channels globally, resulting in more than 800 channels and 4 million videos blocked, including channels tied to RT and Sputnik. Our breaking news and top news shelves on the YouTube homepage have received more than 170 million views in Ukraine, helping people stay connected and informed. And as the largest video-sharing service in Russia, YouTube continues to provide Russian citizens uncensored news and information.

Our report [Fog of War: How the Ukraine Conflict Transformed the Cyber Threat Landscape](#) is based on analysis from our Threat Analysis Group, Mandiant, and Trust and Safety, and provides insights into changes in the cyber threat landscape triggered by the war.

We have also invested in several efforts to tackle the spread of and harms caused by mis- and disinformation. For example, we piloted [Jigsaw's "prebunking" campaign](#) in Poland, Czechia, and Slovakia, to preemptively fight against narratives scapegoating Ukrainian refugees. The initiative proved so effective that we [announced](#) that we will expand it this year to Germany, in partnership with [Moonshot](#) and local experts.

Addressing the Risks and Opportunities of Artificial Intelligence

We have developed artificial intelligence tools to help solve some of society's biggest challenges. AI is embedded in many of our services, such as on Google Maps where we are cutting carbon emissions by [reducing stop-and-go traffic](#). In 2017 we announced our intention to be an “AI-first company”, and we wholeheartedly believe AI has the potential to transform our societies for the good.

AI also presents important challenges that must be addressed clearly, thoughtfully, and affirmatively. In 2018 we set out our [AI Principles](#) and accompanying framework for [responsible AI innovation](#) that describe our commitment to developing technology responsibly and the specific application areas we will not pursue.

The recent momentum behind large-scale machine-learning models (including generative AI) has sparked additional dialogue around the social impacts of AI and surfaced concerns as diverse as misinformation, unfair bias, privacy, security, and safety. One challenge of emerging relevance to systemic risk is the potential for manipulative use of our services by bad actors seeking to use large-scale machine-learning models to scale their spam, scam, or disinformation efforts, or use AI to attack our data and security systems. Another challenge relates to when large-scale machine-learning models do not work as intended or users unintentionally misuse them.

The opportunities and challenges presented by large-scale machine-learning models require global, multi-stakeholder, and collaborative approaches. For this reason we are a founder or active participant in several new initiatives, such as:

- The [Frontier Model Forum](#), a new industry body focused on ensuring safe and responsible development of frontier AI models.
- The Partnership on AI's [Responsible Practices for Synthetic Media: A Framework for Collective Action](#), an initiative to foster best practices in the development, creation, and sharing of media created with generative AI.
- The [White House's Office of Science and Technology Policy](#) initiative to ensure safe, secure, and trustworthy AI.

In this first systemic risk assessment we considered the risks to our services presented by the use of generative AI by bad actors, such as influence campaigns, phishing, and cyberattacks. Over time, it is possible that the development of large-scale machine-learning models will alter the scale and possible severity of some risks, especially those relating to the generation of illegal or harmful content (such as child abuse and exploitation content, terrorist and violent extremist content, and hate speech); misinformation and disinformation relating to elections, civic

discourse, and democratic participation; and digital threats such as account hijackings, phishing attempts, or malware.

We believe our existing mitigations perform well, but we must keep pace with the latest developments in AI technology. Additional mitigations will include integrating [watermarking](#), metadata, and other innovative techniques into our latest generative models and bringing an [About this image](#) tool to Search to give users context about where an image first appeared online. However, large-scale machine-learning models are evolving rapidly, and we expect to assess the impact of these developments across all relevant risks in our future systemic risk assessments.

We have made a [set of voluntary commitments](#)—developed jointly with other leading AI companies and the White House’s Office of Science and Technology Policy—to promote the safe, secure, and transparent development and use of AI Technology. [These commitments](#) will support efforts by the G7, the OECD, and national governments to maximise AI’s benefits and minimise its risks:

Safety:

- 1) Commit to internal and external red-teaming of models or systems in areas including misuse, societal risks, and national security concerns, such as bio, cyber, and other safety areas;
- 2) Work toward information sharing among companies and governments regarding trust and safety risks, dangerous or emergent capabilities, and attempts to circumvent safeguards.

Security:

- 3) Invest in cybersecurity and insider threat safeguards to protect proprietary and unreleased models;
- 4) Incentivise third-party discovery and reporting of issues and vulnerabilities.

Trust:

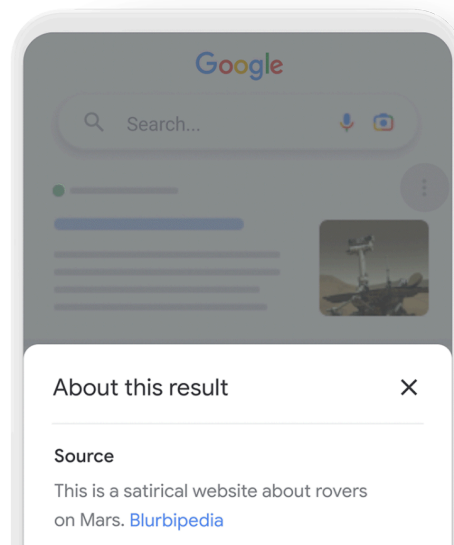
- 5) Develop and deploy mechanisms that enable users to understand if audio or visual content is AI-generated, including robust provenance, watermarking, or both, for AI-generated audio or visual content;
- 6) Publicly report model or system capabilities, limitations, and domains of appropriate and inappropriate use, including discussion of societal risks, such as effects on fairness and bias;
- 7) Prioritise research on societal risks posed by AI systems, including on avoiding harmful bias and discrimination, and protecting privacy;

8) Develop and deploy frontier AI systems to help address society's greatest challenges.

We expect the importance of risks related to AI to increase as the use of these tools expands. Future risk assessments will delve further into the use of AI and our efforts to mitigate systemic risks.

Equipping Users

In addition to raising the visibility of high quality information and fighting misinformation, we aim to equip users with the tools they need to evaluate information they come across on our services.



The 'About this result' feature in Google Search allows users to learn more about the information they are seeing

For example, our “[About this result](#)” feature in Search allows users to learn more about where the information they are seeing is coming from and how our systems determined it would be useful for their query. This feature is available in all languages where Search is available. With this context, users can make more informed decisions about the sites they may want to visit and what results will be most helpful to them. Similarly, [My Ad Center](#) gives users greater control of the ads they see on Google services—like Search and YouTube—by providing options for customising ads, managing privacy settings, and influencing how we determine what ads to show.

[Be Internet Awesome](#) is a Google-created educational program that teaches kids the fundamentals of digital citizenship and safety so they can explore the online world with confidence, such as how to communicate responsibly, discerning between what’s real and what’s fake, and safeguarding valuable information.

On Google Play, the Teacher Approved program identifies apps approved by teachers and children’s education specialists, and then offers a description of the apps’ quality attributes. This information helps families easily review the apps and make informed choices about whether they want their children using an app or game.

Informed users are able to make better use of and see more benefit from our services, a win for both our users and us.

Three: Partnering to Create a Safer Internet

We recognise that the systemic risks associated with VLOSEs and VLOPs are not unique to Google and cannot be addressed by Google alone. We scale our industry-leading practices to help keep users safe online through proactive partnership with experts and organisations to both inform and share our resources and technologies.

Partnering for Information Quality

To effectively combat misinformation, technology companies already collaborate with academics, policymakers, publishers, and NGOs who possess the expertise that helps inform effective methods to address the issue at scale.

For example, with health-specific misinformation, our key partners include the WHO and the Centers for Disease Control and Prevention (CDC). During the height of COVID-19, these partnerships were essential in our efforts to raise information from reliable sources and help people around the world to navigate the pandemic with high quality information.

In 2021, we contributed €25 million to the launch of the European Media and Information Fund to help academics, publishers, and nonprofits launch their own media literacy programs, extend fact-checking initiatives, and conduct vital research into misinformation. In November 2022, we announced a \$13.2 million grant for the International Fact-Checking Network (IFCN) to launch a new Global Fact Check Fund to support their network of 135 fact-checking organisations from 65 countries covering over 80 languages. Building on previous work, this is our single largest grant in fact-checking to date.

During 2023 we [initiated long-term partnerships across Central and Eastern Europe](#), a region considered highly vulnerable to disinformation and propaganda due to its geographic proximity to the war in Ukraine. In the Baltics, we have entered into a long-term partnership with the Civic Resilience Initiative and the Baltic Center for Media Excellence. These two established and well-respected organisations will receive €1.3 million in funding from Google to build on their impactful work towards increasing media literacy, building further resilience and actively tackling disinformation in Lithuania, Latvia, and Estonia. We are partnering with the Charles University in Prague, the main research centre of the Central European Digital Media Observatory (CEDMO) project, and providing €1 million in funding for CEDMO to further expand its research into information disorders (such as misinformation, disinformation, or clickbait), and work to increase the level of media and digital literacy in Poland, Czechia, and Slovakia.

Consulting with Experts

We scale our industry-leading practices to help keep users safe online through proactive partnership with experts and organisations to both inform and share our resources and technologies.

The [Google Safety Engineering Center \(GSEC\) in Dublin](#) is a regional hub for Google experts working to tackle the spread of illegal and harmful content, and a place where we can share this work with policymakers, researchers and regulators. Over the last two years, GSEC Dublin has held approximately 100 public and private engagements to share our experience of managing content risk and hear from experts

across a wide range of topics, including misinformation, ads safety, election integrity, the use of AI in content, and fighting child sexual abuse and exploitation online.

As part of our ongoing support for the people of Ukraine, GSEC Dublin recently conducted several [Fighting Misinformation Online roundtables and summits](#) with local governments, NGOs, and fact-checking organisations across Central and Eastern Europe. We shared Google and YouTube's approach to mis- and disinformation, and learned real-time insights from over 100 local experts and organisations.

In November 2022, Google, the [European University Institute](#), and [Calouste Gulbenkian Foundation](#) convened European policymakers, NGOs, media organisations, academics and tech companies to collaborate and share knowledge about tackling misinformation. Over 900 people in Brussels and online joined the discussion, with talks led by experts from across the misinformation landscape.

YouTube regularly updates its family product experiences and policies in consultation with experts in children's media, child development, digital learning, and citizenship from a range of academic, non-profit, and clinical backgrounds. A key channel for this consultation is YouTube's [Youth and Families Advisory Committee](#), a collection of independent experts that provide advice on the policies and services YouTube offers to young people and families.

YouTube also sponsored the [National Academy of Medicine](#) to convene an independent advisory group to develop principles and attributes to guide digital services companies in identifying and elevating credible sources of health information in their channels. The outcome of this project was a peer-reviewed discussion paper and the use of these principles when [providing content from reliable health sources](#) on Google.

As part of our commitment to image equity and improving representation across our products, we partnered with Harvard professor and sociologist [Dr. Ellis Monk](#) to release a new skin tone scale designed to be more inclusive of the spectrum of skin tones we see in our society. The [Monk Skin Tone \(MST\) Scale](#) is a 10-shade scale designed to be easy to use for development and evaluation of technology while representing a broader range of skin tones. Our research found that people found the Monk Skin Tone Scale to be more representative of their skin tones compared to the current tech industry standard, and this was especially true for people with darker skin tones. The scale will be [incorporated into various Google products](#) (such as in image Search), and we're openly releasing the scale so anyone can use it for research and product development.

Sharing Tools and Technology

We also share tools to help organisations protect platforms and users, including nine safety APIs across child safety, security (such as cyber attacks, malware, and phishing), and information quality (such as misinformation, toxic discourse, and explicit content).

For example, our [Child Safety Toolkit](#) consists of two APIs: the Content Safety API (which classifies previously unseen images of potential child sexual abuse and exploitation) and CSAI Match (which matches known abusive video segments). We offer these APIs to qualifying partners free of charge. Our partners use these technologies to process billions of files, allowing them to evaluate millions of images and videos for abusive behaviour each year.

[Perspective API](#) (which uses machine learning to identify "toxic" comments, making it easier to host better conversations online) and [Harassment Manager](#) (an open source codebase that allows users to document and manage abuse targeted at them on social media) help journalists, activists, politicians, and other public figures document and manage abusive comments on their sites.

Collaborating with Companies and Stakeholders

Many of the risks reviewed in this systemic risk assessment cannot be addressed by a single company acting alone, so we have established and continue to fund and participate in a mix of multi-company and multi-stakeholder efforts that take system-wide approaches to the most intractable problems. This includes sharing signals of illegal and harmful content, collaborating with civil society to gain deeper insights into risk, and sharing best practices across companies.

- **[Global Internet Forum to Counter Terrorism \(GIFCT\)](#)**: In 2017, YouTube co-founded GIFCT with a group of companies dedicated to disrupting terrorist abuse of members' digital platforms. GIFCT provides a formal structure to accelerate and strengthen our work and present a united front against the online dissemination of terrorist content, such as by identifying and sharing signals of terrorist and violent extremist activity via the GIFCT hash sharing database.
- **[Tech Coalition \(TC\)](#)**: In 2006, we joined the Technology Coalition, teaming up with other tech industry companies to develop technical solutions that disrupt the ability to use the Internet to exploit children or distribute child sexual abuse material (CSAM). For example, we have been one of two members to test a system to increase the chances of detecting CSAM videos through [hash matching](#), while our child safety experts also chair or actively participate in half a dozen key working groups of the Tech Coalition.
- **[Global Network Initiative \(GNI\)](#)**: We were a founding member of the GNI in 2008, and since then we have worked closely with civil society, academics, investors, and industry peers to protect and advance freedom of expression and privacy globally, especially when faced with demands from governments that conflict with international human rights standards.
- **[Web Foundation](#)**: We are a funder and partner to the World Wide Web Foundation. Founded in 2009 to advance the open web as a public good and a basic right, the Web Foundation is an independent, international organisation fighting for a world where everyone has affordable, meaningful access to a web that improves their lives and where their rights are protected. We were an active participant in the "[Tackling Online Gender-Based Violence and Abuse](#)" workstream, which brought together tech companies and women from across civil society to gather evidence of online abuse and create policy and product solutions to address online gender-based violence.

Developing Best Practices

We actively participate in efforts to develop best practices that advance responsible and effective approaches to risk assessment across the industry, as well as to develop the field of trust and safety more broadly.

- **[Digital Trust and Safety Partnership \(DTSP\)](#)**: We co-founded the DTSP alongside nine other companies in 2021. The DTSP is committed to developing industry best practices, verified through internal and independent third-party assessments, to ensure consumer trust and safety when using digital services.
- **[Partnership on AI \(PAI\)](#)**: In 2016 we were a co-founder of the Partnership on AI, a non-profit partnership of academic, civil society, industry, and media organisations helping AI advance positive outcomes for people and society. We are also a member of PAI's [Responsible Practices for Synthetic Media: A Framework for Collective Action](#), which is fostering expertise and best practices for responsible practices in the development, creation, and sharing of media created with generative AI.
- **[World Economic Forum Global Coalition for Digital Safety](#)**: We have been an active participant in this effort to accelerate public-private cooperation to tackle harmful content online and exchange best practices. For example, we contributed to the [digital safety risk assessment framework and bank of case studies](#).
- **[Trust and Safety Professional Association](#)**: We are a [founding supporter](#) of the Trust and Safety Professional Association, a non-partisan membership association that supports the global community of professionals who develop and enforce principles, policies, and practices that define acceptable behaviour and content online and/or facilitated by digital technologies.

Setting High Standards for Advertising

We strive to create a healthy, trustworthy, and transparent digital advertising ecosystem that supports users and advertisers. Our advertising policies apply across all Google services and are designed to ensure a safe and positive experience for our users, in part by prohibiting content that is harmful to users and the overall advertising ecosystem.

Our advertising policies and review process help address risks across our VLOSE and our VLOPs and cover 4 broad areas:

- **Prohibited content:** Content that cannot be advertised on the Google Network⁵, such as counterfeit goods, dangerous products or services, and inappropriate content.
- **Prohibited practices:** Practices that advertisers may not engage in, such as misrepresenting the company purchasing ads or the products or services they are offering, or personalised advertising in certain disallowed contexts. For example, our [personalised advertising policy](#) defines certain interests as not eligible for personalised advertising, such as offerings related to personal hardships, identity and belief, and sexual interests. We also prohibit ads that potentially profit from or exploit a sensitive event that may create risk to our ability to provide high quality and relevant information, such as a natural disaster, public health emergency, terrorism and related activities, conflict, or mass acts of violence.
- **Restricted content and features:** Content that can be advertised, but with limitations, such as sexual content, alcohol, gambling and games, healthcare and medicines, financial services, and political content. These limitations stop an ad from showing when it is inappropriate for that context.
- **Editorial and technical:** Quality standards for ads, websites, and apps, such as high editorial standards, destination requirements, and ad format requirements.

We set a high standard of quality and reliability for advertisers. We have processes in place to identify bad ads before they are published on our services and to monitor violations on an ongoing basis. We do not want to make revenue from harmful content or behaviours.

Advertisers may not run personalised ads on content designated as “made for kids”, and we maintain a separate [Ads & made for kids content policy](#) that includes content on topics such as restricted ad categories and prohibited ad content. Advertising that is intended for children or on content designated as “made for kids” must not be potentially harmful to children, must not make use of any third party trackers or otherwise attempt to collect personal information without first obtaining parental consent, and must otherwise comply with all applicable laws and regulations.

⁵ The Google Network refers to all the places where ads can appear, including Google sites, websites that partner with us, and other settings like mobile apps.

To keep ads safe and appropriate for everyone, we have developed and maintain automated classifiers to review ads before publishing to make sure they comply with Google Ads policies. This review covers the content of ads, including the headline, description, keywords, and destination. Ads that do not follow Google Ads policies are disapproved and are not able to run until the policy violation is fixed and the ad is reviewed again. Accounts may be suspended if we detect an egregious violation. [In 2022](#), we removed over 5.2 billion ads, restricted⁶ over 4.3 billion ads, and suspended over 6.7 million advertiser accounts worldwide. This represents an increase of 2 billion ads removed in 2022 compared to the previous year.

For repeat violations of an Ads policy, we [issue strikes](#) to the Google Ads account and penalties progressively increase with each subsequent strike leading up to account suspension. To address the risk of over-enforcement, advertisers can appeal potentially erroneous [ad reviews](#), [strikes](#), and [suspensions](#).

In addition, in the EU, our [Google Ads Transparency Center](#) is a searchable hub of all served ads and is designed to give users more information about the ads they see on Google services. In the Ads Transparency Center, users can see the ads an advertiser has run, find out which ads were shown in a certain region, and learn more about the advertiser.

We are also committed to delivering ads responsibly in ways that respect user privacy, which we seek to achieve by applying the following six privacy principles to our ads business:

1. We never sell your personal information to anyone. This includes for ads purposes.
2. We are transparent about what data we collect and why. We clearly label ads and sponsored content on our services and make it easy for you to understand why specific ads are shown, what information is used, and how you can control your Google ad experience.
3. We make it easy for you to control your personal information. [My Ad Center](#) allows users to customise their ad experiences on Google services. Ads personalization can be turned off altogether, and activity data tied to an account can be permanently deleted at any time.
4. We reduce the data we use to further protect your privacy. We never use sensitive information like health, race, religion, or sexual orientation to tailor ads, and never use the content users create and store in apps like Drive, Gmail, and Photos for ads purposes. We do not allow ads personalization for users where we know that they are under 18.
5. We protect you by building products that are secure by default. We verify advertisers globally and work to detect bad actors and limit their attempts to misrepresent themselves.

⁶ Restricted ads are legally or culturally sensitive and can only run in limited contexts.

- 6.** We build advanced privacy technologies and share them with others. Teams across Google are collaborating with the wider industry to implement the [Privacy Sandbox](#) initiative that aims to make current tracking mechanisms obsolete, and block covert tracking techniques, such as fingerprinting.

To learn more about our commitment to maintaining a responsible advertising service, see our [Ads Safety Report](#).

3. Methodology

Introduction

Article 34 of the DSA requires providers of VLOSEs and VLOPs to identify, analyse, and assess systemic risks in the EU stemming from the design or functioning of their services and their related systems or from the use made of their services. The DSA requires that these systemic risk assessments are undertaken annually and prior to deploying functionalities that are likely to have a critical impact on systemic risks.

We have undertaken a separate systemic risk assessment for each Google service designated as a VLOSE (Search) or VLOP (Google Maps, Google Play, Shopping, and YouTube).

The DSA enumerates four categories of systemic risks to be addressed:

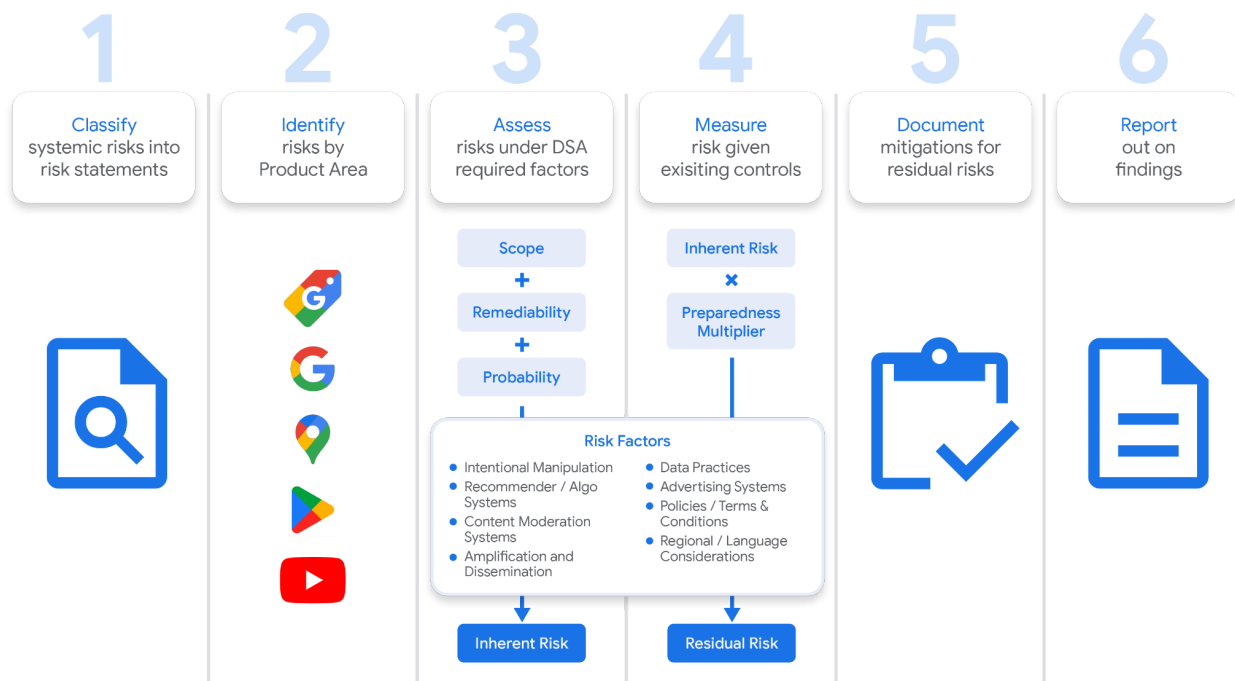
- A.** The dissemination of illegal content.
- B.** Any actual or foreseeable negative effects for the exercise of fundamental rights enshrined in the Charter of Fundamental Rights of the European Union (EU Charter), in particular human dignity, privacy, data protection, freedom of expression and information, non-discrimination, rights of the child, and consumer protection.
- C.** Any actual or foreseeable negative effects on civic discourse and electoral processes, and public security.
- D.** Any actual or foreseeable negative effects in relation to gender-based violence, the protection of public health and minors, and serious negative consequences to the person's physical and mental wellbeing.

We developed our systemic risk assessment methodology by combining the specific systemic risk assessment requirements of the DSA with proven risk assessment methodologies, such as those used to assess enterprise risk, human rights risk, compliance risk, and systemic risk assessments in other sectors. We integrated practices from well-established systemic risk assessments, such as the World Economic Forum's Global Risk Report, as well as our own existing best practices.

To help ensure that our methodology was sound and executed well, we retained the services of two consultancies with expertise in risk assessments of different kinds. Each reviewed and contributed to the development of our risk assessment, bringing points of view from their respective fields. Teams from Business for Social Responsibility (BSR), with extensive experience in the field of human rights assessments, and KPMG, with expertise in systemic risk assessments in the financial, energy, and pharmaceutical sectors, contributed to the development of our systemic risk assessment methodology as well as the execution of the assessment.

Importantly, we designed our systemic risk assessment methodology to identify and prioritise risk to people and society, rather than risk to business objectives.

The methodology we developed has six steps:



Our systemic risk assessment methodology

Step One: Classifying Risk

In this step we established a list of 40 “risk statements” across the four categories of systemic risk. The risk statements are plain-language articulations of the potential adverse impacts for each risk category and provide the core focus for each VLOSE and VLOP systemic risk assessment.

By using specific risk statements we were able to assess related risks that may require different mitigations, or competing risks that need to be balanced against each other. Risk statements are commonly used in risk assessments to clarify the scope of the risk assessment and focus risk assessor responses on specific exposures.

We relied on insights from a range of internal and external sources to generate these risk statements, including human rights due diligence, outputs from external stakeholder engagement (e.g., surveys, dialogue, roundtables), literature review, and discussions with relevant teams, staff, and subject matter experts at Google.

We assured completeness by reviewing the risk statements against all articles in the EU Charter and cross-referencing against rights contained in international human rights instruments.

We created thirty-seven risk statements to apply to each VLOSE and VLOP, and added a small number of VLOP-specific risk statements where unique service features warranted it. Specifically, because Google Play is a service that hosts other apps, we created three additional risk statements to determine whether those app offerings were adequately diverse to serve many demographic groups.

The complete list of the 40 risk statements can be found in Annex A.

Step Two: Identifying Risk

In this step we identified the risk drivers and exposure scenarios that may lead to inherent risk for each risk statement relevant to each VLOSE and VLOP, and pinpointed the quantitative and qualitative insights needed to assess systemic risk. This included establishing clarity on the purpose, function, and features of each VLOSE and VLOP, including the volume and type of content and the service's potential contribution to the virality of content.

The systemic risk assessment included a mix of quantitative factors that were more straightforward to assess (e.g., violative view rates or successful appeals) and qualitative factors that required professional judgement (e.g., the remediability of a privacy violation or the impact of hate speech).

Engaging Stakeholders

Recital 90 of the DSA sets out the expectation that providers of VLOSEs and VLOPs engage with external stakeholders (such as representatives of the recipients of the service, representatives of groups potentially impacted by their services, independent experts, and civil society organisations) when undertaking risk assessments and designing mitigation measures.

We have long engaged external stakeholders to provide expertise related to emerging and evolving issues that intersect with our services. This input is important to the business, which helps inform our decision-making, our due diligence efforts regarding human rights obligations, and our design of mitigation measures. Consistent with Recital 90, these efforts include:

- **Utilising prior relevant engagements:** We systematically catalogued and synthesised insights relevant for each risk statement from sources such as:
 - **Google’s Government Affairs and Public Policy** engagements with independent experts and civil society organisations for due diligence, decision-making, and strategy.
 - **Google’s Human Rights Program** stakeholder engagement on human rights issues with civil society and human rights organisations, government officials, and others.
 - Engagements through the **Google Safety Engineering Center** led by **Google’s Trust and Safety** team with policymakers, researchers, and regulators with an interest in Google’s content policy and its enforcement.
 - Google’s participation in relevant multi-stakeholder and multi-company efforts, such as the **Global Network Initiative**, the **Global Internet Forum to Counter Terrorism**, **Partnership on AI**, the **Tech Coalition**, the **Family Online Safety Institute**, and the **Digital Trust and Safety Partnership**.
 - Engagements that **content policy teams** within VLOSEs and VLOPs have with civil society organisations, academics, and relevant third party experts to inform the review, development, and enforcement of content policy and get ahead of emerging issues.
 - YouTube’s **Youth and Families Advisory Committee**, a collection of independent experts that provide advice on the policies and services YouTube offers to young people and families.

- **User engagements** facilitated by marketing functions and specific product teams to test service features or understand user sentiments about Google and its services.
- **New DSA-specific engagements:** We participated in [two multi-stakeholder convenings](#) hosted by the Global Network Initiative and Digital Trust and Safety Partnership to discuss both methodological questions, such as the definition of systemic risk and key features of risk assessment methodology, and substantive issues, such as fundamental rights, illegal content, civic discourse, and gender-based violence.

Some of the insights gleaned from external stakeholder engagements like these helped determine whether the risk statements developed for the assessment appropriately addressed the categories of systemic risk identified in the DSA and informed our assessment of scope, remediability, probability, and preparedness.

The perspectives of external stakeholders formed an essential part of the overall mix of information that was used to assess systemic risk and design mitigations. The insights gained tended to be qualitative rather than quantitative in nature and were especially useful for assessing remediability and preparedness.

As evidenced in this assessment, we are already deeply engaged with experts, stakeholders, and civil society. In future years, we expect that these engagements will naturally align with the DSA systemic risk areas to better inform our annual assessment.

Step Three: Assessing Inherent Risks

In this step we assessed each risk statement according to the potential severity of the negative effects that the risk could cause and the probability or frequency of the risk's occurrence. Combined, these elements produce an estimate of the **inherent risk**—the risk absent risk reduction efforts by Google. This step is necessarily theoretical, imprecise, and abstract because we have long been dedicated to mitigating all the systemic risks identified in the DSA. That estimate was then used in the later steps as the foundation to review how well we address each risk. This enabled us to understand the **inherent** systemic risks that could stem from the design, functioning, and use of VLOSE and VLOP services, as well as from potential misuses by others.

This step comprised two important elements.

First, we considered whether the following factors set out in the DSA would impact the risk positively, negatively, or both:⁷

- A. Design of recommender systems and any other relevant algorithmic system
- B. Content moderation systems
- C. Applicable terms and conditions and their enforcement (e.g., content policy)
- D. Systems for selecting and presenting advertisements
- E. Data-related practices of the provider
- F. Intentional manipulation of the service, including by inauthentic use or automated exploitation of the service
- G. Amplification and potentially rapid and wide dissemination of illegal and policy-violating content

⁷ See Article 34(2) of the DSA.

We also considered whether linguistic or regional differences could affect the risk or any of the above factors. Second, we used the quantitative and qualitative metrics and insights pinpointed in Step Two (Identification) to assess each risk statement according to the following objective criteria:⁸

- **Severity**, meaning the potential consequences of the risk for people and societies, as defined by two criteria:
 - **Scope**: The number of users and/or persons who could be primarily affected by the risk—for example, we considered whether the risk would impact all users of the service or only a subset, and whether the risk would impact non-users as well as users of the service.
 - **Remediability**: The potential to reverse the impact of the risk were it to occur—for example, we considered the gravity of adverse impacts on physical, mental, or financial wellbeing, and whether a post-hoc remedy could restore those affected to their condition prior to the impact.
- **Probability**: The likelihood and frequency of the risk—for example, we considered the prevalence and potential virality of violative content, the volume of cases or data involved, or the number of successful appeals.

In line with human rights guidance and risk assessment best practices, we used a weighting system so that severity rather than probability would be the predominant factor, meaning that “high severity/low probability” risks received higher prioritisation than “low severity/high probability” risks.

These inherent risks do not actually manifest in our services because each of our services take steps (as described below) to mitigate these inherent risks.

Step Four: Assessing Preparedness

In this step, we reviewed the mitigations (e.g., policies, controls, enforcement practices and other measures) we have in place to address each risk and assessed the level of our preparedness, resulting in an estimate of residual risk (i.e., the risk after mitigation efforts by Google) for each risk statement.

To achieve this estimate, we identified controls and other measures that contribute to our preparedness including (1) the existence and coverage of design decisions, features, policies, processes, metrics, accountability, and formal controls, and (2) other relevant measures, such as participation in industry and multi-stakeholder efforts to address risks. Ultimately, we considered the extent to which the combination of mitigations prevents or significantly addresses adverse impacts of the risk.

Many of the factors the DSA directs to be considered, and which we considered in assessing inherent risk (such as recommender systems, content systems, terms and conditions, and systems for selecting and presenting advertisements) are also important measures for addressing risk, so were also considered in our determination of preparedness.

⁸ Recital 79 of the DSA states: “In determining the significance of potential negative effects and impacts, providers should consider the severity of the potential impact and the probability of all such systemic risks. For example, they could assess whether the potential negative impact can affect a large number of persons, its potential irreversibility, or how difficult it is to remedy and restore the situation prevailing prior to the potential impact.”

A discussion of the most important residual risks for each VLOSE and VLOP is found in the results section below.

Taking a Human Rights-Based Approach

We have long been committed to respecting the rights enshrined in the Universal Declaration of Human Rights and its implementing treaties, and to undertaking human rights due diligence (including human rights assessment) using methods based on the United Nations Guiding Principles on Business and Human Rights (UNGPs).

The systemic risk assessment requirement of the DSA shares a common goal with ongoing human rights due diligence processes undertaken to fulfil our commitment to upholding the UNGPs.

For example, the DSA requirement that a systemic risk assessment consider actual or foreseeable negative effects for the exercise of fundamental rights enshrined in the EU Charter is very similar to the UNGPs expectation that companies assess any actual or potential adverse human rights impacts using internationally recognised human rights as a reference point. Other elements of systemic risk assessment (such as impacts on civic discourse, electoral processes, public security, gender-based violence, public health, and physical and mental wellbeing) are also clearly relevant to ongoing human rights due diligence.

While we designed our systemic risk assessment methodology to meet the requirements of the DSA, we were able to build upon our prior experience undertaking ongoing human rights due diligence based on the UNGPs. This included the generation of risk statements, which was informed by our prior ongoing human rights due diligence, and the creation of assessment criteria, which were based on the notions of severity and likelihood used during human rights due diligence.

Step Five: Identifying Additional Mitigations

In this step, we used the results of the risk assessment to identify where additional mitigations are needed. We identified these additional measures to ensure that there were reasonable, proportionate, and effective mitigations in place to address the specific systemic risks we identified, consistent with Article 35 of the DSA. To ensure successful execution, these enhancements were designed and are being tracked and monitored through established, formal business processes.

This step concluded the systemic risk assessment and mitigations process. The results have been calibrated across our VLOPs and VLOSE to ensure the consistent application of the methodology, and they were approved by central Google stakeholders as well as risk owners and leadership teams from each VLOP and VLOSE. The Independent Compliance Function ensured that the risks were properly identified and reported, and that identified risk mitigations were reasonable, proportionate, and effective.

A discussion of the mitigation enhancements for each VLOSE and VLOP is found in the results section below.

Step Six: Reporting the Results

We disclose the results of the systemic risk assessment in this report. We will publish these reports (subject to removal of confidential information) in due course, consistent with the requirements of Articles 35 and 42 of the DSA.

Some information in this report is confidential or security-sensitive. This includes specific discussion of vulnerabilities, or details of security and programs that could be abused by bad actors. We reserve the right to remove this information from the publicly available version of this report, as contemplated by Article 42(5) of the DSA.

4. Results of the Assessment

Dedicated sections below contain the systemic risk assessment results for each VLOSE (Search) and VLOP (Maps, Play, Shopping, and YouTube). Those VLOP and VLOSE sections each describe:

- The service and its associated systemic risk profile based on its use
- A summary of assessment results, emphasising elevated inherent and residual risk
- The existing mitigations such as content and service design choices that address systemic risk, and new mitigations to address residual risk. Taken in combination, existing and enhanced mitigations are intended to be reasonable, proportionate, and effective for the risk being addressed.

Three important observations emerged across the five systemic risk assessments.

First, the purpose of a service is a primary factor in determining the greatest inherent risks. For example, services prioritising maximum access to information (such as Search) had lower risks to freedom of expression and higher risks associated with potentially harmful content; services oriented toward a narrower purpose (such as Maps) had higher risks to freedom of expression and lower risks associated with potentially harmful content. Product and content policies are tailored to allow or disallow certain types of content and conduct based on this purpose.

Second, the highest evaluations of preparedness (i.e., our existing mitigations) generally correlated with the most significant inherent risks, confirming that we are appropriately allocating resources to the most significant risks.

Third, and despite our existing measures, risk from highly motivated bad actors continues to be of concern in connection with misinformation related to civics, public health, and fraudulent business, as well as external digital threats such as fraud, malware, scams, and malicious sharing of private information. Notable shared characteristics of these areas include the ever-evolving nature of the risks, the determined nature of highly motivated bad actors, and the importance of industry-wide and multi-stakeholder efforts to address the challenges. We also concluded that there are several areas where we can enhance our mitigations, such as new or enhanced user reporting and appeals channels, improved translation and content moderation across languages, and more robust efforts to address disinformation and misinformation.

Google Search

Search

Description of Service and Associated Risk Profile

Google’s mission to organise the world’s information and make it universally accessible and useful starts with Google Search. Search continuously maps the web and connects users to the most relevant and helpful search results for their queries. You can read more in our description of [How Search Works](#).

Search plays an essential role in supporting the enjoyment, realisation, and fulfilment of the right to freedom of opinion and expression. Over 360 million users in the EU⁹ exercise their right to seek and receive information¹⁰ through Search, and web publishers are better able to express themselves and reach interested audiences through Search results.

However, where there is clear user intent to find certain content, returning responsive results that some may find objectionable, offensive, or problematic is not just tolerable, but the right outcome, ensuring users’ access to information they seek. When a user wants to know where on the web a particular piece of content can be found, the user should be able to construct a query that seeks it out, and Search returns responsive information with links to relevant sources, subject to any legal obligations and transparently and clearly defined policies. Content appearing in response to sufficiently clear queries indicates that Search is working as intended. Failure to deliver this content would harm both the rights of the speaker to freedom of expression and the rights of the user to seek and receive information.

This approach is consistent with our understanding of the DSA, which acknowledges the important distinction between search engines and hosting services¹¹ and states that VLOPs and VLOSEs should pay particular consideration to the impact on freedom of expression when mitigating content risks and avoid unnecessary restrictions on the use of their service.¹²

⁹ Average monthly counts based on distinct signed-in accounts of recipients.

¹⁰ Article 19 of the Universal Declaration of Human Rights: “Everyone has the right to freedom of opinion and expression; this right includes freedom to hold opinions without interference and to seek, receive and impart information and ideas through any media and regardless of frontiers”; Article 11 of the EU Charter: “Everyone has the right to freedom of expression. This right shall include freedom to hold opinions and to receive and impart information and ideas without interference by public authority and regardless of frontiers.”

¹¹ For example, Recital 19 recognises “the different nature of the activities” between caching services and hosting services.

¹² Recital 86 states that mitigations should “avoid unnecessary restrictions on the use of their service, taking due account of potential negative effects on those fundamental rights...providers should give particular consideration to the impact on freedom of expression.”

To protect the right to freedom of expression, it is therefore essential that any restrictions we implement be reasonable and proportionate. We [first outlined our approach to freedom of expression in 2007](#), and while we have [refreshed](#) and refined our principles, our philosophy on this issue has remained [largely consistent](#) since then.

Content policies for Search are designed to minimise restrictions on freedom of expression and promote access to information. This design means that risks associated with potentially illegal or “legal but harmful” content will always be present with Search because content may still be discoverable if it is available on the internet. When returning Search results, we take action to avoid surfacing egregious content, such as CSAM, highly personal information, or known non-consensual explicit imagery. We take protective measures to avoid showing shocking or harmful results when a user is not deliberately looking for such content, and provide tools such as SafeSearch to limit results. This includes turning SafeSearch on by default for known minors, and applying the SafeSearch explicit-image blurring by default for all new users. These measures help address risks relating to content that may be objectionable, offensive, or problematic, especially to those who are not seeking it out.

Our approach is informed by several important factors.

First, Search results should not unexpectedly present content that may be objectionable, offensive, or problematic. We deploy a range of measures such as ranking algorithms, quality testing, and content policies (described below) to ensure that results are relevant, helpful, and trustworthy. We acknowledge the risks of problematic content and provide users with relevant contextual information where appropriate.

Second, we are cognizant of the unique concerns around protecting children who use Search, and have implemented features like SafeSearch, described further below, to address those concerns.

Third, our approach distinguishes between core web results (such as links to external pages) and certain other features of Search (such as Autocomplete, Featured Snippets, and Discover). To keep information accessible, we remove content from the core web results that are relevant to a query only in limited circumstances: this includes blocking CSAM, spam, highly personal information (upon request), and content that is subject to valid legal complaints or site owner requests. By contrast, Search features offer additional value—such as providing extra context, helping users formulate queries, or creating a personalised feed—and we understand that users may perceive this content to have higher credibility because of how it is presented. Here we apply content policies that cover a wider variety of issues, including barring harassing, hateful, and violent content. We carefully consider what appears in Search features because our presentation can emphasize and highlight the content in a manner beyond the simpler ordered list we use to display core web results.

Lastly, and critical to understanding the nature of systemic risk on Search, search engines do not have the same relationship with users and user-generated content as hosting providers, including online services. Search engines cannot remove content from the web; only the website owner (and the hosting provider that stores website owners’ content on its servers) can remove the content, and only the website owner can have a direct relationship with the actual user who posted the offending material on its site.¹³

¹³ Recital 50 emphasizes the important role that providers of hosting services play in tackling illegal content online.

Search engines help users access the marketplace of ideas. Placing broad restrictions on the types of content that can be accessed through search engines would interfere with the right to freedom of opinion and expression, including the right to seek, receive and impart information, and the ability to access and hear different views. That's why we remove content from search results only in very limited circumstances, including legal removals, violations of our web search spam policies, or violations of our narrowly-scoped policies that address highly personal information that is rarely if ever in the public interest to display.

Systemic Risk Assessment Results and Associated Observations

In this systemic risk assessment, we considered risks associated with Search and features that appear on Search.

We assessed 37 different risk statements¹⁴ for inherent risk (i.e., risk absent any action taken by Google), preparedness (i.e., the cumulative measures currently in place to mitigate the risk), and residual risk (i.e., risk after mitigation by Google). Residual risk serves as a guide for where further investment may be warranted. The full list of risk statements is found in Annex A to this report.

One important theme for Search is the ongoing presence of residual risk relating to the availability of potentially illegal or harmful content in Search results. This arises from our chosen emphasis on maximising access to information and awareness of the risk of over-broad restrictions to freedom of expression and information; we believe this result to be appropriate given the nature and purpose of Search.

In the next three sections (“Removing Content”; “Investing in Search Information Quality”; “Service Design”) we address each of the categories of systemic risks articulated in Article 34(1) of the DSA. We emphasize where the assessment showed elevated inherent or residual risk for particular risk statements and describe what Search is doing and plans to do about the systemic risk.

As explained above, fundamental rights are closely interconnected and there is a high degree of dependency between different risk statements. The improvement or deprivation of one fundamental right can advance or adversely affect the fulfilment of other fundamental rights, while the controls and measures to address one risk statement may address other risk statements too. With this in mind, we have grouped risks and mitigations together based on how these risks manifest for Search and how they are addressed. This allows for efficient explanation of Search’s existing mitigating practices, as well as improvements consistent with Article 35 of the DSA.

¹⁴ See Methodology Step One: Classification.

Removing Content

Removing Illegal Content

We remove pages from Search results when we have a legal obligation to do so. In many cases, content that is manifestly illegal also violates our policies, so we remove it before we ever receive a legal order to do so. For example, CSAM is generally illegal regardless of the context in which it appears, and so automated methods such as hash matching to detect and remove violating content from search results are effective and enable us to move higher inherent risks into lower residual risks. This is described in more detail below.

Other types of potentially illegal content (such as terrorist and violent extremist content, hate speech, or non-consensual explicit images) either have no standard definition or require contextual understanding, such as whether the subject of the content consents to its availability online or whether the content has educational value, appears as part of a documentary, or represents artistic expression. Deciding whether content is illegal is not always a determination that Google is equipped to make, and we balance taking action against content with respect for the right to freedom of expression and information. Based on the mitigation measures described below, we assess risks relating to potentially illegal content to have lower levels of residual risk. However, risks relating to potentially illegal online activity (such as sharing of non-consensual explicit images) are more complex to address and require notification to Google by affected persons to determine their illegality, and we have assessed these to have more elevated levels of residual risk.

We hold ourselves to a high standard when it comes to our legal requirements to remove content from Search results. We encourage people and authorities to alert us to content they believe violates the law, and we make every effort to appropriately respond to legal notices.¹⁵

Addressing Violations of Intellectual Property Rights

Search responds to clear and specific notices of alleged copyright infringement and delists content and URLs that violate applicable copyright law from search results.

However, a search engine cannot automatically confirm whether a given page hosting content has a licence to do so; we depend on reports from copyright owners. To initiate the process to delist content from Search results, a copyright owner who believes that a URL points to infringing content sends us a takedown notice for allegedly infringing material. When we receive a takedown notice, our teams and automated systems carefully review it for completeness and validity. If the notice is complete and we find no other issues, we delist the URL from Search results.

Because of established frameworks for understanding and mitigating risks associated with intellectual property violations (including copyright), as well as processes for legal review and removal after claims of infringement, our assessment did not find elevated residual risk with respect to these concerns. You can read more in our [Copyright Help Center](#) and the [Content Delistings Due to Copyright](#) section of the Google Transparency Report. The latter provides data for the number of URLs requested to be delisted, the number of unique individuals or entities that have claimed an exclusive right to content specified in copyright

¹⁵ See *supra* at p. 20 (“Handling Government Removal Requests”).

delisting requests, and the reporting organisations, specified domains, and copyright owners who have submitted or been cited in the most requests.

Detecting, Removing, and Reporting CSAM

The mere presence of CSAM (Child Sexual Abuse Material) on a page is illegal in most jurisdictions regardless of context and causes clear harm to victims, so we develop ways to automatically identify that content and prevent it from showing in our results.

We invest heavily in fighting CSAM and exploitation online, and use our proprietary technology to deter, detect, remove, and report offences on all our services. We identify and report CSAM with trained specialist teams and cutting-edge technology, including machine learning classifiers and hash matching technology, which creates a “hash”, or unique digital fingerprint, for an image or a video so it can be compared with hashes of known CSAM. When we find content that appears to be CSAM, we report it to the National Center for Missing and Exploited Children (NCMEC), which liaises with law enforcement agencies around the world.

For many years, we have been working on automated systems to allow us to proactively identify never-before-seen CSAM imagery so it can be reviewed and, if confirmed as CSAM, removed and reported as quickly as possible. In addition to consistently applying it to eliminate CSAM from Search, this technology also powers the [Content Safety API](#), which we developed to help partner organisations classify and prioritise potential abuse content for review. The Content Safety API is one part of our [child safety toolkit](#)—alongside CSAI (Child Sexual Abuse Imagery) Match, YouTube’s proprietary technology for combating CSAI videos online. Every month, our partners use the toolkit to process over 4 billion images and videos, helping them identify problematic content faster and with more precision so they can report it to the authorities. When we help our online partners identify more abusive content, the entire internet benefits.

It’s our policy to block search results that lead to child sexual abuse imagery or material that appears to sexually victimise, endanger or otherwise exploit children. We are constantly updating our algorithms to combat these evolving threats.

We apply extra protections to searches that our systems identify as likely seeking CSAM content. We filter out explicit sexual results if the search query seems to be seeking CSAM. For queries seeking adult explicit content, Search won’t return imagery that appears to include children, to break the association between children and sexual content. In many countries, users who enter queries clearly related to CSAM are shown a prominent warning that child sexual abuse imagery is illegal, with information on how to report this content to trusted organisations like the Point de Contact in France and FSM in Germany. When these warnings are shown, users are less likely to continue looking for this material. Our evaluations of the effect of the prominent warning show a 20-27% increase in queries with no interactions and a 15% reduction in CSAM-seeking follow-on queries.

During 2022, we reported and delisted 921,593 URLs for CSAM from Search by using both automated and manual removals. This is in addition to our efforts to deter CSAM-seeking queries, noted above. You can read more about the scale of our efforts to combat online [CSAM in our transparency report](#).

Removing NCEI and ISPI

Globally, we have policies to remove both [non-consensual explicit images](#) (NCEI) and [involuntary synthetic pornographic images](#) (ISPI) upon receiving a request that meets certain requirements. Online sharing of this type of material can be extremely distressing to the subjects. In some contexts and jurisdictions, this content is not only offensive and harmful, but also illegal to post and distribute.

For people who wish to remove NCEI and ISPI depicting them from Search, we [provide a process to request](#) removal of links to the content from Search results pursuant to our policies against this type of content. People can also submit a separate [legal removal request](#) if they believe the content violates particular laws, such as copyright laws or local laws prohibiting the non-consensual sharing of explicit images.

Once NCEI has been reported, if it meets removal requirements and Search has removed the content from results, then we also begin to block duplicates and filter explicit results on queries that return results similar to the previously reported NCEI content. Further, if we process a high volume of personal information removals involving a site with exploitative removal practices (i.e., sites that require payment to remove content), we demote other content from the site in our results. We also look to see if the same pattern of behaviour is happening with other sites and, if so, apply demotions to content on those sites. We may apply similar demotion practices for sites that receive a high volume of doxxing content removals. We also maintain automatic protections designed to prevent non-consensual explicit personal images from ranking highly in response to queries involving people's names.

While Search has a robust set of policies and tactics to mitigate the risk of this content appearing in Search results, especially for users who are not looking for explicit content, the volume and virality of NCEI results in a higher inherent risk. We are well prepared to address this risk, but highly motivated bad actors and the difficulty of proactively detecting NCEI and ISPI means that there is always room to improve our mitigations. Recognizing the need to continually improve our protections for users, and pursuant to Article 35 of the DSA, Search is refining its removal policy for NCEI to make it even easier for people to report this content.

Investing in Search Information Quality

The systemic risk assessment reviewed several risks relating to a wide variety of harmful content, such as content impacting human dignity, promoting discriminatory beliefs, inciting or glorifying violence, promoting practices harmful to health, inciting gender-based violence, or constituting harassment and bullying. For harmful content, we deploy a wide array of measures to address risk, including Search ranking (such as surfacing credible and high-quality content over lower-quality content in web results) and content policy enforcement, especially in Search features.

However, Search has indexed hundreds of billions of web pages, images, videos, and other content, so Search results might occasionally contain material that some find objectionable, offensive, or problematic.

Content that Search has no legal obligation to remove and is not prohibited by our policies remains available in Search results for users who express an intent to explore that content, even if indicators suggest it is of relatively low quality or potentially harmful. While we believe our approach to be reasonable, proportionate, and effective in the context of a search engine service, the continued availability of this content results in medium levels of residual risk for several risk statements relating to harmful content in the fundamental rights and public health dimensions of the systemic risk assessment. Below we describe some of the mitigating measures we take.

Our automated systems are our first line of defence to limit the appearance of harmful content in search results for the most common queries, but we may also have trained experts who manually review and remove content that violates our Search features content policies.

Addressing Sensitive, Harmful, and Policy Violative Content

We use [ranking systems](#) to sort through hundreds of billions of web pages and other content in our Search index to present the most relevant and useful results in a fraction of a second. Our ranking systems are central to addressing systemic risks relevant to Search.

Search ranks and prioritises content using signals that align with meaning, relevance, quality, usability, and context. Our approach is to raise the ranking of the highest quality information, rather than removing low quality information. Our emphasis on the ranking rather than availability of content allows us to address the risk of harm in a proportionate manner and reduce risks to freedom of expression and information.

Our ranking systems are especially designed to surface high-quality content for what we call “Your Money or Your Life” (YMYL) topics, defined as those that may significantly impact the person who is directly viewing or using the content, other people who are affected by the person who viewed the content, or groups of people or society affected by the actions of people who viewed the content. YMYL topics can directly and significantly impact people’s health, financial stability or safety, or the welfare or wellbeing of society—for example, pages that offer financial advice or information regarding investment, taxes, retirement plans, loans, banking, insurance, or which facilitate purchases or online money transfers. You can read more about our more notable ranking systems in our [guide to Google Search ranking systems](#).

To help us [test and improve](#) our Search algorithms we put all possible changes to Search through a rigorous evaluation process to analyse metrics and decide whether to implement a proposed change. We work with external [Search Quality Raters](#) to evaluate the quality of these automated ranking systems based on the expertise, experience, authoritativeness, and trustworthiness of content. This approach to testing our ranking systems is explained more in [How insights from people around the world make Google Search better](#) and [An overview of our rater guidelines for Search](#).

Informing Users

Our [About this result](#) tool enables users to learn more about the result or feature and where the information is coming from (such as a description of the website), when we first indexed the site, and whether connection to the site is secure) so that users can make more informed decisions about the sites they visit and the results that are most useful to them. We’ll soon be bringing an [About this image](#) tool to Search to provide context about images indexed through Google Search.

The growth of disinformation and misinformation and the emergence of new technologies require our ranking methods to also continue to adapt to address evolving risks. Search continues to invest in methods to raise the ranking of most relevant and reliable information available and ensure effectiveness of our algorithms across languages, countries, cultures, and contexts.

Providing SafeSearch

Keeping people safe on Search also means helping them steer clear of unexpected, shocking results. One way we tackle this is with [SafeSearch settings](#), which help detect and manage access to explicit content like pornography and graphic violence in Search. SafeSearch filtering is turned on by default for Google accounts for people under 18. We also offer options for parents and schools to lock this on for supervised minors.

Over the past months, we have [expanded SafeSearch](#) to help further protect people from inadvertently encountering explicit imagery on Google Search. We have rolled out an additional SafeSearch setting that will blur explicit imagery if it appears in Search results. This new SafeSearch blurring setting has become the default for people who do not already have the SafeSearch filter turned on, with the option to adjust settings to SafeSearch “Filter” (blocking any explicit content) or “Off” at any time.

For Google Accounts for people under 18, we take additional steps to help minors make choices to avoid results that may be shocking or harmful. As part of this, SafeSearch is set to Filter automatically when our systems indicate that a user may be under 18. SafeSearch filtering is turned on for children under 13 (or applicable age of consent in the relevant country) signed in to an account managed with Family Link.

When SafeSearch is “Off,” users find all relevant results for their search, even if they are explicit, but our SafeSearch signals still apply to suppress irrelevant explicit content when the user does not appear to be seeking it out. In fact, every day, our safety algorithms improve hundreds of millions of searches globally across web, image, and video modes.

Supervised users are unable to change their SafeSearch setting—for example, for child and student accounts, parents and schools can lock SafeSearch, while parental controls on an operating system or antivirus software may override an individual’s SafeSearch setting. Parents can use Family Link to set up supervision on a child’s account, with SafeSearch filtering turned on automatically and locked so that the child cannot change the setting. You can read more in [Manage Search on your child’s Google Account](#).

Tailoring our Content Policies

In Search, we take a multi-tiered approach to content policies to balance the need to protect freedom of expression with providing users with high quality information.

Search policies apply to content surfaced anywhere within Search, which includes web results (i.e., web pages, images, videos, news content or other material that Google finds from across the web). Search’s policies cover essential content restrictions such as CSAM, spam, and valid legal requests. We maintain the following three [categories of content policies](#).

Search policies include a [highly personal information policy](#) under which we remove certain personal information that creates significant risks of identity theft, financial fraud, or other specific harms, such as doxing content, explicit personal images, and involuntary fake pornography. These policies were developed following an extensive stakeholder consultation to help inform how we balance taking action to protect user privacy and safety with the right to freedom of expression, and were [enhanced in 2022](#) to include the removal of additional personal information (such as contact information) from Search in cases that do not involve doxing.

Search features policies apply to many of our search features, such as Autocomplete, Featured Snippets, and Google Discover. The presentation of these features emphasizes and highlights the content differently than our relevance-based web results.

Search feature-specific policies explain how certain search features work, and set forth any additional feature-specific restrictions. Examples include prohibiting predictions about medically hazardous health claims on Autocomplete and applying a higher quality threshold for recommending content for YMYL topics on Discover.

Addressing Civics Misinformation

Search aims to enable users' informed participation in democracy by providing high-quality information that is accurate, up-to-date, and protected during elections. One way we ensure reliable information is returned to users in the elections context is through our use of classifiers to identify elections-related queries, so that our systems know that returning high-quality information is especially important. As part of this effort, Search has developed a number of features aimed at ensuring we show users trustworthy elections-related content from reliable third parties. These features are activated during elections and in response to elections-related queries to mitigate the risk of low-quality content and return organised search results pages that include comprehensive authoritative information in over 60 countries, including all EU Member States during their national elections. For the upcoming 2024 EU Parliamentary elections, we will work with Elections Commissions across EU member states to make electoral information available and help people find the info they need to get out and vote. These partnerships enable informational features such as How to Vote, How to Register, Where to Vote information, and Politician Knowledge Panels.

In addition to developing Search features with trustworthy third parties, we also utilise Search Quality Raters, described above, and elections-specific classifiers to ensure Search results surface factual information about key persons or entities associated with elections. For the 2024 EU Parliamentary elections, we will have elections classifiers operating in all official EU languages except Maltese and Irish.¹⁶

Search information quality processes, including ranking and our robust policies that allow us to remove violative content, consistently perform well in preventing conspiratorial information from surfacing to users. For example, in 2021, the Leibniz Institute for the Social Sciences conducted a comparative algorithm audit of presence of conspiratorial information in top search results across five search engines: Google, DuckDuckGo, Yahoo, Bing and Yandex. Their research found that “all search engines except Google consistently displayed conspiracy-promoting results and returned links to conspiracy-dedicated websites in their top results, although the share of such content varied across queries.”¹⁷

¹⁶ Low usage in these languages makes it difficult to adequately train classifiers targeted at a specific set of issues, such as elections.

¹⁷ Urmana A, Makhortykh M, Ullioac R, Kulshresthad J (2021) [Where the Earth is flat and 9/11 is an inside job](#): A comparative algorithm audit of conspiratorial information in web search results: Leibniz Institute for the Social Sciences.

Respecting Freedom of Opinion, Expression, Media Pluralism, and Civic Discourse

The systemic risk assessment reviewed several risks relating to content removal, users making informed decisions about what to view, and media pluralism (e.g., the plurality, polarisation, and diversity of perspectives available). Search's role in making information universally accessible resulted in relatively low inherent risks and low residual risks overall. These results are consistent with Search methods and initiatives like the [Search Quality Evaluator Guidelines](#), [Google News Initiative](#), and [Machine Learning Fairness Approach](#); Search endeavours to play an essential enabling role for the realisation, enjoyment, and fulfilment of these rights.

We believe, and studies have shown, that Search returns relevant and helpful sources with “no evidence of ideological bias” when users are looking for news.¹⁸ Our systems are not designed to favour or disfavour any particular publications based on ideology. Instead, our systems look at signals such as relevance, prominence, freshness, authoritativeness, or trustworthiness to determine the most helpful, relevant content to show users. Search does not take the political viewpoint of a webpage into account when ranking. In the [Google Search Quality Evaluator Guidelines](#), Search instructs evaluators that “[r]atings should not be based on your personal opinions, preferences, religious beliefs, or political views.”

Reputable studies consistently find that Search is fair. For example, two studies by *The Economist* evaluated claims of bias in Google News results and found no evidence of ideological bias, concluding that “Google rewards reputable reporting, not left wing politics”.¹⁹ An extensive study by academics at Stanford University drew a similar conclusion. Over a six-month period, researchers reviewed Search results appearing on the first page for every candidate running for federal office in the 2018 U.S. general election. Four million URLs were scraped from Search and audited, and the researchers found that Search results did not exclude sources from either the left or the right of the political spectrum.²⁰

Google Search is also best-in-class in displaying diverse results, which are indicative of strong support for media pluralism. In 2022, the Hamburg University of Applied Sciences published a paper reporting the findings of a study examining the difference between results retrieved by four major web search engines. Researchers compared the top 10 results from Google, Bing, DuckDuckGo, and Metager, using 3,537 queries generated from Germany and the US. The findings of the study showed that “Google displays more unique domains in the top results than its competitors, and Wikipedia and news websites are the most popular sources overall.”²¹

Addressing Disinformation

We believe that elevating authoritative information and combating misinformation and disinformation are of utmost importance to systemic risks relevant for Search. These efforts are especially relevant to issues such as public health, elections, and civic engagement. While the efforts described above all go to combat disinformation appearing on Search and across our services, there are other efforts that are critical to our holistic approach.

¹⁸ *The Economist* (2019) [Google rewards reputable reporting, not left-wing politics](#).

¹⁹ *Id.*

²⁰ Danae Metaxa, Joon Sung Park, James Landay, Jeff Hancock (2019) [Search Media and Elections: A Longitudinal Investigation of Political Search Results](#), proceedings of the ACM on Human-Computer Interaction, Volume 3 Issue CSCW. Article No. : 129.

²¹ Nurce Yagci, Sebastian Sünkler, Helana Häußler, Dirk Lewandowski (2022) [A Comparative of Source Distribution and Result Overlap in Web Search Engines](#); Hamburg University of Applied Sciences.

We implement a multi-faceted approach to address the complex challenges and risks raised by misinformation and disinformation across our services. While our ranking systems seek to connect people with authoritative sources and are described elsewhere in this report, we are cognizant that these are complex issues that no single actor is able to solve on their own.

For this reason, we have long welcomed the multi-stakeholder approach, including the EU's 2018 [Code of Practice on Disinformation](#) and a [Strengthened Code](#) that we signed in June 2022. As part of the Strengthened Code, we have committed to providing the European Commission with [reports](#) detailing how we have implemented our Commitments under the Code.

Our baseline report under the Code highlighted the breadth of our work across EU Member States to tackle the monetisation of disinformation, provide transparency on political advertising, detect and counter a range of threats to the integrity of our services, empower users, and work with the fact-checking and research communities. The report also provided information about the quantitative impacts of our work at the Member State level.

Following this baseline report, we expect to publish subsequent versions of this report biannually. In addition, we expect to remain a committed and productive member of the Code of Practice's Permanent Task-force.

Service Design

Addressing Unfair Commercial Practices and Fraudulent Content about a Business

We take many actions to mitigate the risks of unfair commercial practices and fraudulent content about businesses, such as prioritising the highest-quality results as part of the ranking process described above, removing policy-violating content from Search features, and removing fraudulent content subject to legal removal requests. However, this type of content will continue to be returned in search results when a user seeks it out, provided it is not the subject of a valid legal removal request or prohibited by Google policies (e.g., spam), and is still available on the internet. For this reason, the assessment found that some elevated residual risk of fraudulent business information appearing on Search remains.

We employ a higher standard and a different approach to address unfair commercial practices and fraudulent content about a business in the advertising context. Our [Ad Policies](#), which apply to ads on Search and our VLOPs, have several policies relevant to mitigating this risk, such as policies prohibiting misrepresentation (e.g., phishing, obscuring charges associated with financial services, misleading claims regarding weight loss or financial gain) and policies prohibiting the sale or promotion of counterfeit goods, dangerous products and services, and products or services enabling dishonest behaviour (e.g., hacking software, fake documents, or academic cheating).

Google Ads does not allow ads that deceive users by excluding relevant product information, such as billing details or charges, interest rates, fees, and penalties, or by providing misleading information about products, services, or businesses. This includes impersonating brands or businesses, concealing or misrepresenting a business identity, and implying endorsement by another individual, organisation, product, or service without their knowledge or consent. For egregious violations (those so serious that they are unlawful or pose significant harm to our users), we will suspend Google Ads accounts upon detection and

without prior warning, and not allow the advertiser to advertise with us again, unless an appeal brings to light compelling grounds for a different outcome.

Respecting Privacy

Privacy is an enabling right, furthering rights such as freedom of expression, association, opinion, religion, movement, and bodily security.²² Once violated, the right to privacy can be challenging to remediate because private or highly personal content can remain in circulation on the internet. Given the role of Search in surfacing information from nearly anywhere on the open web, privacy risks in the context of Search can be important inherent risks.

Search has addressed these inherent privacy risks by (1) ensuring responsible stewardship of user data by refraining from selling user data, constantly refining data collection practices, and providing users with easy-to-use data settings; (2) [providing avenues for highly personal information to be removed from Google](#) (described above) and respecting the “[right to be forgotten](#)”; and (3) complying with requirements under applicable data protection and privacy laws, including minimising the data being collected, purpose limitation, providing transparency to users.

To be responsible stewards of user data, we take a [private-by-design approach](#): we encrypt every search, build controls so that users can choose the privacy settings that are right for them, and never sell personal information to anyone. Search also offers privacy controls so that users can decide what to save to their Google Account and can turn on auto-delete to automatically delete data on an ongoing basis.

Since 2014 we have been responding to requests to delist content under European privacy law, which provides individuals with the right to ask search engines like Search to delist certain results for queries on the basis of a person’s name if the links in question are “inadequate, irrelevant or no longer relevant, or excessive.” We evaluate each request on a case-by-case basis, and may not delist content where there is an overriding public interest in the information remaining available in search results. Our [requests to delist content under European privacy law report](#) provides information and data about the volume of requests, the URLs delisted, the individuals submitting requests, and the content of websites and URLs identified in requests. Since 2014, we have received around 1.5 million requests to remove around 5.6 million URLs. We take our responsibility to ensure compliance with European privacy law seriously while being committed to providing access to information, and carefully balance these commitments when assessing each request. As a result, we have refused to delist around 50% of the requests we have received to date; a large majority of those refusals are sustained when challenged in court or before data protection agencies.

These measures are typically sufficient to reduce many privacy risks to much lower levels of residual risk. However, our assessment concluded that some elevated residual risks remain for Search, most notably the unintentional or malicious sharing of private or highly personal information in Search results. This information can be challenging to verify and requires that we be informed of and verify a privacy violation before removing content from search results.

²² [UN Special Rapporteur on the right to privacy](#).

Protecting Children's Rights

Our services provide vital opportunities for learning, communication, and social interaction, and can be formative for a child's cognitive and social development. However, these opportunities are accompanied by risks to which children are particularly vulnerable given their unique stages of development, nascent digital literacy, and evolving cognitive abilities and decision-making skills. It is important for us to address these risks with mitigations, such as user guidance and parental controls, that help children navigate their online experiences now and over the course of their lifetimes.

The systemic risk assessment reviewed several risks relating to children's rights, such as behavioural addictions in children, use of children's data for ads targeting, and unnecessary or disproportionate limitations on children's access to Search. We found the highest inherent risk to be the risk that children under a defined minimum age access services that they should not be able to, and may be exposed to harmful, hateful, or age-inappropriate content or conduct. Based on the mitigation measures described below, we concluded that Search (and our VLOPs) are taking actions that significantly reduce residual risk for children's rights.

The following protections apply horizontally across all our services, and thus pertain to Search and our four VLOPs. These protections will be described here and cross-referenced in the VLOP sections. Our efforts in this space must balance adults' rights to access services and information with a reasonable level of privacy, and the need to protect children from accessing services and information that are not appropriate for their age.

Obtaining Age Assurance

We require users to manually enter their date of birth (without pre-populated options, referred to as a neutral age-screen) during Google Account sign up to help determine which users are likely under the age of 18 so that we can apply heightened privacy, content, and safety protections. To reduce the burden on our users and in accordance with data minimisation principles, these processes are carried out at the [Google Account level](#) so that the results can then be used in connection with all signed-in services (including Search) that are accessed by the user.

Depending on what birth date a user provides at the time of Google Account creation, we apply different protections.

- If a user provides an age that is under [the minimum age to have a Google Account in their country](#) we require approval from a parent/guardian before continuing with account creation, and the account must be supervised until the user attains the minimum age (see further information on Family Link below).
- If a user provides an age under 18, we apply a number of default protections to the account, and we disable access to age-restricted content across some of our services. Parents/guardians who manage their child's account through Family Link may choose to change some of these default settings if a different approach works best for their family.

We independently assess whether or not a user is likely an adult, both for users who sign into their accounts and those who do not. We use a variety of signals, such as the types of sites a user has searched for or the categories of videos a user has watched on YouTube, as well as indicators like the longevity of an account. For example, searches for mortgage lending sites or tax assistance might be signals that the user is likely an adult. Once our model has sufficient signals about a user's age, it sends a signal to our services to automatically set appropriate default settings and protections, such as by turning SafeSearch filtering on for those under 18. This approach does not involve collecting additional information from users.

Enabling Parental Control

Family Link parental controls are available in the Family Link app and also via web browsers. Parents/guardians of minors [under the applicable minimum age](#) can create Google Accounts for their children and must manage those accounts using Family Link parental controls. Family Link parental controls are also available for parents/guardians to supervise minors over the applicable minimum age, but consent from these minors is required before supervision may be enabled.

Family Link helps parents/guardians stay informed about and manage their child's experience on compatible Android and ChromeOS devices. For example, Family Link empowers parents/guardians to set digital ground rules for their family by managing the apps their child can use, keeping an eye on screen time, or setting a bedtime and daily limits for their child's device. These controls help parents/guardians manage their child's experience in ways that make sense for their family.

Providing Ads Protections

We prohibit age-sensitive ad categories from serving to users under 18, including ads that feature adult or sexually suggestive content, alcohol, or gambling and games. We also prohibit the display of personalised ads based on age, gender, or interests to any users we determine to be under the age of 18. Ads shown to these users must meet our under-18 ads policies, and may only be served based on non-personalised contextual information, such as the content on the current site a user is visiting.

Providing SafeSearch on by Default

Using age-appropriate default settings is one way that we incorporate "safety by design" into our products. Specific to Search, [SafeSearch](#) filtering is set on by default for Google accounts for children younger than 18, and parents and schools have the option to lock it on for supervised minors. As described above, SafeSearch filtering blocks explicit content (like sexual activity and graphic violence) from Search results across images, videos, and websites - when the filter setting is on, explicit results will be filtered even when they might be relevant for the query.



Google Maps

Maps

Description of Service and Associated Risk Profile

Google Maps is a service that helps users navigate and explore the world. The service also includes accurate and reliable information about places, business, and experiences, and helps businesses build an online presence, engage with customers, and grow their business. Some of the key elements that make the service compelling include satellite/aerial views, digital street maps, information about places and business, 360° interactive panoramic views of streets, real-time traffic conditions, and route planning for driving, walking, cycling, public transportation, and flying. The information about places and businesses includes some user-generated content, such as content from [consumer users](#) and [merchant users](#) (including ratings, reviews, and photography) and content provided by the merchant users interested in [listing](#) or [advertising](#) their business on Maps.

Google Maps is free, available in over 100 languages, and used by around 275 million users in the EU every month.²³ You can read more in [Google Maps Help](#) and our [Maps 101](#) blog series.

The primary purpose of Google Maps is to help users navigate from place to place and explore the world, with elements such as images, reviews, and information about places being in service of that goal. This systemic risk assessment validated that the most important risks are not intrinsic to the primary purpose of maps—helping users to get from A to B—but associated with the various features designed to enhance the user experience when fulfilling this purpose.

The service emphasizes being a source of reliable information and a reflection of genuine user experiences. For this reason we lean towards user-generated [content policies](#) that are designed to maximise the quality, accuracy, and authenticity of information for consumer and merchant user contributions. We go to great lengths to make sure content published by our consumer and merchant users is helpful and reflects the real world, recognising that this means accepting some attendant limitations to freedom of expression. Google Maps is designed for a low likelihood of content going viral, reducing inherent risks associated with illegal and policy violative content.

Risks relating to conducting a business (e.g., unfair commercial practices, such as paying, incentivizing, or encouraging the posting of positive or negative reviews that do not represent a genuine experience) are important to address given the role of Maps in connecting potential customers with businesses (e.g., helping users find the best restaurant in town or a reliable auto repair shop).

²³ Average monthly counts based on distinct signed-in accounts of recipients.

Finally, the locational nature of the Maps service, combined with the existence of user-generated content, makes it important to address privacy-related risks (e.g., data practices and risks to personal information, such as unintentional or malicious sharing of private or personal information), as we describe in the “Service Design” section below. The privacy of our users is of utmost importance to us, and while we welcome content that illustrates our world, it's critical to do so in a manner that respects users' right to privacy.

Systemic Risk Assessment Results and Associated Observations

We assessed 37 different risk statements²⁴ for inherent risk (i.e., risk absent any action taken by Google), preparedness (i.e., the cumulative measures currently in place to mitigate the risk) and residual risks (i.e., risk after mitigation by Google). Residual risk serves as a guide for where further investment may be warranted. The full list of risk statements is found in Annex A to this report.

This systemic risk assessment surfaced important themes relating to the inherent and residual risk. Overall, we found that the highest assessments of Google Maps' existing mitigations were correlated with the highest inherent risks, confirming that we have been prioritising action on the most significant risks.

The two most significant inherent risk themes for Maps are directly related to the nature of Maps: risks associated with information about businesses shown on Maps (e.g., unfair commercial practices, disinformation) and risks associated with the locational nature of Maps (e.g., privacy-related risks). We have long recognised these risks and our efforts to address them result in much lower residual risks; however, we did assess that bad actors who may engage in the malicious sharing of private or highly personal information continue to pose medium levels of residual risk.

We also found higher inherent risks relating to illegal and harmful content, but concluded that these have much lower residual risk given emphasis on the quality, accuracy, and authenticity of information, our substantial efforts to enforce content policy, and the low likelihood of content going viral on Maps. However, optimising this way does result in medium levels of residual risk to freedom of expression.

Article 34(1) of the DSA encompasses a range of systemic risks that are interconnected and cannot be dealt with in isolation; our policies and practices for Maps often address multiple risks at the same time. To provide a comprehensive understanding of Maps' existing mitigating practices and align with Article 35 of the DSA, we have categorised specific manifestations of systemic risks into groups for efficient explanation of existing mitigations and discussion of planned improvements.

In the next two sections we consider content on Maps, including the development and enforcement of content policies ("Content Moderation"), and explore service design choices that target risks associated with Maps' functionality, including privacy ("Service Design").

Taken in combination, these two sections address each of the four systemic risk categories outlined in Article 34(1) of the DSA. We emphasize areas where the assessment has identified elevated inherent or residual risks, elucidating the measures already implemented by Maps to tackle these risks, as well as any future plans to address systemic risks as appropriate.

²⁴ See Methodology Step One: Classification.

Content Moderation

In this section we show how Maps has designed and enforced its user-generated content policies to address the systemic risks articulated in Article 34(1) of the DSA. We detail some risks in the assessment with elevated inherent and residual risk, and describe what Maps is doing and plans to do about those systemic risks.

Removing Illegal Content

Google Maps has clear policies in place prohibiting illegal content on the service through its policies related to “[regulated, dangerous, and illegal](#)” user contributed content. This includes images or any other content that infringes on anyone else’s legal rights, including copyright, as well as content that relates to terrorism, sexual abuse imagery or sexualization of children, dangerous or illegal acts (such as rape, organ sale, or human trafficking), or illegal products and services (such as endangered animal products and illegal or diverted drugs). We also disallow potentially illegal online activity such as doxxing or content that contains a specific threat of harm or depicts illegal activity.

Despite elevated inherent risk related to illegal content, such as CSAM, terrorist content, and illegal activity, the results of the systemic risk assessment showed Maps’ preparedness—such as taking legal action to stop fake review scams and tackling fake contributions—enabled it to achieve low levels of residual risk for illegal content and activity. As explained earlier in this report,²⁵ we have a well-developed process for responding to legal orders to remove content, and the efforts described below to enforce our policies ensure the removal of content that is illegal or that violates our policies.

Addressing Content that Violates our Policies

The systemic risk assessment reviewed several risks relating to a wide variety of potentially harmful content, such as content impacting human dignity, promoting discriminatory beliefs, inciting, praising, or glorifying violence, promoting practices harmful to health, inciting gender-based violence, or constituting harassment and bullying. However, the enforcement of Google Maps’ user-generated content policies, which favour authoritative information and genuine experiences, lowers residual risks.

Developing Content Policy

We have created strict policies to make sure that user-generated content is based on real-world experiences and to keep irrelevant and offensive comments off Google Maps. Our policies against topics like fake engagement, misrepresentation, and misinformation continually evolve in responding to changing threats.

Our [user-generated content policy](#) describes our overall approach while our [prohibited and restricted content policies](#) clearly set out what is not allowed on Google Maps, covering civic discourse, deceptive content, mature content, regulated goods and services, dangerous and illegal content, and low-quality information. These user-generated content policies are more restrictive than for many other Google services, reflecting our increased emphasis on relevant, authoritative information and genuine experiences for Maps. Our “off-topic” and “fake engagement” policies are good examples of Maps’ unique approach to

²⁵ See *supra* at p. 20 (“Handling Government Removal Requests”).

content. These policies have evolved over time to guard places and businesses from violative and off-topic content when there's potential for this type of content to lead to harmful and targeted abuse. For example, when governments and businesses started requiring proof of COVID-19 vaccination before entering certain places, we put extra protections in place to remove Google reviews that criticised a business for its health and safety policies or for complying with a vaccine mandate.

Other policies relevant for content on Maps include our [Local Guides Program Terms and Conditions](#), which set out who is qualified to be a Local Guide and appropriate conduct standards; [Google-Contributed Street View Imagery Policy](#), explaining how we treat inappropriate content and the criteria we use for publishing Street View imagery to Google Maps; [My Maps Content Policy](#), which sets out policies for creating and sharing custom maps; and [Guidelines for representing your business on Google](#), setting out guidelines for Business Profiles.

Once a policy is written, it's turned into training material—both for our operators and classifiers—to help our teams catch policy-violating content and behaviour.

Enforcing Content Policy

Contributions to Google Maps should accurately represent the location in question. Where user-generated contributions distort truth we remove content, including reviews, photos, or videos not related to the location or business where they are tagged. If user-generated content is inaccurately placed on the map, or is associated with an incorrect listing, the contribution may be rejected. When a user submits a review, we automatically send it to our system to make sure the review doesn't violate any of our user-generated content policies before posting the review. Given the volume of reviews we regularly receive, we have found that we need both the nuanced understanding that humans offer and the scale that automated detection provides to help us moderate contributed content.

Undertaking Automated Detection and Removal

Automated detection is our first line of defence because automated systems are good at identifying patterns that help determine if the content is legitimate. This includes whether the review contains offensive or off-topic content, whether the Google account has a history of suspicious behaviour, such as a history of posting violative content, and whether there has been uncharacteristic activity, such as many reviews over a short period of time. The vast majority of fake and fraudulent reviews are removed before anyone sees them because all reviews are run against classifiers before being posted.

Our human operators regularly run quality tests and complete additional training to remove bias from the machine learning models. By thoroughly training our models on all the ways certain words or phrases are used, we improve our ability to catch policy-violating content and reduce the chance of inadvertently blocking legitimate reviews from going live. We review and update our classifiers, including review for quality and accuracy across language, gender, ethnicity, and religion, and our assessment identified this as a priority for continuous improvement over time.²⁶

²⁶ See *supra* at p. 22 (“Evaluating Content Across Languages”) for further discussion of how Maps, and Google as a whole, are addressing this identified residual risk for Maps.

If our systems detect no policy violations, then the review can be posted within a matter of seconds. However, our automated systems continue to analyse the contributed content and watch for questionable patterns, such as a group of people leaving reviews on the same cluster of Business Profiles or a place receiving an unusually high number of 1- or 5-star reviews over a short period of time.

In addition, we make it easy for people using Google Maps to flag any policy-violating reviews, with [businesses](#) and [consumers](#) both able to report reviews and [flag inappropriate user profiles](#).

Undertaking Human Review

A team of human operators work alongside automated systems to remove reviews that violate our policies, and when appropriate suspend user accounts. We deploy thousands of trained operators and analysts globally who help with content evaluations that might be difficult for automated systems, such as understanding reviews that include local slang.

Undertaking Enforcement Proactively

In 2022 we launched a significant update to our machine learning models that helped us identify novel abuse trends many times faster than in previous years. For example, our automated systems detected a sudden uptick in Business Profiles with websites that ended in .design or .top, which our team of analysts quickly confirmed to be fake. They were therefore able to quickly remove the Business Profiles and disable the associated accounts.

These new machine learning models, paired with our advanced automated and manual techniques, helped us block or remove over 115 million policy-violating reviews globally in 2022 (the majority of them caught before they were ever seen by a user, and over 20% more fake reviews than in 2021), and block over 20 million attempts to create fake Business Profiles globally (8 million more than in 2021). We also put protections in place for more than 185,000 businesses after detecting suspicious activity and abuse attempts. In addition, we blocked or removed over 200 million photos and 7 million videos that were blurry, low-quality, or violated our content policies. The success of these improvements was a primary factor in our assessment that Maps has controls in place for risks related to policy violative content that result in much lower residual risk in these areas.

Additionally, we [took legal action](#) to fight malicious actors who violated our policies, and are sharing our best practices with government agencies to find lasting solutions for the whole industry. We filed a lawsuit that successfully took down a group of fraudsters who were impersonating Google through telemarketing calls and attempting to sell fake reviews online, building on our previous legal action against internet scammers and malware operations. Going forward, Maps will continue to invest in new technologies and processes to keep information on our services helpful and reliable.

Posting Restrictions for Repeat Violators

When we find that user contributions for certain types of places are consistently unhelpful, harmful, or off-topic, we may limit or suspend user-generated content for those places. Maps has developed a measured response regarding [posting restrictions](#):

- Short-term restrictions, when posting may be turned off for a particular place for a short period of time to help protect the place or area from a spike in irrelevant or offensive content.
- Long-term restrictions, when posting on a particular place may be turned off for a longer period of time if its category or geographic area has experienced a continuous pattern of low value or off-topic posts.
- Partial or full restrictions, when, depending on the volume and pattern of policy violating content, a particular place may have posting restrictions on some or all of the types of user-generated content (including Text Reviews, Ratings, Images, and Videos).

We found that policies such as our posting restrictions greatly reduce the opportunity for repeat offenders to manipulate our systems through inauthentic use, reducing residual risk across the board.

Assessment Results for Specific Content Risks

Protecting Civic Discourse

We assessed the risk that misinformation and disinformation relating to elections, civic discourse, democratic participation, or civil unrest may be available on Google Maps. While this risk may occur in the context of user-generated content, Google Maps is designed for a low likelihood of content going viral, and more severe outcomes (such as influencing the result of an election) are highly unlikely. Our preparedness for elections and prohibition of any information that may be deceptive or misleading about civic processes, newsworthy events, or civic discourse significantly reduce residual risk.

In addition to reviewing flagged content, our team proactively works to identify potential abuse risks, which reduces the likelihood of successful abuse attacks. For instance, when there's an upcoming event with a significant following—such as an election—we implement elevated protections for the places associated with the event and other nearby businesses that people might look for on Maps. We continue to monitor these places and businesses until the risk of abuse has subsided. To avoid the spread of election-related misinformation, we prevent people from editing the phone numbers, addresses and other information for places like voting sites.

Protecting Consumers and the Freedom to Conduct a Business

Google Maps provides information to users that enable them to find and navigate to a business. For this reason, prominent inherent risks include the risk that disinformation, misinformation, or fraudulent content about a business is available or that unfair commercial practices take place on Google Maps. This is typically driven by intentional manipulation of the Google Maps service and might include positive or negative fake reviews, “review bombing” by competitors, fraudsters creating false business listings, or “predatory removals,” which occur when a bad actor demands payment for the removal of fake reviews. These risks can disproportionately impact less technologically literate users and newly opened businesses, which are typically more vulnerable than established brands.

We work to stay ahead of scammers and protect small businesses by continuously monitoring for fraudulent content on our products, using a combination of people and technology. One of the best tools we have to fight back is our understanding of inauthentic use patterns on Google Maps, which informs our classifiers. These classifiers detect and remove policy-violating content across a variety of languages, and also scan for signals of abnormal user activity.

Our teams and protections are built to fight two main types of bad actors: content fraudsters and content vandals.

Fraudsters, who are ultimately motivated by money, try to trick people with scams like fake reviews to attract customers or fake listings to generate business leads. To deter them, we preemptively remove opportunities for them to profit from fake content, and have focused efforts on detecting content coming from click farms where fake reviews and ratings are being generated. Through better detection of click farm activity, we are making it harder to post fake content cheaply, which ultimately makes it harder for a click farm to sell reviews and make money.

Content vandals, who may be motivated by social and political events or simply want to leave their mark online, often post fake reviews or edit the names of places to send a message, or add off-topic photos as pranks. Content vandalism can be more difficult to tackle than fraud as it is often random. Impeding content vandals requires anticipation and quick reaction, and as certain places become more prone to vandalism, we adjust our defences—such as when we modified our algorithms to preemptively block racist reviews when we observed anti-Chinese xenophobia associated with COVID-19.

These risks are further mitigated by the implementation of our Ads policies, such as the [Misrepresentation Ads Policy](#), which disallows ads that deceive people, and the [Restricted Businesses Policy](#), which restricts certain kinds of businesses with products prone to abuse. These ads policies are complemented by relevant Google Maps User Contributed Content policies, such as the [Misrepresentation Policy](#), which doesn't allow users to mislead or deceive others, and the [Impersonation Policy](#), which doesn't allow users to impersonate any person, group, or organisation.

We have long recognised these inherent risks as priorities, and our wide range of measures to remove policy violating reviews, stop fake Business Profiles, and protect targeted businesses serve to reduce these to much lower residual risks. As described above, our classifiers are primarily responsible for the successful reduction of risk in this area. Our actions here also address equality and non-discrimination risks, since this activity can disproportionately be targeted at those under-represented as content contributors, such as minority businesses.

Respecting Freedom of Opinion, Expression, and Media Pluralism

The systemic risk assessment explored several risks relating to content removal, users reporting potentially violating content, and media pluralism (e.g., the plurality, polarisation, and diversity of perspectives available). Maps' focus on providing topical and authoritative information and genuine experiences rather than being a forum for dialogue lowered freedom of expression and media pluralism as an inherent risk. Maps does remove high volumes of content that violates our policies against off-topic or misleading information about locations and businesses, and for this reason the systemic risk assessment found that some medium residual risk remains for over-moderation of user-generated content. We believe that this approach is reasonable, appropriate, and proportionate given the nature, purpose, and intended use of Maps.

While Google Maps' merchants have long been able to appeal potentially erroneous account and listing suspensions, we have not in the past offered appeals for user-generated content. This resulted in lower preparedness evaluation for the risk of erroneously removing content. Appeals offer users a path to redress and also gives Maps better information for improving in the first instance. Accordingly, and pursuant to Articles 20 and 35 of the DSA, Maps is now establishing and implementing a new appeals mechanism for users to challenge potentially erroneous removals of their content. We believe these mitigation steps will increase our preparedness and lower residual risks to freedom of expression over time.

Service Design

Respecting Privacy

Some prominent inherent risks for Maps relate to privacy, reflecting the locational nature of the Maps service, the existence of user-generated content, and challenges associated with reversing privacy impacts once they have occurred. Our privacy risks and mitigations cover three dimensions: users of Maps; contributors to Maps; and images shown on Maps that may involve users, non-users, and contributors.

Protecting Users of Maps

Google Maps uses location data to make its service functional and useful for users. Real time [location information](#) plays a very important role for Google Maps, such as assisting in providing accurate driving directions, the latest transit status, and useful search results. The [Google Privacy Policy](#) governs how user data is collected and used by Maps and other Google services and is designed to ensure that we collect data only where it is necessary for the user's intended purpose. In addition to the use of real time location data, users may turn on Location History in their Google account settings to opt into preserving precise historical location data. Location History is off by default. On Maps, real time location data is used even when Location History is off, and people who use our services can also choose to share (or not share) their real time location with others regardless of their choice of settings for Location History. Our well-established policies, procedures, and options for users result in low residual risk for Google collecting, processing, aggregating, or sharing more user information than is necessary for the stated purposes. This reflects a deliberate investment in our highest inherent risks.

Protecting Contributors to Maps

The nature of user-generated content tied to locations on Google Maps opens up the potential for unintended or malicious disclosure of private or highly personal information about users attached to a specific location. While the complexity of data choices, and the link between content and location, means that the unintentional sharing of information will always remain a risk, we do have many well-established measures such as blurring and user choice that result in much lower levels of residual risk. However, we assessed that bad actors who may engage in the malicious sharing of private or highly personal information continue to elevate the level of residual risk.

Addressing Risks Relating to Images on Maps

We take several steps to protect the privacy of individuals when Street View imagery is published to Maps. We have developed cutting-edge face and licence-plate blurring technology that is designed to blur

identifiable faces and licence plates within Google- and user-contributed imagery in Street View. If we do not automatically or completely blur an image, users and non-users can [request that Maps do so](#) if their face or licence plate requires additional blurring, or if they would like us to blur an entire house, car, or body.

Protecting Children’s Rights

The nature and purpose of Google Maps (helping users navigate from A to B, making available accurate and reliable information about places, business, and experiences) results in lower levels of inherent risk for children’s rights. With policies designed to ensure accuracy and relevance of content we are able to place fewer restrictions on children when compared to other Google services.

The Google Maps experience is largely the same for children, except that those under 13 (or the minimum age in their country) cannot contribute content (including photos, ratings, and reviews), publish public place lists, add or edit places on the map, or turn on Location History. Children under 13 (or the minimum age in their country) whose accounts are managed with Family Link can only share their real-time location with their parents, and won’t see where they went with their devices or get recommendations based on visited places. You can read more in [Google Maps and your child’s Google Account](#).

As described in more detail above, Google Ads policies including the specific policies on Ads allowed on “made for kids” content and “ad-serving protections for teens” apply to ads shown on Google Maps. We prohibit personalised ads to any users determined to be under the age of 18, for whom ads may only be served based on non-personalised contextual information, such as the content being viewed.

Because of the underlying nature of the service design, the safety functionality built into Google Maps resulted in lower residual risk for children’s activity.



Google Play

Description of Service and Associated Risk Profile

On Google Play, users find and download their favourite apps, games, movies, books, and more. Google Play provides millions of apps and games to over 280 million users²⁷ in the European Union. Google Play [ranks and organises apps in order to](#) help users discover the most relevant apps for them on Google Play through features such as categories, For You, and recommended for you. Ads and sponsored content are clearly marked.

Google Play also connects over three million developers to billions of users worldwide and invests in the platform, tools, services, and marketing opportunities that support developers. This investment allows small or nascent developers to benefit from economic opportunity and contribute to a healthy, competitive ecosystem. In fact, 97% of developers pay no service fees to benefit from Google Play. We believe higher numbers of active developers, subject to compliance with our consumer-protection policies, results in wider choices for users.

This report will primarily focus on apps and games (collectively referred to as apps) as the main drivers of systemic risk relevant to Google Play. Our assessment also examined the other forms of content on Play, such as movies and books, and found they posed less risk to users because of robust policies (some of which are addressed below), and more standardised content lacking the dynamic data and user-generated challenges inherent in apps.

We take our responsibility to provide a safe and trusted experience for all users very seriously and provide a platform for developers to deliver apps safely to billions of people worldwide. To help achieve this, we establish and seek to enforce clear expectations via our [Google Play Developer Program Policies](#), which cover topics such as restricted content, privacy, malware, and monetization. We also help keep users safe by building protections into Google Play, requiring developers to follow high safety standards. You can read more in our description of [How Google Play Works](#).

Google Play is a “platform of platforms.” Many of the apps available through Google Play are also platforms themselves; in these instances, the app hosted on Google Play is the front door into a user experience controlled by the third-party app or game developer. This structure creates two dimensions of risk, which you will see reflected in our systemic risk assessment: risks related to the Google Play platform (e.g., hate speech in a review left on Google Play) and risks created by third-party apps on the Google Play platform (e.g., hate speech in a post within a social media app).

²⁷ Average monthly counts based on distinct signed-in accounts of recipients.

This separation reflects the appropriate allocation of systemic risk among Google Play and the apps that appear on the Google Play platform. While Google Play's risk assessment references both dimensions of risk, it is focused on our role in the mitigation of risks to the Google Play platform, with app-level mitigations most appropriately taken by the developers of those apps. As Recital 27 of the DSA notes, requests or orders related to the removal of illegal content should be "directed to the specific provider that has the technical and operational ability to act against specific items of illegal content, so as to prevent and minimise any possible negative effects on the availability and accessibility of information that is not illegal content."

Systemic Risk Assessment Results and Associated Observations

We assessed 40 different risk statements²⁸ for inherent risk (i.e., risk absent any action taken by Google) preparedness (i.e., the cumulative measures currently in place to mitigate the risk) and residual risks (i.e., risk after mitigation by Google). Residual risk serves as a guide for where further investment may be warranted. The full list of risk statements is found in the Annex A to this report.

This systemic risk assessment surfaced important themes relating to the inherent and residual risk.

In the first of the two sections that follow (“Content Moderation”) we consider risks and mitigations relating to content moderation on Google Play, which primarily pertains to apps themselves as content. While there is inherent risk of illegal or harmful content appearing on apps, this section explains Google Play’s app review and moderation program, which results in much lower levels of residual risk. We also discuss how we address other types of user-generated content, such as reviews on Google Play.

In the second of the two sections that follow (“Platform Design”), we consider risks and mitigations related to the way Google Play functions. Three important inherent risk themes that emerged during the assessment were related to privacy, security, and child rights. These risks reflect Google Play’s role in the overall app ecosystem, and this section explains the actions we take that result in much lower levels of residual risk.

Taken together, these two sections address the four broad categories of systemic risks articulated in Article 34(1) of the DSA and the specific manifestations of those systemic risks that we evaluate. This report emphasizes those risks for which the assessment showed elevated inherent or residual risk, and describe Google Play’s current risk mitigation practices as well as improvements pursuant to Article 35 of the DSA.

Content Moderation

Removing Illegal Content

Google Play has appropriate policies in place prohibiting illegal content on the platform through policies related to [restricted content](#), [intellectual property rights](#), and [other policies preventing fraudulent or malicious apps](#). As discussed previously in this report, Google has a developed process for evaluating government requests to remove content.²⁹ Additionally, Play has reporting channels for users to report illegal content or content that violates Google Play policies. There is extensive overlap between content prohibited by Google Play’s product policies and content that is illegal, meaning that our policy development and enforcement efforts work to mitigate the risks of both illegal and policy-violative content. These efforts are described in detail in the next section.

The Google Play risk assessment identified a range of relevant illegal content-related inherent risks, including risks such as CSAM, terrorist content, apps infringing intellectual property rights, and illegal

²⁸ See Methodology Step One: Classification.

²⁹ For more information on how YouTube and Google respond to government requests to remove content, see *supra* at p. 20 (“Handling Government Removal Requests”).

activity like scams. However, the assessment found that robust policies that are binding on apps (as described below) and enforcement of these policies resulted in much lower levels of residual risk.

As noted above, illegal content appearing within apps available on Google Play is primarily the responsibility of developers, though (as described below) Google Play takes a variety of enforcement actions against developers with multiple or egregious policy violations.

Addressing Content that Violates our Policies

The systemic risk assessment reviewed numerous risks relating to a wide variety of harmful content, such as content impacting human dignity, promoting discriminatory beliefs, inciting, praising, or glorifying violence, promoting practices harmful to health, inciting gender-based violence, or constituting harassment and bullying.

There were several factors that caused us to conclude that these risks are of much lower residual risk for Google Play. There are high costs associated with developing an app as compared to a single piece of user-generated content, thus well-designed policies ensure developers are effectively disincentivised from spending time and resources developing apps that clearly violate Google Play Policies. Additionally, when app developers submit their apps to Google Play, we use a combination of automated processes and human review to assess these apps before they can be published for distribution on the Play Store. The automated processes—which include static and dynamic components—scan an app’s code, app images, the developer profile, and the app description.

We review millions of apps submitted to Google Play each year, including technical reviews of code for malware. If we identify policy violations at this pre-publication stage, we reject the developer’s submission and give the developer an explanation of the policy issues along with instructions on how to correct them. Once the issues are addressed, the developer can resubmit their app or app update for further review. If we find no policy violations, we publish the app or app update to the Google Play store. As explained below, we also have robust developer education processes to keep apps in compliance with evolving policies and enforcement mechanisms when they fall out of compliance.

These estimates of residual risk also rest on the distinction raised earlier: between user-generated content moderation that Google Play can undertake and content moderation responsibilities held by app developers, who may themselves be providing a user-generated content platform.

For example, we establish clear requirements around robust, effective, and ongoing [user-generated content \(UGC\) moderation](#) in apps. But only an app developer operating a UGC platform can remove specific pieces of content (e.g., a specific post in a social media app; a video from a streaming app) from its platform. Google can only remove the UGC platform app in its entirety—including all legitimate content within it—from Google Play. This limitation raises questions of proportionality, fairness, freedom of expression, and user impact, all of which must be balanced against the risks that may be posed by the specific underlying content.

Maintaining Developer Policies

Our Google Play [Developer Policies](#) set out what developers can and cannot provide users on the Google Play platform and are the foundation upon which Play delivers apps and games safely to billions of people worldwide. These policies cover areas such as restricted content, impersonation, monetization and ads,

privacy, malware, and mobile unwanted software, and are relevant across most of the risk statements included in this systemic risk assessment.

Our [User Generated Content \(UGC\) Policy](#) requires providers of apps that contain UGC services to implement ongoing UGC moderation and sets out requirements in the areas of informed consent, defining objectionable content and behaviours on the app, and undertaking reasonable UGC moderation that is consistent with the type of UGC hosted by the app. It also includes requirements for reporting channels and an in-app system for blocking UGC and users. We believe that requiring providers of apps that contain UGC services to implement these policies—and requiring developers to play a role in enforcing them—provide an appropriate, proportionate, and effective approach for Google Play.

Developing Policy

We update the [Developer Policies](#) over time to reflect insights into new and emerging risks, and conduct regular reviews of our policies based on developer feedback, external media, expert and stakeholder feedback, and internal enforcement data. The practice of constantly updating our policies based on emerging threats also reduces many or most of the systemic risks considered in this risk assessment. Over the years, Google Play has taken strides in developing our policies such that the risks posed by apps and the content that appears on apps have been appreciably lowered.

For example, we created the UGC Policy mentioned above in response to the evolution of social media platforms, which included significant fleeting and/or real-time content, and new types of UGC, which had resulted in heightened societal concern about user safety. After market research, user studies, and collecting insights from developers in global markets, Google Play established a developer policy requiring in-app moderation for all UGC apps. Google Play's UGC Policy reflects the belief that users should have a direct means to contact social media platforms, which should be held accountable for consistently moderating content.

When specific user harm concerns arise, Google Play's policy development team goes through a rigorous process to understand the issue, develop guardrails, internally test those guardrails, and then introduce new policies into our ecosystem. After introduction, we monitor the impact of our policies to refine or expand protections, as needed.

For example, in 2021, we introduced our “Personal Loans” policy. This policy was developed in response to user feedback from India and several Southeast Asian markets that developers were charging high and often illegal interest rates to users, and that some developers were blackmailing users with the permissions they had obtained through apps. We mandated declaration and disclosure of financial agreements between users and financial apps or their developers, so that we can verify the legitimacy of the loan agreements and make sure the loan terms are clear to users. We later expanded the policy to prohibit these apps from accessing sensitive data, such as photos and contacts, which were being used to verify credit worthiness in markets without formal credit-scoring mechanisms.

In 2020, we added a [“Stalkerware” policy](#) to address code within apps that collects personal or sensitive user data from a device and transmits the data to a third party. Our policy requires prominent disclosure and consent for a narrow set of permissible uses and prohibits these apps for all other uses. Only apps designed and marketed for enterprise management or for parents to monitor their children's activities are

allowed to have such functionality. Google Play prohibits apps used to track anyone else, even with their knowledge and permission.

We recognise that smaller developers may have fewer resources to help them understand our policies or keep up with changes, so over the last three years we have expanded our education and support efforts. We now offer the [Google Play Academy](#), where developers can take courses to better learn our platform, and [PolicyBytes videos](#) about policy updates. We stream global webinars throughout the year where we make major policy announcements, and we offer the [Google Play Developer Help Community](#) for developers to get advice from other expert developers.

Well developed policies that are constantly being reexamined and updated, and which are binding on developers, were a significant factor in lowering the content-specific residual risks on Google Play.

Enforcing Policy

If an app does violate any of our policies, we take appropriate, necessary, and proportionate action pursuant to our [enforcement](#) processes. These actions may include app rejection (for apps and app updates submitted for review prior to being made available on Google Play), app removal (for existing apps), app suspension, limited visibility, limited regions, and account termination (for multiple suspensions or suspensions for egregious policy violations). Additionally, Google Play users can [report an app policy violation](#) and flag individual app reviews as inappropriate through a link on the Google Play listing. We [offer an appeals mechanism](#) for developers who believe there has been an enforcement error.

In order to protect developer's rights, when we deploy new policies on Google Play, developers generally have at least 30 days from the announcement of the new policy to make changes to their apps, and longer if the updates are likely to require significant technical changes. Because app removal can negatively impact users and developers, in addition to giving time for compliance, we invest heavily in efforts to educate developers about our policies and how to comply. Education lessens the need for enforcement and keeps well-intentioned developers and apps on Google Play. Enforcement of, and education about, our policies are key aspects of Google Play's moderation and user safety program that resulted in lower residual risks for much of our assessment.

Addressing Specific Content-Related Risks

Specific content-related risks feature less prominently as inherent or residual risk for Google Play, since individual apps (rather than Google Play) have a greater determining role in creating and managing these risks. For example, social media apps are available on Google Play, but those apps are primarily responsible for enforcing their own UGC policies. Google Play does face risks when it comes to UGC in the form of app reviews, such as efforts to influence the visibility of apps, either positively or negatively, with inauthentic reviews (known as "review bombing"). However, our efforts to address this risk, combined with the fact that users are often searching for a specific app, reduces the residual risks considerably.

Preventing Review Bombing and Ensuring Rating and Review Integrity

While Google cannot undertake content moderation for in-app content of apps available on Google Play, we work to ensure the integrity of app, game, books, and movie reviews on Google Play.

There are several ways that we work to moderate ratings and reviews. Both qualitative comments and quantitative ratings (i.e., one to five stars) are monitored, especially to detect coordinated campaigns to either artificially boost or downgrade a listing's rating. We deploy specialised algorithms to identify signals that may indicate coordinated attacks (e.g., duplicate or repeat reviews), which are then reviewed by humans. And as discussed earlier in this report, we are improving our systems' ability to detect violative content across different languages.³⁰

We also work to ensure app ratings present an accurate picture of the current user experience by calculating ratings using a percentage of the most recent reviews, not the average of all the ratings, to determine the overall rating for an app. This methodology protects against impact spikes that sham ratings can have on an app's rating and helps lead instead to more accurate ratings that reflect true user sentiment towards app experiences. We believe that the work we do to root out sham ratings leads to a more transparent app ecosystem, which ultimately supports the visibility and availability of a diversity of viewpoints and content on our platform.

Protecting Civic Discourse

We conducted a comprehensive evaluation of systemic risks associated with civic discourse. Some inherent risks relate to (1) the risk of apps engaging in misinformation and disinformation relating to elections, civic discourse, or democratic participation and (2) digital threats such as targeted account hijacking, phishing, and targeted disinformation campaigns. While significant mitigations are in place to address these inherent risks (described below), room for improvement with respect to these dynamic threats remains.

We have made significant investments in addressing civic discourse misinformation and disinformation risks through the introduction and enforcement of clear Google Play Policies.

For example, we introduced [minimum requirements](#) that apps must meet prior to being classified in the News category, including transparency requirements about the source and ownership of in-app news content, requirements applicable to news subscription services, and requirements regarding the use of affiliate marketing and ad revenue.

To protect integrity in elections, our [Deceptive Behaviour policy](#) prohibits apps from making misleading claims or providing false information about the app, including demonstrably deceptive or false content about an app's capabilities or functionality that may interfere with voting processes. For example, an app that misleads voters into believing they can cast their vote through the app would violate these policies. The policy additionally prohibits apps that promote or help create false or misleading images, video, and/or text, and requires apps that manipulate or alter media to prominently disclose or watermark the altered media.

Because apps contain layers of their own hosted content, may collect user data, and are constantly being updated and evolving, we found that app development and usage were of the most relevance for the systemic risk assessment. However, the assessment also considered other types of content, such as books on Google Play. Our [Publisher Content Policies for Google Play Books](#) are specific to book publishers and set out what books publishers can and cannot distribute to users on the Google Play platform. These

³⁰ See *supra* at p. 22 (“Evaluating Content Across Languages”) for further discussion of how Play, and Google as a whole, are addressing this identified residual risk for Play.

policies cover areas such as hate speech, child safety, misleading content, and copyright. Because of well developed and enforced policies in this area, we did not identify a significant residual systemic risk with respect to offerings on Play other than apps, such as books.

Platform Design

Protecting Privacy

Reflecting the fact that Google Play exists in the app ecosystem and offers apps in categories that are likely to involve the use of personal data (e.g., banking or government services) some of the highest inherent risks for users who access content through Google Play relate to privacy.

While we strive to maintain an open and accessible Google Play and maximise user choice, we also enforce safety standards for apps through our developer policies, ensuring we provide a more safe and secure environment for app users than would exist without Google Play. These measures are typically sufficient to lower residual risk considerably, such as risks relating to the collection and use of sensitive personal data without consent; however, the determined and constantly evolving nature of bad actors caused us to conclude that some elevated residual risk for phishing, malware, and malicious apps remains. An explanation of this elevated residual risk and related Article 35 mitigations are below.

Our developer policies create consistent safety standards for apps that appear on Google Play, and generally give users additional transparency and control over their personal data. These policies include heightened protections for [personal and sensitive user data](#), which prohibit developers from selling personal and sensitive user data, and require developers to limit the access, collection, use, and sharing of personal and sensitive user data acquired through the app to purposes reasonably expected by the user. This year, we also introduced [data deletion requirements](#), requiring developers to delete associated data when they receive an account deletion request unless the user indicates they want their data preserved or certain other exceptions apply.

We also offer a “Data safety section” for apps. The Data safety section provides developers with a transparent way to show users if and how they collect, share, and protect user data, before users install an app. Developers are required to tell us about their apps' privacy and security practices by completing a form in Google Play Console. After a developer completes and submits the Data safety form, Google Play runs automatic checks on key elements of the information provided as part of the app review process. This information is then shown on the [app's store listing on Google Play](#). We cannot wholly know what data a developer collects and shares and so compliance remains the responsibility of the developer; however, if we become aware of a discrepancy between app behaviour and this declaration, we may take appropriate action, including enforcement action. With strengthened platform protections and policies, and developer outreach and education, we prevent submitted apps from unnecessarily accessing sensitive permissions.

Apps that are deceptive, malicious, or intended to abuse or misuse any network, device, or personal data are strictly prohibited. You can read more about our approach to topics such as user data, permissions, misrepresentation, and deceptive behaviour in [Privacy, Deception and Device Abuse](#). However, motivated bad actors are constantly evolving their tactics to circumvent known protections on Google Play, so we assessed this as having some medium levels of residual risk remaining. In recognition of this challenge, Google Play plans on putting additional mitigations in place to address these risks as per Article 35 of the DSA. More specifically, we are planning to increase the number of pre-publication app bans, focusing efforts on top threat vectors, including phishing scams and malware. We are also enhancing our “Know Your Developer Program” to increase the verification of developers and decrease the number of developers causing policy violations.

Other privacy concerns relate to the use of software developer kits (SDKs). App developers often rely on third-party code, or SDKs, to integrate key functionality and services for their apps. We are clear with developers that our existing privacy and security requirements apply in the SDK context and are designed to help developers safely and securely integrate SDKs into their apps. In 2022, we launched the [Google Play SDK Index](#) to help developers evaluate an SDK’s reliability and safety and make informed decisions about whether an SDK is right for their business and their users. You can read more about our approach in [SDK Requirements](#).

We believe these policies are investments to protect our users and help developers meet consistent standards. When it becomes apparent that a developer is not meeting our established requirements for privacy and user safety, we take action to remove offending apps or developers from Google Play. We have prevented policy-violating apps that were submitted for publishing from appearing on Google Play with improved security features and policy enhancements. Google Play Commerce prevented over \$2 billion in fraudulent and abusive transactions in 2022 by banning bad accounts.

Protecting Children’s Rights

The systemic risk assessment reviewed several risks relating to child rights, and found the highest inherent risks to include the risk that children under a defined minimum age may access services that they should not be able to, that children’s data may be used for ads targeting, that apps may not function equitably for children of varied learning styles, and that apps primarily directed at children may not be of an adequate quality across languages, markets, and age groups.

However, the systemic risk assessment concluded that the combination of platform and service design measures and policies that Google Play has in place are reasonable, proportionate, and effective mitigation measures. The assessment reinforces our view that we provide a more safe and secure environment for app users than would exist without Google Play.

Several important features of Google Play address child safety risks on the platform:

Maintaining Additional Policies for Minors

We have additional requirements for apps that are targeted at children under the age of 13. Before an app is published on Google Play, the developer must certify whether children under the age of 13 are part of the target audience and, if so, the app must comply with the [Google Play Families Policies](#) (in addition to the standard Google Play Developer Program Policies). While developers generally are in the best position to

identify the correct audience for their apps, in some instances, we may disagree with a developer's stated age designations and redesignate the app.

The Google Play Families Policies establish heightened obligations for developers regarding age-appropriate content, data practices (e.g., not making use of precise location data and no personalised ads for users known to be under 18), and social app features. Apps subject to these policies must also disclose in greater detail how they use the user data they collect. A dedicated enforcement workstream that uses both automated protections and human reviewers enforces the Google Play Families Policies.

Providing a Teacher-Approved Program

Google Play makes it easy for families to find quality content for children. Google Play's Teacher Approved program is a quality review program for apps that specifically target children under the age of 13. It collects ratings from teachers, children's education specialists, and media specialists, who rate and approve apps based on a range of quality criteria (i.e., whether apps are fun and inspiring, age-appropriate, and thoughtfully designed). Approved apps are included in Google Play's Kids Tab, along with a description of their quality attributes, to help families easily review the apps and make informed choices for their children. The program provides an additional layer of review, insight, and quality control on top of the Google Play Families Policies.

Obtaining Age Assurance

Users can view Google Play on the web without being signed into a Google Account, but must sign in to download, purchase, or install content on Google Play, whether on the web or on the mobile store. Adult content is not available in a signed-out state and is blocked for signed-in users under the age of 18. As described further above, Google utilises age assurance technology, along with a neutral age-screen in the Google Account sign up process, to help determine which recipients are likely under the age of 18. Recipients identified as likely being under 18 are subject to heightened privacy, content, and safety protections. To reduce the burden on our recipients and in accordance with data minimisation principles, these processes are carried out at the Google Account level, so that the results can then be used in connection with all Google services, such as on Google Play.

As a part of age assurance during Google Account sign up, if a user is under 13 (or the minimum age in their country) then a parent, guardian, or caregiver's consent is needed to continue to sign up for or use the Google Account.

When a child reaches their country's [minimum age to manage their own Google Account](#), the child can choose to continue their current parental supervision settings or manage their own account. Family Link facilitates a range of parental controls on Google Play for [supervised Google Accounts](#), including purchase controls, approving or blocking apps, and filtering content based on content ratings.

Enforcing Content Ratings and Content Restrictions

We incorporate official content ratings from the International Age Rating Coalition (IARC) into Google Play ratings. The IARC is administered by a group of participating regional ratings agencies. IARC ratings are designed to help developers communicate locally relevant content ratings to recipients. Ratings are assigned by a regional authority based on a rating questionnaire completed by the developer and displayed in each app's listing page on Google Play. IARC ratings may be updated when developers make changes to their app's content or features that affect issues in the IARC questionnaire.

IARC ratings are used to aid parental controls and to restrict access by recipients under the age of 18 to mature-rated content where legally required. For supervised Google Accounts, parents can filter or block content based on IARC ratings (i.e., limit their child to seeing content rated PEGI 16 or below). Unrated apps are treated as high-maturity apps for the purpose of parental controls until they receive a rating.

Google Play blocks the purchase or download of mature-rated content in the EU, unless we have signals providing sufficient confidence that the recipient is an adult. In some circumstances, we require users to provide additional verification (e.g., by providing evidence of a government ID or credit card) of their age. We might require such verification if a user is trying to access mature-rated content or services, and we cannot otherwise establish with sufficient certainty that they are an adult, or if our model has classified the user as under 18 but the user wishes to verify eligibility to access such content.



Shopping

Shopping

Description of Service and Associated Risk Profile

Shopping helps users discover and learn about the products they are interested in, whether from a big-box retailer, direct-to-consumer brands, or the local store. Users use Shopping to search for products and compare prices between different merchants. They then buy products directly from the merchant on the merchant's website or at their physical store, not on Google. Our mission is to democratise e-commerce by supporting an open network of retailers and shoppers, help businesses get discovered, and give users more options when they are looking to buy.

Shopping uses a variety of factors to determine which products are displayed in search results, including a product's price, availability, and relevance to the user's query. Users can filter Shopping results by price, brand, and other criteria.

Our [Shopping Graph](#) is a dynamic, AI-enhanced, and real-time dataset of product listings, sellers, brands, reviews, product information, and inventory. Listings are updated constantly based on information retailers share directly via Google Merchant Center or from what retailers and brands post across the web. The Shopping Graph makes those sessions more helpful by sorting through a vast set of products to connect people with over 35 billion listings globally across the web. Shopping is used by around 70 million average monthly users in the EU.³¹

Online merchants use the [Merchant Center](#) to connect with customers on Shopping and either use free product listings or ads to promote their products. All ads are clearly marked as "Sponsored" or "Ad."

In addition to the content promoted by merchants, Shopping includes user-generated content in the form of product and merchant reviews and ratings. Google collects some reviews and ratings directly through [Google Customer Reviews](#), a free program that merchants enable to allow Google to collect feedback on their behalf. Shopping also features reviews and ratings collected using a merchant's own UGC service or a third party service working in a software as a service model (e.g., Yotpo, Avis Vérifiés).

You can read more in [How Shopping Works](#), [How Merchant Center Works](#), and [Shopping Graph](#).

³¹ Average monthly counts based on distinct signed-in accounts of recipients.

Systemic Risk Assessment Results and Associated Observations

We assessed 37 different risk statements³² for inherent risk (i.e., risk absent any action taken by Google), preparedness (i.e., the cumulative measures currently in place to mitigate the risk), and residual risks (i.e., risk after mitigation by Google). Residual risk serves as a guide for where further investment may be warranted. The full list of risk statements is found in Annex A to this report.

This systemic risk assessment surfaced important themes relating to the inherent and residual risk. Because Shopping operates on limited types of content that are directly related to products available on Shopping, many of the risks covered by the systemic risk assessment (such as CSAM, illegal hate speech, and election misinformation) have a lower likelihood of appearing.

However, risks relating to privacy, the freedom to conduct a business³³, consumer protection³⁴, and intellectual property³⁵ feature more prominently given the role of Shopping in presenting and raising the visibility of products sold by merchants. For these themes the systemic risk assessment identified several areas of important inherent risk that are being appropriately addressed, as seen through high preparedness evaluations, resulting in much lower levels of residual risk.

In the following two sections we consider the risks and mitigations relating to illegal and policy violating content (“Content Moderation”) and the design and functioning of Shopping (“Service Design”), though in practice there are several interactions and relationships between the two.

Taken together, these two sections address the four broad categories of systemic risks articulated in Article 34(1) of the DSA and the specific manifestations of those systemic risks that we evaluate. This report emphasizes those risks for which the assessment showed elevated inherent or residual risk, and describe Shopping’s current risk mitigation practices as well as improvements pursuant to Article 35 of the DSA.

Content Moderation

Removing Illegal Content

Identifying and Blocking Illegal Products and Services

One risk associated with Shopping is that merchants may promote or attempt to sell illegal products and services through Shopping. Here, we assessed that determined bad actors seeking to use Shopping services for the sale of illegal products or services constitute higher levels of inherent risk, but our effective mitigations result in much lower levels of residual risk. For example, at the time of writing, we have identified around 1.1 million products as live animals, 5.2 million products as illegal drugs, and 1.9 million as other illegal products that we have blocked from appearing on Shopping.

³² See Methodology Step One: Classification.

³³ Article 16 of the EU Charter: Freedom to Conduct Business.

³⁴ Article 38 of the EU Charter: Consumer Protection.

³⁵ Article 17 of the EU Charter: Right to Property.

Shopping has a robust set of policies that prohibit the sale of illegal products and services, including those relating to [gambling](#), [abuse of the network](#), [local legal requirements and safety standards](#), [dishonest behaviour](#), and [healthcare and medicines](#).

Prohibiting and Detecting Violations of Intellectual Property Rights

Shopping maintains policies that address the sale of goods that infringe on the intellectual property rights of others, such as our [counterfeit](#), [trademark](#), and [copyright](#) policies.

Shopping prohibits the sale or promotion of counterfeit products. Malicious actors may attempt to leverage Shopping to disseminate counterfeit goods, but our robust reactive and proactive enforcement scheme means Shopping is well prepared to address this risk, resulting in lowest residual risk.

Shopping uses well-established proactive detection measures for counterfeit violations, which include techniques like keyword matching and detection of signals that may indicate merchants are promoting trending products with unrealistically low prices. Additionally, trademark owners can report merchants offering counterfeit goods in a dedicated reporting channel. Where a merchant is identified as promoting counterfeit goods, its Shopping account is typically suspended.

Trademark owners can also report Shopping content that uses their trademarks in a way that is likely to cause confusion about the origin of a product. Our teams review each notice carefully, including confirming that the reporter has valid trademark rights. Where the notice is complete and we determine that the content violates our trademark policies, we remove the content from Shopping.

We provide a simple and efficient mechanism for copyright owners from countries/regions around the world. To initiate the takedown process, a copyright owner who believes content is infringing sends us a takedown notice for that allegedly infringing material. When we receive a valid takedown notice, our teams carefully review it for completeness and check for other problems. If the notice is complete and we find no other issues, we remove the content from our services.

Addressing Content that Violates our Policies

Maintaining Shopping Policies

We have two categories of policies—[Free Listings Policies](#) and [Shopping Ads Policies](#)—that outline what is and is not permitted on Shopping, including for product listings pulled from what retailers and brands post on their websites.

The Free Listings Policies and Shopping Ads Policies prohibit content that is harmful to customers or the overall shopping and advertising ecosystem.

Both sets of policies cover four broad areas:

1. Prohibited content, meaning content that is not allowed to be listed, such as counterfeit products, dangerous products, and inappropriate content;
2. Prohibited practices, meaning things merchants cannot do if they want to list products, such as misrepresentation of content;

3. Restricted content that can be listed with limitations or in certain locations only, such as adult-oriented content, alcoholic beverages, and healthcare-oriented content; and
4. Editorial and technical content, meaning website standards, such as irresponsible data collection and use.

In addition, Shopping enables users to [report listings and ads](#) that violate policies and/or contain illegal content, and enables brand and trademark owners to [report merchants misusing their brand or trademark](#). A dedicated team reviews and actions these incoming complaints.

Maintaining Guardrails for User-Contributed Content

User-contributed product and seller reviews are intended to enhance the user experience by helping users discover and select products and online sellers on the basis of opinions and feedback from other customers. We have developed [user-contributed content policies](#) and [product rating policies](#) covering content such as hateful content, misrepresentation, and fake reviews to help ensure everyone who views user-generated content has a positive experience.

An automated system processes reviews before they show up on Google to remove spammy or inappropriate language. Spammy content includes reviews with the same content posted multiple times or from multiple accounts.

After a review is published, it cannot be modified or updated by Google and we are not able to contact reviewers or ask reviewers to update what they wrote.³⁶ However, we may take down reviews that are flagged to us, in order to comply with legal obligations.

Google also enables users to report user reviews that may violate the law and to provide feedback on user reviews to improve the user experience. In the near future, and pursuant to Article 35 of the DSA, Shopping plans to introduce additional reporting functionality to allow users to report policy-violating content as well.

Preventing Unfair Commercial Practices

We strive to create a healthy digital shopping ecosystem that is trustworthy and transparent. Customers should feel confident about the offers they are browsing and the businesses they are purchasing from. Unfair commercial practices—such as scams or representing products inaccurately—pose an inherent risk to the overall ecosystem. Shopping’s policy enforcement processes significantly address this risk, resulting in low levels of residual risk.

For example, our [policy on misrepresentation](#) requires merchants to be upfront, honest, and provide users with the information that they need to make informed decisions. We disallow promotions that represent products in ways that are not accurate, realistic, and truthful. In addition, merchants are encouraged to take part in user-generated content programs to help shoppers review “real world” feedback (from Google and external sources) about product and merchant quality.

Our [abuse of the network policy](#) bans malicious content, sites that offer little unique value to users and are focused primarily on traffic generation, retailers who attempt to gain an unfair advantage in Shopping campaigns, and retailers who attempt to bypass our review processes.

³⁶ Shopping has collected some reviews from the EU via the Google Customer Reviews program, and in this case users are able to delete their own reviews.

We also maintain a list of certain kinds of businesses with products prone to abuse. This list informs prioritisation in risk management and is regularly updated based on Google reviews, feedback from users, regulators, and consumer protection authorities.

Preventing Fraudulent Business Information

We assessed the risk that disinformation, misinformation, or fraudulent content about a business, such as fake reviews, are discoverable on Shopping. This is an area with medium levels of residual risk that we have identified for enhanced mitigation going forward pursuant to Article 35 of the DSA. We are in the process of introducing several new mitigations, including reviewing and verifying merchant identity-related signals, such as VAT information and identity verification. For some time, we have been using signals such as IP location, social media presence, and third party consumer research sources to mitigate these risks.

Regarding fake reviews, Shopping has automated content checks that focus on content quality, and we employ intermittent analyses aimed at identifying anomalous review contributions. We run these checks both on an individual level (a specific merchant) and on the review source level (a review aggregator). Examples of what we might investigate further include elevated levels of 1 star or 5 star reviews or an unusual number of reviews provided by a single user or for a specific entity. We also deploy teams of trained operators and analysts who audit reviews and ratings.

Service Design

Respecting Privacy

The use of sensitive data in eCommerce (such as credit card numbers, user names, and passwords) results in critical inherent risks relating to data collection and use, and data sharing. There are two dimensions to privacy risk and mitigation on Shopping: the privacy practices of merchants, and our own privacy practices.

We do not process payments and are not involved in shipping products, so we do not collect sensitive payment data or pass it onto merchants, nor do we control the actions of merchants and retailers. However, we do set high expectations for merchant and retailer [data collection and use](#), and prohibit unsafe collection or use of personal information, and misuse of personal information. Under this policy, merchants may not collect data for unclear purposes, use personal information in ways customers have not consented to (e.g., re-selling users' contact information), or without appropriate security measures in place (e.g., not obtaining certain data over non-secure SSL server connections). We have also established [checkout requirements](#) covering aspects such as accurate pricing, user information, and language use.

Where our privacy practices are concerned, storage of signed-in user data by Google is controlled by [Web & App Activity](#) and the collection and use of data is controlled by the Google Privacy Policy and [personal results settings](#). To block specific advertisers or opt into personalised ads, users can visit [My Ad Center](#). By default, Shopping ranks product listings based on relevance to a user's current search terms.

Vetting Merchants

All products and merchants go through in-depth reviews before they can list on Shopping.³⁷ These reviews use a combination of automated and/or human evaluation to ensure compliance with our policies, with the more complex, nuanced, or severe cases often reviewed by specially trained experts. Thanks to the [Shopping Graph](#), our dataset of the world's products and sellers, our automated systems can quickly review whether a business is legitimate, whether the products shoppers see are accurate, and whether merchant content follows our policies. This automated vetting process has helped us more efficiently and accurately review a massive amount of products. In January 2023, globally, we stopped over 100 million product offers from being shown and disapproved nearly 300,000 accounts for having quality issues or not following our policies.

Sometimes we make mistakes in our decisions when vetting merchants and enforcing our policies, which may result in the unwarranted removal of products or accounts from our services. For this reason, and pursuant to Article 35 of the DSA, Shopping plans to enhance our appeals process by (1) creating an appeals path for content removed based on counterfeit complaints, (2) creating appeals paths for all content removals, and (3) enabling merchants suspected of fraudulent activities to submit their EU VAT ID as an additional data option during appeal, which increases likelihood of successful account reevaluation.

Monitoring Merchants and Listings

Our safety efforts do not stop once a product listing goes live. Our automated systems are always monitoring for violating activity, and our team of human reviewers is on standby to review issues that might need a more nuanced perspective, such as a sudden drop in prices, a significant shift in product mix, or a change in business information. After they are onboarded, we review merchants and their listings, making sure nothing has suspiciously changed since they first came to Google. We take different types of actions when we see odd behaviour, such as removing listings that violate our policies, or suspending a merchant's Shopping account. In most cases (all except sanctioned accounts), these actions can be appealed by the merchant.

Protecting Children's Rights

Shopping leverages measures applied to all Google services for age assurance for signed-in (including centralised Google Account solutions) and signed-out recipients.

In addition, Google ensures that adult and non-family safe [ads](#) on Shopping are restricted from minors and recipients for whom we do not have an inferred or declared age. When we have insufficient signals to indicate that a user is an adult, we err on the side of turning on children's protections because of the critical importance of protecting minors. In this Shopping context, this makes sure they cannot access products (e.g., adult products) which may not be safe for their age.

³⁷ Shopping does not physically inspect products. Product review is limited to a review of virtual signals that may indicate violations of our policies.

Lastly, Google has a suite of automated and manual processes aimed at scalably identifying and preventing content that depicts harm to children, such as CSAM. For product images that we get from merchants directly we use automated tools and human reviews to identify and block instances of CSAM. Due to the way in which we source user reviews (directly from merchants and from third-party aggregators), Google expects those third parties hosting the content to [moderate it before it reaches our service](#); however, we still run our own protections for images and remove and report any CSAM we find on our service.



YouTube

Description of Service and Associated Risk Profile

YouTube's mission is to give everyone a voice and show them the world. We believe that everyone deserves to have a voice, and that the world is a better place when we listen, share, and build community through our stories. From music to education, from comedy to news, YouTube touches every corner of society, offering access to information in the video format to anyone with an internet connection.

The Internet is a force for creativity, learning, and access to information, and supporting the free flow of ideas has always been and remains at the heart of YouTube's video-first mission.

YouTube allows users to watch, upload, and share videos. YouTube is available to all EU users free of charge, but users can also opt to pay for a premium subscription that removes paid ads and offers other features. Both are covered in this systemic risk assessment.

YouTube's focus on voice, stories, and community means that the service potentially impacts a wide range of rights afforded by the EU Charter, such as freedom of expression and information, media pluralism, freedom of the arts and sciences, freedom to conduct a business, as well as broad civic participation rights.

We strive to make YouTube as open as possible and empower users to easily access, create, and share information. In addition to providing a service for users to express their creativity and ideas, we are also a source of economic opportunity for creators, with whom we share revenue from ads that are served on their video content. Yet, as with all open internet services, there are inherent challenges and risks that arise which we must also address, including those from users that upload violent or dangerous content, sensitive and graphic content, and misinformation. Bad actors actively seek to exploit open services like YouTube for their own nefarious purposes, even as we continue to invest in robust systems designed to stop and deter them.

Over the years, we have worked tirelessly to develop policies and products that protect the YouTube community. As reflected in our Community Guidelines (policies broadly covering spam, deceptive practices, violent and dangerous content, misinformation, sensitive content, and regulated content) and Legal Removals processes (procedures to ensure we comply with legitimate user and government requests to remove content), YouTube is committed to keeping the service safe, for our users, advertisers, and society at large, while balancing open and free creative expression across the service. Beyond removing harmful content, we also leverage our recommendations system and monetization tools to promote a healthier ecosystem.

YouTube's prominent role as an online video-sharing service means that we naturally have a responsibility to protect the service from harmful content that may be uploaded, as well as other abuses of the service.

YouTube's business model only works when our viewers, creators, and advertisers have confidence that we are living up to our responsibility as a business. In other words, responsibility is a business imperative: users do not want to see harmful content, advertisers do not want to be associated with it, and creators and YouTube depend on each other to attract users and advertisers alike.

Systemic Risk Assessment Results and Associated Observations

We assessed 37 different risk statements for inherent risk (i.e., risk absent any action taken by YouTube), preparedness (i.e., the cumulative measures currently in place to mitigate the risk), and residual risk (i.e., risk after mitigation by YouTube). Residual risk serves as a guide for where further investment may be warranted. The full list of risk statements is found in Annex A to this report.

At a high level, the sorts of risks addressed in this report can be divided into two sets: risks posed by the presence of a particular type of illegal or policy-violating content, and risks posed to users based on the design and functioning of a service. Thought of another way, some risks are mitigated by preventing, removing, or raising visibility of certain types of content, i.e., content moderation, and others are mitigated by changing the design or functioning of the service or the way users interact with the service, i.e., service design.

Important inherent risks identified in this assessment include risks associated with the presence of illegal or potentially harmful content, which we address via content moderation. Our investments enable us to achieve much lower levels of residual risk; however, given the complexity of giving everyone a voice while addressing harmful content, some medium level residual risk remains in relation to misinformation, disinformation, civic discourse, and public health.

Other notable inherent risks are associated with the design and functioning of a service, such as privacy, security, and child rights. Below we explain the service design choices which significantly lower residual risks related to the way YouTube functions, such as privacy and security measures and protections for children using YouTube. However, while we have made significant investments in the safety of our younger users, such as preventing access to age-inappropriate content, the limitations of existing research into the existence or nature of a link between service use, the types of content being viewed, and addiction results leads us to conclude that some elevated levels of residual risk exist.

The structure of the below follows this division. We first address content moderation on YouTube, explaining YouTube's content policy development, enforcement, and the measures, like the Violative View Rate (VVR), which we use to gauge the efficacy of our moderation practices. The second section explains service design choices, which address risks related to the way YouTube functions, such as privacy risks or protections for children using YouTube.

Most of the systemic risks addressed in Article 34(1) of the DSA are related to fundamental rights, which are indivisible and interdependent. Because these rights (and associated risks) are interrelated, the practices YouTube employs to ensure users' rights frequently address more than one, or many, rights and risks articulated in Article 34(1) of the DSA. With this in mind, we have gathered together specific manifestations of systemic risks into groups that allow for efficient explanation of YouTube's existing mitigating practices, as well as improvements consistent with Article 35 of the DSA.

Content Moderation

Removing Illegal Content

YouTube is one of the world's largest open video-sharing services. It is not surprising that bad actors work to upload illegal content on YouTube in violation of our express prohibitions (such as child sexual abuse material, terrorist and violent extremist content, hate speech, and non-consensual intimate images). This is why illegal content is one of our most critical inherent risks, and the reason we have invested significantly to address the same—both alone and in collaboration with others, as described below. These investments and partnerships result in significantly lower estimates for illegal content residual risk.

As discussed previously in this report, YouTube, and Google more broadly, have a robust process for evaluating government requests to remove content.³⁸ But in the absence of an order to remove content or a valid complaint from a rightsholder (as in the case of content infringing on intellectual property rights), YouTube enforces its Community Guidelines. There is extensive overlap between content prohibited by our Community Guidelines and content that is illegal, meaning that our enforcement efforts work to mitigate the risks of both illegal and policy-violative content.

Two Examples: Terrorist or Violent Extremist Content and Child Sexual Abuse Material (CSAM)

Terrorist or violent extremist content, and CSAM are examples of the overlap between illegal and policy-violative content. Enforcement efforts in both areas make use of signal sharing and hash matching (i.e., digital fingerprinting) to identify potentially violative content. Although CSAM is always illegal, the legal status of violent and extremist content varies widely according to context (based on the jurisdiction and the way the content is presented, as in the case of a documentary).

Identifying and Removing Violent Extremist Content

Content that violates our policies against terrorist and violent extremist content includes material produced by designated terrorist organisations, content glorifying violent acts, and recruiting or fundraising on behalf of extremist groups, even if the content is not affiliated with a designated terrorist organisation. YouTube also prohibits violent or gory content intended to shock or disgust viewers, or content encouraging others to commit violent acts.

YouTube is committed to identifying and removing content that promotes terrorism or violent extremism on our service. Over the years, we have heavily invested in human review and machine learning technology that helps us quickly detect, review, and remove this content. Content that is removed is also used to train our automated classifiers for better coverage in the future. In the rare cases users do see a video they believe is violative of our policies, we provide users with the option to flag, including for videos that "promote terrorism."

³⁸ For more information on how YouTube and Google respond to government requests to remove content, see *supra* at p. 20 ("Handling Government Removal Requests").

We're also a founding member of the Global Internet Forum to Counter Terrorism (GIFCT), where we work with other tech companies to keep terrorist and violent extremist content off the web and train and provide resources to smaller companies. In 2016 we created a hash-sharing database with industry partners where we share hashes (a type of "digital fingerprint") of terrorist content to inhibit its further spread. Today, this shared database is formally operated by GIFCT, which consists of 28 [member companies](#) (and growing), and the hash-sharing database [contains](#) hashes corresponding to more than 370,000 distinct images, videos, and textual items. This industry-wide collaboration helps address the systemic risk that illegal terrorist and violent extremist content spreads across services and supports smaller companies facing similar challenges. YouTube also uses these hashes for its own detection purposes and to test pertinent policies.

Whether violent extremist content is first detected by our own classifiers, by a GIFCT hash, or by a user flag, these moderation decisions are fed back into our machine learning technology to improve future detection.

Detecting, Removing, and Reporting CSAM

Similarly, we have heavily invested in engineering resources to detect CSAM in ways that are precise and effective, and have long used this technology to prevent the distribution of known CSAM videos on YouTube. This is an area where Google as a whole has been an industry leader, and this report previously addressed other company-wide efforts to combat the distribution of child sexual abuse material.³⁹ We have always had clear policies prohibiting content on YouTube that sexualises or exploits children. We use machine learning systems to proactively detect violations of these policies and have human reviewers around the world who quickly remove violations detected by our systems or flagged by users and our priority flaggers.

While some content featuring minors may not violate our policies, we recognise that the minors could be at risk of online or offline exploitation. This is why we take an aggressive approach when enforcing these policies, including for a feature like comments (considered a minor and ancillary feature pursuant to Recital 13 of the DSA), which accounts for less than 1% of users' time spent on YouTube. Our automated systems help to proactively identify videos that may put minors at risk and apply our protections at scale, such as restricting live features, disabling comments, and limiting video recommendations.

Our proprietary [CSAI Match technology](#),⁴⁰ which we licence to several other technology companies free of charge, allows us to detect known CSAM images and videos. In cases where a video contains CSAM or a user solicits CSAM through comments or other communications, our team reports it to the National Center for Missing and Exploited Children (NCMEC), who then liaise with global law enforcement agencies such as Interpol and Europol.

Once we have identified a video as illegal and reported it to NCMEC, the content is hashed (given a "digital fingerprint") and used to detect matching content. This hashing and scanning technology is highly precise at detecting known CSAM and enables us to detect illegal content more quickly. We maintain a database of known CSAM hashes and any content that is matched against this list is removed and reported to NCMEC.

³⁹ See *supra* at p. 53 ("Detecting, Removing, and Reporting CSAM").

⁴⁰ CSAI is child sexual abuse imagery and is a subset of content that can be considered CSAM.

Prohibiting and Detecting Infringement of Intellectual Property Rights

While losses due to copyright infringement can be serious, the risk of intellectual property (IP) infringement produced only medium levels of inherent risk because of the relatively small number of users that are primarily affected, and because the related harms are not as difficult to remediate as, for example, serious physical harm or harms to vulnerable populations. Additionally, considering the suite of options protecting both institutional and individual rightsholders—such as the copyright webform, Copyright Match Tool, and Content ID, which are explained in more detail below—our assessment produced high preparedness ratings for our industry-leading protection tools.

All rightsholders have access to the YouTube copyright removal request webform, which is a streamlined and efficient way to submit copyright removal requests, and is available in 80 languages. It is designed for infrequent use by creators who hold few copyrights and rarely find their content on YouTube. For the vast majority of rightsholders, the webform is the only tool they need. Nevertheless, creators who have used the webform to remove videos from YouTube have access to powerful features, including the ability to ask YouTube to automatically prevent copies of the removed videos from being reuploaded.

For creators who experienced a higher amount of reposting of their copyrighted content and needed to submit more frequent copyright removal requests, we built the Copyright Match Tool to facilitate those creators' attempts to protect their intellectual property rights. The Copyright Match Tool is available to any YouTube user who has submitted a valid copyright removal request through the webform. Once a takedown request is approved, the Copyright Match Tool starts scanning YouTube uploads for potential matches to the videos reported in the removal request. The tool surfaces these potential matches to the claimant so they can decide what action to take next. For creators in the [YouTube Partner Program](#), the tool automatically scans for potential matches on other channels, maintains a log of those matches for review, and through an easy-to-use interface allows the creator to archive the match, submit a takedown request, or contact the user. As of December 2022, over 2.5 million channels on YouTube have access to the Copyright Match Tool.

Efforts such as the Copyright Match Tool equip creators with the resources they need to protect their content, and are evidence of YouTube's high level of preparedness to prevent the non-authorized use of copyright protected materials. You can read YouTube's bi-annual [Copyright Transparency Report](#) for a description of the other means by which YouTube protects rightsholders.

Addressing Content that Violates our Policies

Our open service embraces a wide diversity of voices to entertain, teach, showcase talents, advocate, and build businesses. YouTube's commitment to free expression enables this diversity. But free expression on an open service can create tension with other fundamental rights, such as the right to security or the right to privacy. YouTube is available in over 80 languages, with billions of monthly active users worldwide (over 400 million in the EU), and more than 500 hours of video content uploaded every minute. Because of YouTube's scale, even a relatively small number of bad actors can have systemic impacts on the service. YouTube acknowledges that its scale and extensive reach necessitate that we carefully balance the fundamental rights of users with any potential harms that may arise from misuses of our service.

Content that may technically be lawful but still creates harm (such as content impacting human dignity, promoting discriminatory beliefs, inciting, praising, or glorifying violence, promoting practices harmful to health, inciting gender-violence, or that constitutes harassment and bullying) can be uploaded to YouTube's open service, rendering it one of our most critical inherent risks. We have developed thoughtful and comprehensive approaches for addressing this type of content (described below), but several factors make these risks more challenging to address than illegal content, including the need to take proportionate and reasonable measures that respect the right to freedom of expression and information, the need to consider the likelihood of real world harm, and the need to consider context to determine whether content actually violates policy. This results in some remaining medium residual risk, such as in relation to content promoting practices harmful to health (including health misinformation), misinformation and disinformation related to civic discourse, and harassing and bullying content.

Below we describe some of YouTube's efforts to identify and respond to all policy violative content. We then address some of the specific types of content violations we examined in this risk assessment, and discuss both current and planned mitigations for these specific risks.

Developing Policy

Preventing systemic risk related to content starts with YouTube's [Community Guidelines](#). These "rules of the road" allow creative expression while prioritising the protection of the YouTube community from harmful content. As explained previously in this report,⁴¹ YouTube's policy development process is robust, involving extensive internal analysis before implementation, regular reviews and updates, and engagement with internal and third-party experts to address issues before they reach, or become widespread, on our service. You can read more in [How YouTube Works](#) and in our [blog post on policy development at YouTube](#).

It is not a coincidence that our Community Guidelines closely mirror many of the potential systemic risks addressed by Article 34 of the DSA. For years, YouTube has been attuned to these same risks. Our policies cover areas (such as hate speech, harassment, child safety, and violent extremism) across five broad categories: spam and deceptive practices; sensitive content; violent or dangerous content; regulated goods; and misinformation. Dozens of individual policies fall under these five categories, and our assessment concluded that YouTube's policies provide excellent coverage of the risks identified both in Article 34(1) of the DSA and the Recitals elaborating on those systemic risks. Some of the most relevant Community Guidelines are explained below, such as those related to election misinformation or harassment and bullying. But one can find Community Guidelines that correlate to any of the DSA systemic risks. For example, YouTube's Community Guidelines related to sensitive content, violent and dangerous content, and regulated goods provide complete coverage of the illegal content related concerns detailed in Recital 12 of the DSA.

Providing EDSA Exceptions

We recognise that some content that may otherwise violate our Community Guidelines but nonetheless provides compelling educational, documentary, scientific, or artistic value should remain available for viewers. We call this the "[EDSA](#)" (Educational, Documentary, Scientific, or Artistic) exception, and it is a critical way to make sure that important speech stays on YouTube, while protecting the wider YouTube ecosystem from harmful content. To educate creators, we include information about EDSA in [our Help Center](#). To help determine whether a video might qualify for an EDSA exception, we look at multiple factors,

⁴¹ See *supra* at p. 17 ("Designing Appropriate Content Policies").

including the video title, descriptions and the context provided in the video's audio or imagery, as well as the public interest of the content. These decisions are nuanced and context is important. Examples include hate speech that is condemned in a documentary about war, content targeting minors with insults that might appear as part of an educational anti-bullying campaign, or nudity that has scientific value or constitutes artistic expression.

Enforcing Policy

In addition to developing robust policies, we use a wide range of tools to enforce these policies. By combining multiple methods and approaches, some of which are listed and described below, YouTube continually improves the service for our viewers and creators.

Undertaking Automated Detection and Removal

Automated detection of problematic content enables YouTube to enforce our Community Guidelines at scale. Sophisticated automated systems are our primary tool. In Q1 2023, 93.7% of videos and 99.4% of comments that were removed were first detected by automated means. This increased scale and efficiency allows YouTube to keep up with the more than 500 hours of video content uploaded every minute.

And automated detection isn't just identifying huge amounts of problematic content; it's doing so quickly, before the impact is widespread. For example, in Q1 2023, 72.3% of the videos we removed had ten views or fewer.

Maintaining a Priority Flagger Program

While we facilitate and encourage flags by users, generic user flags typically have low actionability rates. . We do rely on organisations and experts in our [Priority Flagger program](#) (formerly called Trusted Flaggers) to complement our automated systems and help spot potentially problematic content. We developed the YouTube Priority Flagger program to streamline the reporting processes for government agencies and non-governmental organisations (NGOs) that are particularly effective at notifying YouTube of content that likely violates our Community Guidelines. The program provides these partners with dedicated reporting processes and a channel for ongoing discussion and feedback about YouTube's approach to various content areas. The program is part of a network of more than 300 government partners and NGOs that bring valuable expertise to our enforcement systems. Participants in the Priority Flagger program receive training in enforcing YouTube's Community Guidelines, and because their flags have a significantly higher action rate than the average user, we prioritise them for review. However the size of the program compared to YouTube's scale meant that in Q1 2023 Priority Flaggers accounted for only 0.6% of videos removed from the service.

Enforcing a Three-Strike System for Repeat Violators

YouTube recognises creators' significant investments in their video content. For that reason, we already provide creators with proportionate due process when we think it is necessary to take enforcement action against a creator or their content. We have consistent penalties for violating our policies, but provide opportunities to cure violative conduct, exemplified by our [three-strike system](#). Generally, after one Community Guidelines violation, the user gets a warning, but with subsequent violations the user begins to accrue strikes. Strikes carry increasing penalties when a channel receives them within a 90-day period:

- 1st strike - 1 week suspension;
- 2nd strike - 2 week suspension; and
- 3rd strike - channel termination

We developed our three-strikes policy to balance terminating bad actors who repeatedly violate our Community Guidelines with the need to make sure people have an opportunity to learn our policies and appeal decisions. At the same time, we work hard to make these policies as understandable and transparent as possible, and we enforce them consistently across YouTube. In general, we know that about 98% of users never break our Community Guidelines. And 94% of people who receive a first strike never get a second one. We do not hesitate to issue strikes and terminate channels whose content repeatedly violate our policies, irrespective of whether the channel has a large audience.

While legitimate users get three strikes, we directly terminate egregious offenders such as uploaders of CSAM or channels dedicated to posting spam. Every quarter between Q3 2022 and Q1 2023, we have terminated more than 5M channels for spam.

Strikes, terminations, and content removals are only a few pieces of a larger puzzle. These complex problems necessitate multifaceted solutions, and dealing with material such as misinformation or potentially sensitive content on YouTube is no exception. While our Violative View Rate (described in the following section) shows that YouTube has made strides in removing clearly violative material, more nuanced harms are not solely addressed by removals under our Community Guidelines. These areas of harmful content, which often brush up against our policy lines, require a comprehensive approach that includes raising authoritative content and rewarding creators who meet the higher bar required for our partner program.


Measuring Success: Violative View Rate

As described above, automated classifiers allow for the quick detection of problematic content. YouTube strives to remove content that violates our Community Guidelines before users are exposed to it. In Q1 2023, we removed 40.4% of violative videos before they had a single view, and 31.9% of videos when they had one to ten views.

To measure our progress on removing violative videos before they are viewed, we developed a metric called Violative View Rate (VVR), which has been publicly available since 2021. This metric, updated and made publicly available quarterly, estimates the percentage of total views on YouTube that are of violative videos (i.e., videos that are inconsistent with our Community Guidelines).

VVR data gives critical insight into how well we are protecting our community. Although metrics like the turnaround time to remove a violative video or the number of takedowns are important, those statistics do

not fully capture the actual impact of violative content on viewers. The VVR is a better measure because it tells us how widely violative videos have been disseminated before they are taken down. Two videos could be removed from YouTube within 24 hours, but one may have 100 views while the other has 1 million views. This is a 100% takedown rate within 24 hours, but that metric obscures the most important information. Because we care most about the potential for harm to users, and potential harm can arise by actual exposure to violative content, we have chosen to focus attention on a metric that specifically measures user exposure. We believe the VVR is the best way for us to understand the extent to which harmful content may reach viewers, and to identify where we need to make improvements. We are committed to being transparent about this metric and working to continue to reduce it over time, as we have since 2017.



Graphical depiction of the YouTube Violative View Rate (VVR)

Calculating VVR serves a second purpose: it helps us gain insight into the type of content we should remove but sometimes miss. Our methodology for calculating the metric allows us to do this. We calculate VVR by taking a sample of videos on YouTube and having content reviewers gauge which videos violate our policies and which do not. By sampling, we gain a more comprehensive view of the violative content that evades our detection and enforcement systems. With that understanding we can improve those systems and, over time, further decrease the VVR.

Over the years, we have seen the VVR fluctuate—both up and down. For example, immediately after we update a policy, this number may temporarily rise as our systems ramp up to catch content that is newly classified as violative. Our methodology for this reporting mechanism [has been validated](#) by MIT Sloan professor of statistics Dr. Arnold Barnett as “thoroughly sensible and statistically sound.”⁴²

Our VVR reports indicate that violative views today are around 0.1% of all videos viewed (i.e., out of every 1,000 views on YouTube, just one is of violative content). We recognise that even if the prevalence of violative content is low, it might still represent a large volume of content in absolute terms, and significant

⁴² Arnold Barnett (2021) [YouTube’s Violative View Rate Methodology](#), Massachusetts Institute of Technology.

investments are required to maintain these low levels. This report describes the efforts we undertake to prevent users from seeing violative content, and identifies potential areas for improvement.

Elevating Authoritative Sources

Removal of violative content is not the only way that YouTube makes adjustments to balance the freedom of expression with other rights such as safety and security. Over the past several years we have invested significantly in the systems that take authoritativeness of the channel into account when making recommendations. Our systems are trained to elevate authoritative sources higher in search results, particularly in sensitive contexts. We raise high-quality information from authoritative sources in search results, recommendations, and information panels, in turn helping people find accurate and useful information. Whether you are searching for something evergreen or a current event, YouTube aims to surface videos from sources like public health authorities, research institutions, and news outlets, within the top search results. These efforts are particularly important when it comes to connecting people with information from high-quality sources at the moments that matter most—for example when learning about a breaking news event, or searching for health information. So when a user in France searches for, "arreter de fumer," the top listed videos are from top hospitals, cancer centres, and Sante Publique France, the national public health authority.

Providing Information Panels with Topical Context

For YouTube search queries or videos related to topics prone to misinformation, such as COVID-19 and climate change, we surface information panels which provide content sourced from independent third parties. These panels give viewers additional context and help them make more informed decisions about what they are watching. Depending on the topic, the panels will point to information from sources like health authorities, Wikipedia, Encyclopedia Britannica, and the United Nations. These information panels will show regardless of what opinions or perspectives are expressed in a video. YouTube also uses information panels to inform users when content has been uploaded by a news outlet funded in whole or in part by a government.

Addressing Specific Content Risks

Addressing Misinformation and Disinformation

Risks related to intentional manipulation of the service (e.g., dis/misinformation impacting civic discourse or such content promoting practices harmful to health) are complex, constantly evolving, and societal-wide issues. They necessitate a multifaceted approach, combining policy enforcement and content evaluation with real-time context and information for users. YouTube has many measures in place (described below) but must still contend with determined and highly-motivated bad actors constantly evolving their techniques, resulting in some medium levels of residual risk.

YouTube has developed measures to respond to situations where the risk of misinformation is at its greatest. Where misinformation violates our policies we are quick to remove such content, and we will terminate a channel for egregious or repeated offences. But we employ other techniques to combat misinformation in addition to simply taking down content or channels. First, we have a higher bar for monetised content, so that creators are not incentivised to create untrustworthy clickbait or other low-quality, misleading videos. Another way we prevent the spread of misinformation is to raise up authoritative content. For example, the "Breaking News Shelf" is available in 42 countries, including 16 EU member states, and appears on the YouTube homepage automatically when there is a significant news

event unfolding in a specific country. We know that bad actors take advantage of fast-breaking situations and will post unreliable or false information to capture users attention and serve their own malign purposes. Content appearing on the Breaking News Shelf is from authoritative new sources to counteract this trend. Additionally, during these events, YouTube's recommender systems emphasize authority in results. We also demonetise channels that do not meet heightened standards related to the reliability of their content or make misrepresentations about who they are or the purpose of their channel. In egregious cases, we terminate channels that make misrepresentations like this.

Addressing Public Health Related Violative Content

In our assessment, we identified critical inherent risk in relation to content promoting practices harmful to health (e.g., self-harm, anorexia, health misinformation) and assessed our preparedness at "effective," in part due to the viral nature of this content.

Our methods for addressing public health-related misinformation exemplify the multifaceted approach described above. We raise information from authoritative health sources, as determined by external experts and externally published [principles](#), and provide context on the sources of health information for users via information panels below each of those videos. For example, we recently launched an [updated approach to eating disorder-related content](#) that was informed by third-party experts and seeks to create space for community, recovery, and resources, while continuing to protect viewers. This new approach involves an expansion of our Community Guidelines (e.g., prohibiting content about eating disorders that feature imitable behaviour, or behaviour that we worked with experts to determine can lead at-risk viewers to imitate), age-restricting certain videos, and surfacing crisis resource panels with videos discussing eating disorders.

Crisis resource panels are an important part of the suite of health products on YouTube. These are information panels that help easily connect users with authoritative and helpful information in times of crises. YouTube's crisis resource panels allow users to connect with live support from recognised crisis service partners. The panels may surface on the Watch page, or in YouTube search results. Currently, topics covered include suicide, self-harm, eating disorders, and topics related to certain health crises or emotional distress.

Another recent policy change with inherent challenges and trade-offs concerns medical misinformation, with scientific understanding evolving all the time and important topics (such as vaccines) being a source of fierce debate, notwithstanding consistent guidance from health authorities about their effectiveness. Our Community Guidelines already prohibited certain types of medical misinformation, but we worked with experts to expand them, introducing [new guidelines](#) on currently administered vaccines that are approved and confirmed to be safe and effective by local health authorities and the World Health Organization.

Addressing Civic-Discourse-Related Violative Content

Article 34(1) of the DSA directs YouTube to conduct an assessment of systemic risks to civic discourse and electoral processes. In the assessment, YouTube evaluated the inherent risk of civics misinformation to be higher, leaving elevated levels of residual risk despite our preparedness, largely due to the dynamic and viral nature of misinformation in politics, during elections, and at times of crisis and civic unrest.

With users around the world coming to YouTube to learn about politics and develop informed opinions about current events, we have a responsibility to support an informed citizenry and foster healthy political discourse. We provide a range of resources for civics partners such as government officials, candidates, civics organisations, and political Creators to ensure a broad range of voices are heard.

Among other items, our Community Guidelines prohibit content that has been technically manipulated or doctored in a way that misleads users and may pose a serious risk of egregious harm, content that aims to mislead people about voting processes, and content encouraging others to interfere with democratic processes, such as obstructing or interrupting voting procedures. Policies that are relevant during elections include:

- **Voter suppression:** Content aiming to mislead voters about the time, place, means, or eligibility requirements for voting, or false claims that could materially discourage voting.
- **Candidate eligibility:** Content that advances false claims related to the technical eligibility requirements for current political candidates and sitting elected government officials to serve in office. Eligibility requirements considered are based on applicable national law, and include age, citizenship, or vital status.
- **Incitement to interfere with democratic processes:** Content encouraging others to interfere with democratic processes. This includes obstructing or interrupting voting procedures.
- **Impersonation:** Content intended to impersonate a person or channel.
- **Deceptive practices:** Spam, scams, or other deceptive practices that take advantage of the YouTube community.
- **Harassment & cyberbullying policies:** Content that threatens individuals, including content that incites others to harass or threaten individuals on or off YouTube.

In addition to our robust policies about what is not allowed on YouTube, we also devote significant resources to systems that raise the visibility of authoritative content, as described above.⁴³ These techniques are designed to ensure that users find the trustworthy content they are looking for on topics that can be targets for manipulation by bad actors.

⁴³ See *supra* at p. 103.

Detecting and Removing Harassment and Bullying in YouTube Comments

Experience also shows that comments on YouTube are sometimes misused to directly and indirectly threaten the wellbeing of creators and other users particularly at risk of being the targets of harassment and abuse. YouTube's automated moderation systems are specifically and proportionately designed to mitigate these risks to user and creator safety. With these protections, we removed nearly half a billion comments in 2022 for violating our Community Guidelines prohibiting harassment and cyberbullying. YouTube's automated detection systems are removing this content at scale, but the overall number of bullying and harassing comments produced one of the higher inherent risk ratings in our assessment, so we believe there is more work to be done to protect our users.

For a minor and ancillary feature like YouTube comments, which account for less than 1% of the time users spend on the YouTube service, much of the violative content is directed at creators (i.e., users that upload videos), who are the heart of YouTube. Consistent with Recital 40, YouTube has developed appropriate and proportionate strategies for detecting and removing offensive content in comments aimed at creators, including those creators particularly at risk of being subject to hate speech, sexual harassment, or other discriminatory actions. YouTube deliberately casts a wide net, using automated technologies optimised to identify and remove any comments appearing under videos directed at creators at particular risk of being subject to harassment, discriminatory actions, or bullying.

In Q3 2022, about 62% of actioned comments were removed because they were spam (i.e., deceptive, high-volume commercial content that harms the user experience). The remaining approximate 38% were removed for other important user safety reasons: 15% were removed for harassment and cyberbullying; 15% for child safety; and 7% were removed for hateful and abusive content (including content that targets vulnerable populations).

While our size influences our risk profile, so does the format of the content with which users engage across our service. Users come to YouTube to create, share, and view audiovisual content, i.e., videos. Comments are a secondary feature, which creators can opt to enable and permit users to contribute additional textual feedback under creator videos. Our Community Guidelines apply to all content on the service, regardless of its format. But when it comes to how those policies are enforced and the corresponding consequences for users who post secondary text content in comments, there are critical differences as compared to video content. In other words, comments occupy a fundamentally different place in YouTube's video-first ecosystem.

European users spend less than 1% of their time on YouTube engaging with comment functionality (as of Q4 2022). Viewed another way, both globally and in the EU, users spent over 120x more time watching videos than they did engaging with comments. On an average day in Q4 2022, fewer than 2% of daily active users posted a comment, in the EU as well as globally.

Moreover, a user's investment in commenting on a video does not compare to a creator's investment in terms of time, effort, and resources in creating the video content available on YouTube. Creators often take many steps to create a YouTube video: research, scripting, filming, editing, audio-mixing, thumbnail creation, and search engine optimisation. They use multiple devices and software applications and can spend many hours of production work creating a single video, costing time and money. By contrast, commenting requires very little effort: a few keystrokes amounting to much less time and effort than video creation.

There are other aspects of YouTube comments that make them unlike videos:

- Creators have control over whether comments on their videos are enabled or not, can remove any comments under their videos for any reason, can edit comments under their videos, can create block lists for words or phrases permissible in comments under their videos, and can block specific users from commenting on their videos.
- Comments are not searchable, recommended,⁴⁴ nor accessible via the YouTube Homepage.
- Comments are not a factor in a creator's ability to monetise their video content.
- Comments are intrinsically tied to the video to which they relate, and not independent pieces of hosted content. If a video is removed or taken down, the comments associated with it are automatically taken down. The same is not true when a comment is moderated, which has no impact on the availability of the video on the service.
- Comments are not enabled on all versions or interfaces of YouTube. Additionally, comments are not available on certain types of videos featuring minors, YouTube Kids, or for embedded videos.

Given the above considerations, the moderation of comments on YouTube does not pose the same risks to freedom of expression or information present in video moderation.

Prohibiting and Removing Hate Speech

Our assessment found lower levels of residual risk. YouTube's hate speech policy outlines clear guidelines prohibiting content that promotes violence or hatred against individuals or groups based on certain attributes. We enforce this policy rigorously and regularly report on the removal of hateful content from our service. For example, in Q1 2023, we removed over 177,000 videos for violating our hate speech policies.

We have made significant progress in our work to quickly remove hateful content from our service. In 2019, we updated our hate speech policy, resulting in an increase of the number of daily hate speech comment removals by 46x, and with a 5x spike in the number of hate videos removed.

A 2020 [report by the Institute of Strategic Dialogue](#) showcased the efficacy of our hate speech policy update: "Following YouTube's change of hate speech policies we found a significant reduction of such content on the platform... an analysis of the volume of these mentions over time reveals a dramatic drop in content around spring 2019, demonstrating the effectiveness of YouTube's ban on Holocaust denial content."⁴⁵

Additionally, all our policies, including our hate and harassment policies, include penalties for creators who repeatedly brush up against the line, including [removal from the YouTube Partner Program](#).

⁴⁴ Comments are not recommended by an algorithm to increase engagement, but they can be sorted chronologically or by most engagement depending on settings.

⁴⁵ Jakob Guhl, Jacob Davey (2020), [Hosting the 'Holohoax'](#), Institute for Strategic Dialogue.

As outlined on our help center page, under YouTube’s hate speech policy, we may remove content or issue other penalties—such as terminating an account—when a creator repeatedly targets, insults and abuses a group based on attributes such as race, ethnicity, sexual orientation, or gender identity and expression, across multiple uploads.

Additionally, YouTube is a founding signatory to the EU Code of Conduct on Countering Illegal Hate Speech. Each year, YouTube takes part in the annual monitoring exercise, responding to flags from NGOs specialising in hate speech. Results from the 7th monitoring in 2022 show that YouTube reviewed over 80% of flags in 24 hours and removed over 90% of the content flagged. [As the report shows](#), YouTube was the only signatory to remove more illegal hate speech content in 2022 than in 2021.

Service Design

Respecting Privacy

YouTube’s main source of revenue is advertising—a portion of which is shared with creators participating in the YouTube Partner Program—and we use the information we collect for the purposes described in our Privacy Policy, including to provide the service, customise services, provide recommendations, personalise search results, and serve relevant ads. We also take our responsibility to protect user information seriously, and while advertising makes YouTube free of charge for everyone, we do not sell personal information to anyone. YouTube’s data practices turn a significant inherent risk into a much lower residual risk.

Our [Privacy Policy](#) and [YouTube’s Help page on privacy](#) provide transparency over what information we collect, why we collect it, how we process it, and how users can manage their information. [Your data in YouTube](#) is a powerful, easy-to-use tool designed to give users control over the privacy settings that are right for them, and provides further information on the data we collect and use across our services.

Now users have more control over the collection and use of watch history data. As of August 2023, if a user turns YouTube watch history off and has no significant prior watch history, features that require watch history to provide video recommendations will be disabled.

Protecting Children’s Rights

In this assessment and consistent with our obligations under DSA Article 28 regarding “Online Protection of Minors”, YouTube considered numerous risks particular to children. As described below, these include the risk that children access or are exposed to content they should not see, or conversely that their access to content is overly restricted; the risk that YouTube stimulates behavioural addictions in children; and the risk that children’s data are used to target ads.

YouTube is heavily invested in the safety of its younger users. As described above, we have well developed and advanced tools to quickly detect and remove illegal content. We collaborate with industry partners, and make available first-in-class tools to allow other services to remove illegal content at scale as we do. Our policies provide additional protection, under which we remove harmful but legal content. Below we describe the policies and protections that go beyond content removal, and ensure that our service is designed in a way that is aimed at keeping children safe.

YouTube pursues many policies and programs to protect children on the service, and we seek the input of experts to shape those efforts. [Our Youth and Families Advisory Committee](#) is made up of experts in children’s media, child development, digital learning, and citizenship from a range of academic, non-profit and clinical backgrounds, and provides advice when we update our family product experiences and policies. Other components include rules and guidelines for when children appear in content, restricting access to mature content, and protecting minors at risk. You can read more in [Fostering Child Safety](#).

Maintaining Guardrails for Children’s Access to Content

We assessed the risk that children under a defined minimum age access YouTube services that they should not be able to or are exposed to harmful, hateful, or age-inappropriate content. Without action by YouTube, children could readily have access to content that is age-inappropriate; however, our wide range of measures (such as minimum age requirements, signals for estimating the age of users, “Made for Kids” content, parental controls, and granular age categories in YouTube Kids) significantly address this risk. Although this risk will never be eliminated—children with access to the internet will seek to view content available on an open service—our measures result in a significantly lower residual risk profile. The mitigations described below also resulted in a much lower residual risk of children’s access to content being over or under restricted.

We are always looking at ways to create an appropriate environment for family content on YouTube, so we invest heavily in the policies, technology, and teams that help provide families with the best protection possible. Our holistic child rights approach has several important components.

We age-restrict content that does not violate our policies, but is nonetheless inappropriate for viewers under 18. This includes videos containing adults participating in dangerous activities that children may imitate or videos related to regulated substances, sexually suggestive content, or violent and vulgar content. Videos that are age-restricted are not viewable by signed-out users either.

[YouTube Kids](#) is a separate app built from the ground up to be a safer and simpler experience for kids to explore, with tools for parents and caregivers to guide their journey. The app is a filtered version of YouTube and has a much smaller set of content available than YouTube’s main app and website. This is because we work to identify content that is age-appropriate, adheres to our quality principles, and is diverse enough to meet the varied interests of kids globally.

[Supervised experience on YouTube](#) is for parents who decide their tween or teen is ready to access YouTube through a supervised Google Account. Videos a child can watch depend on the content setting their parent selects when setting up a supervised experience. We have disabled a number of standard features normally available in YouTube, like comments, uploads, purchases, and live chat. To reinforce healthy screen time habits, reminders for breaks and bedtime are set to “on” by default and YouTube’s autoplay feature is set to “off” by default.

When we find that a user is under 13 and using the service unsupervised, we terminate that account. YouTube employs classifiers to find signals on YouTube channels that indicate that the channel may be owned by a user under the age of 13. These classifiers rely on content signals to find such channels, which are then flagged for a team to review more closely when they appear owned by an underage user. Users identified as potentially underage are sent through the account recovery process and are given two weeks to provide evidence that they are 13 years or older or to obtain parental consent and establish parental supervision. Accounts flagged as potentially underage are disabled and thereafter deleted if the user

does not prove they are 13 years of age or older or establish parental supervision within two weeks of the initial notice.

YouTube also employs a classifier to determine whether young minors are livestreaming themselves without supervision. These accounts are further reviewed by a team to determine whether to disable the account.

Addressing Potentially Addictive Behaviour in Children

We also assessed the risk that the interface, design, or features of YouTube stimulate behavioural addictions in children using the service. Google has many measures in place to address this risk (such as parental controls, the unique experiences designed for kids described above, and surfacing high quality content), but the general lack and limitations of existing research into the existence or nature of a link between service use (e.g., screen time), the types of content being viewed, and addiction results in some elevated levels of residual risk.

Despite inconclusive research, we take steps to tamp down excessive use of our service by children. For users that declare themselves to be under 18 when they create their Google Account, autoplay on YouTube is turned “off” by default. The autoplay setting on YouTube Kids and YouTube Supervised Experience is also turned “off” by default and the parent has the ability to [disable](#) autoplay so that their child cannot change this control. This step is aimed at encouraging more active choices by recipients about how they want to spend their time online. [Take A Break](#) and [Bedtime](#) reminders are turned “on” by default. These are aimed at reinforcing healthy screen time habits.

Additional protections apply for YouTube Kids and YouTube Supervised Experience. YouTube Kids allows parents to set the amount of time their child can spend on the service. Additionally, Family Link accounts allow parents to control the time children spend with their device or with specific apps, including YouTube Kids and YouTube Supervised Experience.

Protecting Children’s Data

[Personalised ads are prohibited on YouTube Kids](#), as well as for users in a supervised experience on YouTube, consistent with our obligations under Article 28(2) of the DSA. This means the ads that appear are matched to videos being watched based on the content, not the specific user watching. For videos 'Made for Kids', we limit data collection and use, and as a result, we restrict or disable some service features. For example, we do not serve personalised ads on content 'Made for Kids', and some features are not available on these videos, like comments and notifications. All Creators are required to indicate whether or not their content is 'Made for Kids'. Accordingly, our assessment resulted in lower residual risk of children’s data being used to target ads.

Protecting Children’s Safety in YouTube Comments

For years, YouTube has been attuned to the pernicious threat of predatory conduct towards children in comments. To combat this threat, we have continually refined classifiers that automatically detect and remove potentially predatory comments. It is critical that classifiers related to child safety be designed to cast a wide net because they must detect comments that are often facially innocuous. In egregious cases, we terminate the account and report the content to the National Center for Missing and Exploited Children (NCMEC), an organisation that works with global law enforcement agencies to protect children. We have developed and launched increasingly effective classifiers in order to ensure that YouTube remains a safe

space. For a minor and ancillary feature like comments specifically, these automated classifiers are deliberately designed to cast a wide net so as to identify and remove as much material in comments as possible that may potentially be harmful to children (e.g., sexualisation of minors, information regarding minors, CSAM). We also use machine learning classifiers to identify hundreds of millions of non-violative videos depicting children and automatically turn off comments to avoid any chance of the child being the subject of harassing or predatory comments. We choose to err on the side of safety to protect this vulnerable population from exploitation, and continuously work to refine our automated approaches for identifying and removing any content in comments that potentially threaten the safety of children.

Although less than 1% of the time users spend on YouTube is spent using comments functionality, YouTube commits resources and has developed reasonable and proportionate enforcement strategies optimised to detect and remove as many comments threatening the safety of children as possible, consistent with the express obligations under DSA Article 28. As noted above, this content is the most challenging for any service hosting user-generated content. Much of this content appears innocuous to many viewers but may still be used in ways that YouTube wants to prevent (for example, by individuals seeking sexual gratification). Because the content may be posted innocently and omit objectively problematic content, the challenges of addressing potential misuse are significant. YouTube has developed machine-learning tools and content policies to identify this and similar types of content that may appear innocuous or humorous, but may put minors in potentially risky situations.

As an example, in early 2019, YouTube learned that some innocuous videos (such as a home video of a young girl jumping into a pool) could potentially appeal to bad actors. These videos do not sexualise or endanger minors, and thus do not violate YouTube's content policies. However, bad actors could present some risk of engagement by viewing or commenting on the video. Accordingly, YouTube developed a comprehensive approach to address these issues: combining machine-learning tools and content policies to remove violative comments and apply restrictive measures to the discoverability of this type of content.

Where an egregious violation occurs, we terminate the commenter's account. The poster of an egregious comment removed for child safety reasons receives a notice of the termination and is given the right to appeal. In the unlikely event termination was the result of a false positive, the account is reinstated and the poster's fundamental rights are protected. The vast majority of removed comments, however, do not result in the suspension or termination of a user's account, and the uploader of an actioned comment remains free to use the service, similarly protecting the users' fundamental expression rights.

We believe our tailored approaches to moderating comments and videos strike the appropriate and proportionate balance. Data, comparatively minor user engagement, and the structure of our service place comments in an ancillary position to videos. But comments pose an outsized risk of harm to creators and young users. We choose to err on the side of safety above other considerations in this narrow context because the weight of the competing interests clearly demands it.

Promoting Equity

In 2020, we established a dedicated Racial Justice, Equity, and Product inclusion [team](#) to explore practices, policies, and norms that could reproduce bias and inequity on YouTube. We recognised the need to address safety concerns from historically underrepresented creators (above and beyond YouTube's [Community Guidelines](#)) and have since launched new features to more easily moderate comments that may be personally offensive. Such features include an optional setting that allows creators to increase strictness for comments and hold them for review in YouTube Studio, the ability to filter comments that may be more hurtful in a separate section of the creators' review tab, and the rollout of Channel Guidelines which allow creators to communicate what is and is not acceptable in their comments section. Our [Creator Safety Center](#) offers creators information on how to navigate strategies for dealing with safety concerns such as bullying.

One of the core programs under the Racial Justice, Equity, and Product inclusion team is YouTube's [Inclusion Working Group](#) (IWG). This group works to institutionalise inclusion and equity across YouTube's products, content policies, and business - prioritising equity considerations prior to product launch. Members include executive sponsors, a dedicated product inclusion lead, and representatives from employee resource groups across YouTube. Since its inception in 2020, the IWG has partnered in over 900 projects to understand and consider equity and inclusion early in the development process. The IWG's work has improved how we detect racially hateful comments and prepared teams to identify and respond to new forms of online hate.

5. Conclusions

Our mission is to organise the world's information and make it universally accessible and useful. But for information to be helpful, it must also be reliable. That's why we take our responsibility seriously to provide access to trustworthy content, deliver reliable information, and partner to create a safer internet.

Fulfilling this responsibility is essential to achieving our founders' commitment to developing services that significantly improve the lives of as many people as possible, and reflects our belief in the potential of technology to have a profoundly positive impact.

This responsibility is global, and finds expression in our global [AI Principles](#), [Human Rights commitment](#), and content policies for every Google service. It is in this context—our global mission, responsibility, commitments, and policies—that we submit this first EU DSA systemic risk assessment report.

We have long developed and implemented methodologies to assess our services and features prior to launch and throughout their use. Our DSA risk assessment builds upon our existing and well-established global approaches to risk assessment and mitigation, while identifying areas of additional work to comply with specific requirements of the Act.

We have already begun planning for future annual systemic risk assessments, and expect the following:

- Further embedding the DSA systemic risk assessment process into our broader risk assessment frameworks and systems, including finding synergies with regulatory requirements arising in other jurisdictions and our pre-existing human rights due diligence work.
- Maintaining the timeliness of systemic risk assessment by addressing new technologies, such as generative AI and synthetic media, and the opportunities and challenges they present.
- Increasing alignment between the systemic risk assessment process and our engagements with independent experts, civil society organisations, and other stakeholders.
- Further testing of mitigation measures with the users of Google services and those with insight into the interests of users and communities impacted by mitigation measures.

The world is entering a phase of significant innovation in technology regulation. In this context we encourage policymakers to design and deploy regulatory approaches that are aligned with existing international standards, frameworks, and best practices for risk assessment and management. The UN Guiding Principles on Business and Human Rights, the OECD Guidelines for Multinational Enterprises on Responsible Business Conduct, and the various multi-company and multi-stakeholder initiatives referred to in this report provide an essential foundation for principled, global, and scalable approaches that can be tailored for each service and context.

We welcome the opportunity to discuss our analysis further with the European Commission and other stakeholders, and to build upon this foundation in subsequent reports.

Annex A: Full List of Risk Statements

Illegal Content

- Risk that Child Sexual Abuse Material and other illegal content relating to child sexual exploitation is available on an online platform or search engine
- Risk that illegal terrorist and violent extremist content is available on an online platform or search engine
- Risk that illegal hate speech is available on an online platform or search engine
- Risk that intellectual property (e.g., copyright protected material, patents, trademarks) is available on an online platform or search engine in ways that violate legal protections (e.g., counterfeit products)
- Risk that an online platform or search engine is used for illegal online activity (e.g., sharing of unlawful non-consensual private images or unlawful online stalking)
- Risk that illegal products and services (e.g., illegal, dangerous, non-compliant, and banned products and services; illegal sale of live animals; illegal offers of accommodation; illicit drugs) are promoted or available for sale on an online platform or search engine

Fundamental Rights

Freedom of Expression and Information

- Risk that an online platform or search engine removes content that does not constitute a necessary or proportionate removal of content with a legitimate purpose
- Risk that users are not able to report potentially violating content on an online platform or search engine
- Risk that users are not able to appeal content removals on an online platform or search engine
- Risk that the users' ability to make autonomous and informed decisions about what they view on an online platform or search engine is impaired by limited transparency or options

Pluralism in the Media

- Risk that the visibility of content on an online platform or search engine adversely impacts the plurality of voices, opinions, and perspectives in the media
- Risk that the visibility of content on an online platform or search engine promotes a polarisation and/or siloed segmentation of voices, opinions, and perspectives in the media
- Risk that the visibility of content on an online platform or search engine disfavors content from populations that have historically been underrepresented in the media

Privacy and Data Protection

- Risk that an online platform or search engine collects, processes, aggregates, and/or shares more user information than is necessary for a stated purpose or without the informed consent of users
- Risk that private or highly personal information of users is unintentionally made available on an online platform or search engine
- Risk that sensitive personal data are used to target paid speech at users of an online platform or search engine without the informed consent of the user
- Risk that private or highly personal information about users or others is maliciously made available on an online platform or search engine
- Risk that content or applications enabling phishing, malware, data breaches, or other digital threats is available on an online platform or search engine

Human Dignity

- Risk that degrading, harmful, discriminatory, or exploitative content impacts human dignity or the physical and emotional wellbeing of some users of an online platform or search engine

Consumer Protection

- Risk that unfair commercial practices take place on an online platform or search engine

Child Rights

- Risk that children under a defined minimum age access online platform or search engine services that they should not be able to and/or are exposed to harmful, hateful, or age-inappropriate content or conduct
- Risk that children's access to and/or use of an online platform or search engine is limited more than is necessary or proportionate for a legitimate purpose
- Risk that the interface, design, or features of an online platform or search engine stimulate behavioural addictions in children using the online platform or search engine
- Risk that children's data are used by an online platform or search engine for ads targeting in ways that have adverse impacts on children's rights, including their right to be protected from economic exploitation
- Risk that applications on an online platform or search engine do not function equitably for children with varied learning styles, learning challenges, or disabilities.
- Risk that applications primarily directed at or predominantly used by children on an online platform or search engine are not of adequate quality across languages, markets, and age groups and have adverse impacts on children

Equality and Non-Discrimination

- Risk that content that has a negative impact on human dignity or promotes discriminatory beliefs and values or harmful stereotypes is available on an online platform or search engine
- Risk that online platforms or search engines select organic content or paid speech based on factors that result in discrimination
- Risk that applications on an online platform or search engine are not of adequate quality across languages, markets, and age groups
- Risk that some populations are under-represented as content contributors on online platforms and search engines, with adverse impacts on minority businesses
- Risk that algorithms on an online platform or search engine are less well trained in some languages, dialects, and vernaculars than others

- Risk that applications on an online platform or search engine do not function equitably for users with disabilities

Freedom to Conduct a Business

- Risk that disinformation, misinformation, or fraudulent content about a business (e.g., fake reviews) is available on an online platform or search engine

Civic Discourse

Civic Discourse and Elections

- Risk that misinformation and disinformation relating to elections, civic discourse, democratic participation, or civil unrest are available on an online platform or search engine
- Risk that digital threats such as account hijackings, phishing attempts, or disinformation campaigns are targeted at users of on an online platform or search engine during election times and other important civic discourse milestones

Public Security

- Risk that content with value as evidence in legal process and access to remedy is removed and/or deleted by an online platform or search engine
- Risk that content inciting, praising, or glorifying violence or content that is legal but harmful, dangerous, or hateful is available on an online platform or search engine

Public Health

Public Health

- Risk that content that promotes practices harmful to health (e.g., self-harm, anorexia, health misinformation) is available on an online platform or search engine

Gender-based Violence

- Risk that targeted gender-based harassment, bullying, or prejudice or content inciting, praising, or glorifying gender-based violence (including sexual, physical, mental and economic harm), and threats of violence, coercion, and manipulation appear on an online platform or search engine

Physical and Mental Wellbeing

- Risk that content targeting individuals with prolonged or malicious insults based on intrinsic attributes (such as protected group status or physical traits) or that constitutes harassment and bullying is available on an online platform or search engine

Annex B: List of Mitigations

Background

This annex contains specific mitigation measures being put in place pursuant to Article 35(1) of the DSA. As part of the systemic risk assessment for each VLOP and VLOSE, we evaluated our existing mitigation measures for each risk statement. As explained in this report, we have long invested in efforts to address user safety and have thus already put in place an extensive array of mitigations. These existing mitigations are discussed in the report, where relevant to the reporting of the results. Please see below for a list of new or enhanced mitigations being put in place pursuant to Article 35(1) to address the salient residual systemic risks identified in the Article 34 assessment.

Article 35 Mitigation Types

Mitigation Type	Full Article 35 Mitigation Description
Adapting the design, features or functioning of services	Adapting the design, features or functioning of their services, including their online interfaces
Adapting terms and conditions and its enforcement	Adapting their terms and conditions and their enforcement
Adapting content moderation processes	Adapting content moderation processes, including the speed and quality of processing notices related to specific types of illegal content and, where appropriate, the expeditious removal of, or the disabling of access to, the content notified, in particular in respect of illegal hate speech or cyber violence, as well as adapting any relevant decision making processes and dedicated resources for content moderation
Testing and adapting algorithmic systems	Testing and adapting their algorithmic systems, including their recommender systems
Adapting advertising systems and adopting targeted measures	Adapting their advertising systems and adopting targeted measures aimed at limiting or adjusting the presentation of advertisements in association with the service they provide
Reinforcing internal processes, resources, testing, documentation and supervision	Reinforcing the internal processes, resources, testing, documentation, or supervision of any of their activities in particular as regards detection of systemic risk

Initiating or adjusting cooperation with trusted flaggers	Initiating or adjusting cooperation with trusted flaggers in accordance with Article 22 and the implementation of the decisions of out-of-court dispute settlement bodies pursuant to Article 21
Initiating or adjusting cooperation with other online platform providers	Initiating or adjusting cooperation with other providers of online platforms or of online search engines through the codes of conduct and the crisis protocols referred to in Articles 45 and 48 respectively
Taking awareness-raising measures	Taking awareness-raising measures and adapting their online interface in order to give recipients of the service more information
Taking targeted measures to protect the rights of the child	Taking targeted measures to protect the rights of the child, including age verification and parental control tools, tools aimed at helping minors signal abuse or obtain support, as appropriate
Ensuring that information is distinguishable through prominent markings	Ensuring that an item of information, whether it constitutes a generated or manipulated image, audio or video that appreciably resembles existing persons, objects, places or other entities or events and falsely appears to a person to be authentic or truthful is distinguishable through prominent markings when presented on their online interfaces, and, in addition, providing an easy to use functionality which enables recipients of the service to indicate such information

Mitigations Applicable to Multiple Services

DSA Article 35 Mitigation Type	Mitigation	Description
Testing and adapting algorithmic systems	Incorporating signals out of recent phishing incidents	We incorporate signals from incidents to continuously improve machine learning models, internal human review guidelines, and investigation methods. This includes suspending bad actors from Google services.
	Updating manual reviews and machine learning models to stay ahead of new abuse behaviours and patterns	We regularly update machine learning models to flag phishing ads and accounts at their source, learning from the latest methods adversarial actors use to circumvent systems. While systems are constantly improving, attackers swiftly shift tactics in an attempt to game the systems. This is an adversarial space, so as the systems learn about new fraud patterns, they can better detect and action ads and accounts. Additionally, we continuously update internal human review guidelines based on new abuse behaviours and patterns.
Adapting advertising systems and adopting targeted measures	Expanded scope of business verification	We require advertiser verification in multiple key verticals (e.g. elections, financial services) and look to maximise verification generally. Currently, the vast majority of ad impressions in the EU (and globally) are from verified advertisers. We are further investing in scaling advertiser verification.
Reinforcing internal processes, resources, testing, documentation and supervision	Algorithms and language action function	By improving the translation capabilities that support content moderation, we are able to improve the accuracy of our moderation systems. Google Translate will be investing in improving general translation quality between English and German, French, Italian, Portuguese, Dutch, Polish, Turkish, Arabic, Russian, and Spanish. As we improve our translation technology with these languages, we will be bringing these improvements to more European languages.
	Increasing due diligence around personal information	We are continuing to improve access restrictions regarding personal and serving data to improve data security.

Google Maps

DSA Article 35 Mitigation Type	Mitigation	Description
Adapting the design, features or functioning of services	Language expansion for content moderation	We will continue to focus on expanding the breadth of languages supported in content moderation with the EU official languages as priority.
Adapting terms and conditions and its enforcement	Improvements to information policies	We are evaluating potential updates to our personal information policies related to user-generated content.
Adapting content moderation processes	New notifications and appeals channel	We launched new notifications and improved existing ones for content removal and created new appeal paths for users to challenge removal and access restriction decisions, including removal of user reviews. We will monitor and improve the appeal processes.

Google Play

DSA Article 35 Mitigation Type	Mitigation	Description
Adapting content moderation processes	Enhancements to pre-publication enforcement actions	We are expanding existing capabilities to block top threat vectors from entry to the Play Store to prevent user harm, including threat vectors specific to phishing and/or malware.
Reinforcing internal processes, resources, testing, documentation and supervision	Know Your Developer Program	We are enhancing our “Know Your Developer Program” to increase verification of developers and decrease the number of developers causing policy violations. These enhancements include implementing an Enhanced Due Diligence program including developer interviews, questionnaires, and testing requirements that will raise the bar for entry to the Play Store.
	Malicious app removals	We are expanding the application of existing capabilities to more quickly identify targeting predatory loan apps and apps using deceptive functionality (e.g., cleaner apps).

Google Search

DSA Article 35 Mitigation Type	Mitigation	Description
Adapting the design, features or functioning of services	Enhanced election protections	Election protections are aimed at preventing inaccurate or misleading information from the open web surfacing in Search features. We plan to continue to enhance protections that have proven effective so that we are ready for overlapping priority elections next year.
Adapting terms and conditions and their enforcement	NCII removal policy & tools	We recently launched a new removal policy for private explicit images as well as a new reporting flow within our Help Center to make it easier for users to report personal content and request removal from web results. We will be launching protections that demote sites that have a high volume of non-consensual explicit imagery removal requests approved, as well as improvements to Images, Web & Video modes to protect users from involuntary synthetic pornographic content.
Adapting content moderation processes	Improved offensive content contextualization	We are continually improving our systems' ability to identify and reduce contextually offensive content in Search. This includes content that reinforces harmful stereotypes.
	Personal hardship coping support improvements	We are working on product improvements to features like Related Questions in Search that intend to help users to cope with personal hardship and distress by promoting helpful resources (e.g., grieving, financial troubles, severe health conditions etc.).
	Reduction in oversexualization and suggestive imagery	In previous years we announced that we applied BERT to decrease oversexualization and suggestive imagery for queries about women and protected groups by 30%. We will continue to invest in the ongoing reduction of oversexualization in image based web results.
	Updating content reviewers to take action against hateful content	We continually update our training materials for the Trust and Safety teams to offer authoritative, global sources to ensure we can identify, detect, and protect against hateful content.
Testing and adapting algorithmic systems	Updating health misinformation	For Health topics, we place a particular emphasis on expertise and trustworthiness, so we surface information from authoritative sources. Our policies for

	machine learning classifiers	our Search features prohibit content that contradicts or runs contrary to scientific or medical consensus and evidence-based best practices, as well as content that promotes harmful health claims . We update our ranking systems and machine learning classifiers regularly to account for changing real-world circumstances.
Taking awareness-raising measures	Increased voter transparency	We will launch features that raise awareness of authoritative election information. This includes information about voter registration, voting processes, and election dates. We will continue to provide factual information about election candidates through Knowledge Panels.
Taking targeted measures to protect the rights of the child	CSAM reporting and detection tools	We are streamlining the notice-and-takedown process to help simplify the process for third parties like NCMEC and Internet Watch Foundation to report child sexual abuse and exploitation material.
		We are planning to launch new technical measures to detect age-indeterminate content, i.e., content where the age is ambiguous and could be potentially confused with CSAM.

Shopping

DSA Article 35 Mitigation Type	Mitigation	Description
Adapting terms and conditions and its enforcement	Enhanced appeals processes for businesses	We plan to enable merchants suspected of fraudulent activities to submit their EU VAT ID as an additional data option during appeal, which increases likelihood of successful account reevaluation.

YouTube

DSA Article 35 Mitigation Type	Mitigation	Description
Adapting the design, features or functioning of services	Adapting current safeguards for behavioural addictions in children	We will continue to improve our current safeguards by tailoring Take a Break and Bedtime reminders to different age groups. Additionally, by default, "autoplay" will continue to be turned off for supervised accounts and in the YouTube Kids App, and parental settings include a timer feature that can set a limit for the time spent in the YouTube Kids app.
Adapting terms and conditions and its enforcement	Expanded medical misinformation policies	We have policies that prohibit misleading or deceptive content with serious risk of egregious harm, including medical misinformation. These policies are constantly evolving, as real-world circumstances change and in response to changes to global or local health authorities' guidance.
	Improvements targeting adversarial abuse	We will make continued improvements targeting adversarial abuse.
	Updating hate speech, harassment, and cyberbullying policies	We will continually evaluate the need to update our Hate Speech and Harassment and Cyberbullying policies to address emerging threats.
	Updating misinformation policies	We will continually evaluate the need to update our Misinformation policies to address emerging threats.
Testing and adapting algorithmic systems	Medical misinformation classifier development	We use a combination of people and machine learning to detect potentially problematic content, including medical misinformation. Our machine learning classifiers are updated regularly to account for changing real-world circumstances.
Reinforcing internal processes, resources, testing, documentation and supervision	Updates to warnings and strikes system to include user education	Beginning in Q3, creators will have the option of taking a specialised training course when they receive a Community Guidelines warning. Completion of the course will remove the warning from a creator's channel—provided they don't violate the same policy for 90 days. We believe this update will help the vast majority of creators achieve their goal of making content in accordance with our policies, and will further

		help ensure that YouTube remains a safe and responsible platform for everyone.
Taking awareness-raising measures	Educational media literacy campaign	In November 2022, we launched our ‘Hit Pause’ media literacy campaign, and as of June 2023, the campaign is live in all EEA Member States. We have plans to launch another media literacy campaign in H2 2023, and will continue to explore opportunities to develop additional campaigns.

Annex C: List of Consultations

Recital 90 of the DSA sets out the expectation that VLOSEs and VLOPs engage with external stakeholders when undertaking risk assessments and designing mitigation measures, such as representatives of the recipients of the service, representatives of groups potentially impacted by their services, independent experts, and civil society organisations .

Prior Stakeholder Engagement

We already undertake significant engagement with external stakeholders to run the business, such as to inform decision-making, conduct human rights due diligence, and design mitigation measures. For this reason, relevant prior engagements were a primary source of external stakeholder input into this assessment and accompanying mitigation measures. These engagements include:

- Google’s Government Affairs and Public Policy and Google’s Human Rights Program engagements with government officials, law enforcement agencies, independent experts, and civil society organisations for reasons of due diligence, decision-making, and strategy. This includes groups working on child safety, privacy, freedom of expression, non-discrimination, civic engagement, hate speech, violent extremism, mis- and disinformation, and gender-based violence.
- Engagements through the Google Safety Engineering Center led by Google Trust and Safety teams with policymakers, researchers, and regulators with an interest in our content policy and its enforcement.
- Foundational research led by the Google Trust and Safety’s Research organisation on ecosystem-wide technology risks and policy issues, including information generated via direct user feedback and insights. The team’s portfolio includes over 300 reports (averaging 60 per year) that encompass time series studies, deep investigations, and secondary research for rapid decision-making.
- Engagements that content policy teams in VLOSEs and VLOPs have with civil society organisations, academics, and relevant third party experts to inform the review, development, and enforcement of content policy and get ahead of emerging issues.
- User engagements facilitated by marketing functions and specific product teams to test service features or understand user sentiments about Google and its services.
- Discussions with YouTube’s Youth and Families Advisory Committee, a collection of independent experts that provide advice on the policies and services YouTube offers to young people and families.

For example, our child safety efforts are enhanced by our active membership of several coalitions, such as the [Tech Coalition](#), the [WeProtect Global Alliance](#), [INHOPE](#), and the [Fair Play Alliance](#), that bring companies and NGOs together to develop solutions that disrupt the exchange of CSAM online and prevent the sexual exploitation of children. Together, we fund child safety research and share tools and knowledge, such as our insights into transparency reporting, in-product detection, and operational processes.

Our approach to risk assessment and design of mitigation measures is also informed by our participation in relevant multi-stakeholder and multi-company efforts, such as the [Digital Trust and Safety Partnership](#) (DTSP), the [Global Internet Forum to Counter Terrorism](#) (GIFCT), the [Global Network Initiative](#) (GNI), [Partnership on AI](#), and the [Tech Coalition](#).

Our counterterrorism and violent extremism efforts are informed by independent experts and civil society organisations via our participation in various GIFCT workstreams. [GIFCT's Working Groups](#) bring together individuals and organisations from diverse stakeholder groups, geographies, and disciplines to offer advice on critical themes related to countering terrorism and violent extremism online, while GIFCT's research arm regularly publishes insights and reports. GIFCT is advised by an Independent Advisory Committee made up of representatives from civil society, government, and intergovernmental organisations, and in 2021 published an [independent human rights impact assessment](#).

The recent momentum behind large-scale machine-learning models has sparked additional dialogue around the social impacts of generative AI and surfaced concerns as diverse as misinformation, privacy, security, and safety. Here our approach to promoting trustworthy information (such as watermarking, metadata, and other innovative techniques) is informed by our participation in the Partnership on AI's [synthetic media working group](#).

Much of our multi-company and company-specific stakeholder engagement takes place in confidential or Chatham House Rule settings. This approach supports trusted relationships, fosters candid feedback, and protects the wellbeing of stakeholders, but can also restrict our ability to share the names of the organisations we consult. Designing stakeholder engagements to facilitate public disclosure about participants may impact the type of organisations we engage and the quality and nature of actionable feedback we receive.

Stakeholder Engagement to Inform the Systemic Risk Assessment

We actively participated in [two multi-stakeholder convenings](#) hosted by the Global Network Initiative and Digital Trust and Safety Partnership to discuss both methodological questions, such as the definition of systemic risk and key features of risk assessment methodology, and substantive issues, such as fundamental rights, illegal content, civic discourse, and gender-based violence.

These convenings informed and validated key systemic risk assessment design decisions, such as how we assessed inherent and residual risk, the topic areas covered by our risk statements, and how we utilised approaches based on well-established human rights assessment methodology. These discussions provided very helpful insight into the expectations that independent experts and civil society organisations have for the methodology and output of systemic risk assessments.

More information on the convenings, including a summary of the key points raised during the discussion, is found in the formal convening outputs:

- [Agenda](#)
- [Discussion Summary](#)
- [Additional Reference Material](#)

Participants attended under the Chatham House Rule so none of the comments or observations in the discussion summary can or should be attributed to any participant. However, the full list of participants is published, including the following non-company participants: Access Now; AWO; Brainbox Institute; BSR; Center for Democracy and Technology; Center for Security and Emerging Technology; Centre on Regulation in Europe; CERRE; CIPESA; Danish Institute for Human Rights; Digital Trust and Safety Partnership; ECNL; Global Forum for Media Development; Global Media Registry; Global Network Initiative; Global Partners Digital; Graz University of Technology and Complexity; GW Law School; Information Society Law Center; Institute for Strategic Dialogue; Integrity Institute; Internet Freedom Foundation; InternetLab; Internews; ITS; Justitia; Micova; National Law University Delhi; Office of the UN High Commissioner for Human Rights; Paradigm Initiative; Reset Tech; Science Hub Vienna; Stiftung Neue Verantwortung; Taraaz; UCLA ITLP; University of Amsterdam; University of East Anglia; University of Geneva; University of Namur; Women of Uganda Network; and WZB Berlin Center for Social Science.