

Australian Code of Practice on Disinformation and Misinformation
Redbubble Inc.
Annual Transparency Report
1 January 2025 - 31 December 2025

Summary

Redbubble is a global artist marketplace dedicated to giving independent artists a meaningful way to sell their creations. Redbubble hosts user-generated content uploaded by artists and provides them with online tools to upload their art and to design and sell products printed with their art to their customers worldwide.

Redbubble recognises that the spread of misinformation and disinformation can lead to real-world harm and is committed to preventing, detecting and removing such information from its platform. In doing so, Redbubble strives to balance the fostering of artistic freedom with the goal of preventing the spread of harmful disinformation and misinformation through its marketplace. Redbubble policies prohibit misleading or false information that harms or significantly threatens public health and safety, or where the intent is to cause fear and suspicion about a topic that can cause real-world harm.

By deploying global, scalable content moderation technologies like duplicate detection, image matching, keyword and text detection, machine learning, and artificial intelligence, Redbubble is well positioned to efficiently tackle problematic content on its platform.

This report provides transparency into the measures that Redbubble takes to prevent, detect and remove harmful disinformation and misinformation on its marketplace.

Notable Trends

During 2025, conspiracy theory-related content accounted for the majority of moderation activity throughout the year, driven primarily by ongoing uploads related to Qanon and 9/11. Moderation volumes were highest in Q1 and declined across subsequent quarters, with a notable reduction from Q2 to Q3, reflecting both the continued effectiveness of proactive detection tools and routine updates to internal processes that resulted in the removal of older listings. These updates were driven by efforts to continuously improve policies, processes and procedures to ensure they adequately address evolving user behaviour and current events.

Uploads of medical-related dis-/misinformation remained at consistent levels throughout the year as seen in *Figure 1A*.

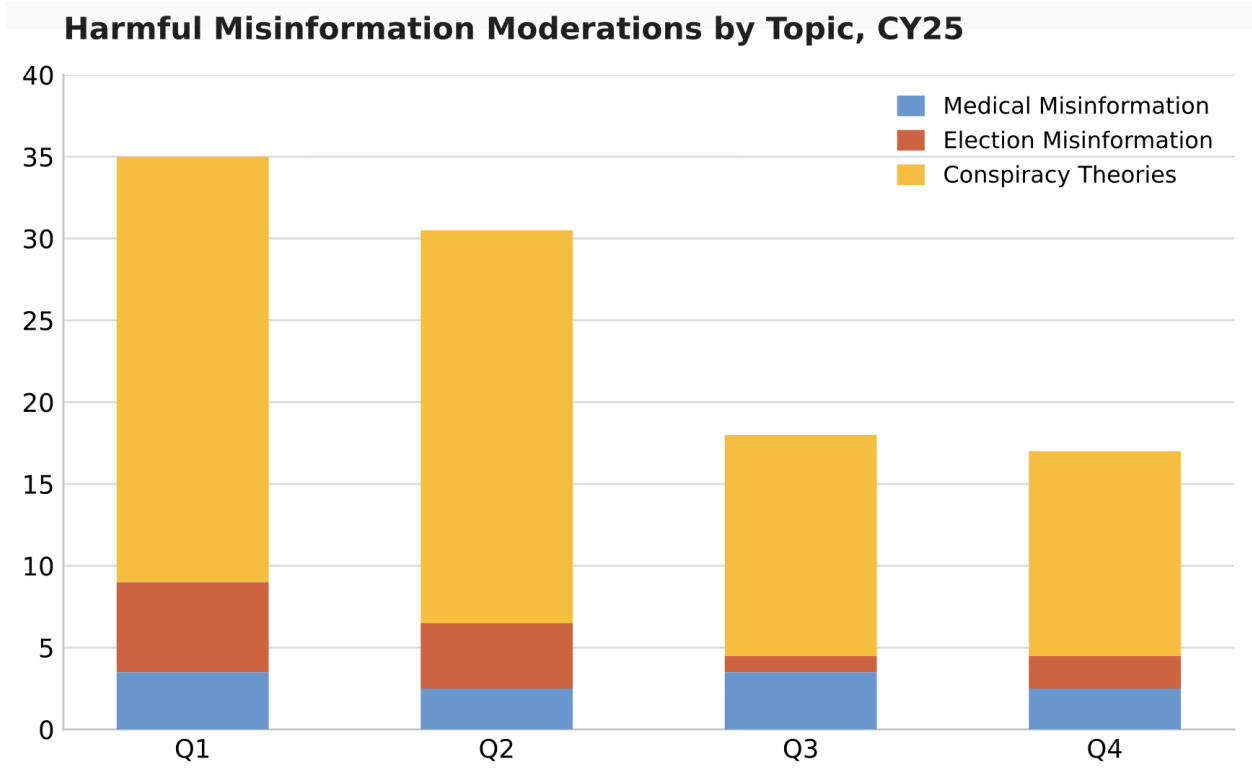


Figure 1A: Shows the total amount of content uploaded by Australian users that was moderated for mis-/disinformation in 2025.

The sustained volume of 9/11 and Qanon-related content is consistent with broader global trends. After appearing to fade from the spotlight in recent years, Qanon content experienced a resurgence in 2025, fueled in part by the return of Donald Trump to the US presidency and ongoing public interest in the Epstein files. Similarly, 9/11-related conspiracy content persisted across all quarters, a pattern consistent with the enduring nature of that topic online.

In Australia specifically, the 2025 federal election is likely to have contributed to the broader misinformation environment during the year. Elections have consistently been associated with increased online misinformation activity, and the 2025 federal election would have been no exception. Despite this, election misinformation moderation on Redbubble remained low throughout 2025.

More broadly, research from the Australian Communications and Media Authority (“ACMA”) found that misinformation related to conspiracy theories, health and medical topics declined in Australia from 2024 to 2025, which is consistent with the downward trend in overall moderation volumes seen across the year.

Harmful Misinformation Moderations, YoY




Figure 1B: shows the total amount of content uploaded by Australian users that was moderated for mis-/disinformation over 2021- 2025.

After reaching a low point in 2023, harmful misinformation moderations have risen steadily but remained relatively flat, reaching approximately 100 in 2025 — the highest level since 2022. This uptick is likely driven by a resurgence in conspiracy theory content, particularly QAnon and 9/11-related uploads, as well as ongoing improvements to detection and removal tools that have increased the platform’s ability to identify and remove violating content.

Case Study: QAnon Content Moderation in 2025

QAnon-related content was among the most frequently moderated categories on Redbubble in 2025. A typical example involved product listings using text-based slogans or imagery associated with QAnon narratives, detected through Redbubble's moderation tools. Upon detection, Redbubble reviewed the listings against applicable policy and removed those found to violate the platform's harmful misinformation policies. In several instances, the accounts responsible were found to be part of coordinated upload networks and were subject to account-level action, including permanent termination of networks of connected accounts.

Redbubble's Commitments Under the Code

Outcome 1a: Reducing harm by adopting scalable measures	Opt In
Outcome 1b: Inform users about what content is targeted	Opt In
Outcome 1c: Users can easily report offending content	Opt In
Outcome 1d: Information about reported content available	Opt In
Outcome 1e: Information about recommender engines	Opt In
Objective 2: Disrupt advertising and monetisation incentives for disinformation	Opt In
Objective 3: Work to ensure the integrity and security of services and products delivered by digital platforms	Opt In
Objective 4: Empower consumers to make better informed choices of digital content.	Opt In
Objective 5: Improve public awareness of the source of political advertising carried on digital platforms	Opt Out
Objective 6: Strengthen public understanding of Disinformation and Misinformation through support of strategic research	Opt In
Objective 7: Signatories will publicise the measures they take to combat Disinformation	Opt In

Reporting Against 2025 Commitments

Outcome 1a: Reducing harm by adopting scalable measures

Redbubble prohibits users from uploading harmful disinformation and misinformation to the marketplace in the Redbubble User Agreement, published at <https://www.redbubble.com/agreement>, and the Redbubble Community and Content Guidelines at: <https://help.redbubble.com/hc/en-us/articles/202270929>.

The vast majority of user-generated images and text uploaded to the Redbubble platform must pass through one or more scalable detection technologies. Continuous improvements to tooling allow Redbubble to more accurately and efficiently detect and remove harmful misinformation and disinformation from its marketplace. Current scalable technologies include:

- Duplicate detection identifies previously moderated images that users may try to re-upload
- Image matching detects content similar to images known for disinformation or misinformation
- Keyword detection in text-based user-generated fields, like titles, tags, and descriptions to catch keywords linked to disinformation and misinformation
- Text-in-image matching spots text-based misinformation and disinformation within images
- Machine learning pinpoints user accounts linked to networks known for violating Redbubble policies
- Artificial intelligence recognises users' keyword tagging patterns associated with disinformation and misinformation

The Redbubble Content Safety Team proactively screens the marketplace for potential misinformation and disinformation detected by Redbubble's tools, and the team removes content that it determines violates our User Agreement or Community Guidelines.

When the Content Safety Team makes decisions relating to content that potentially includes misinformation or disinformation, the team considers Redbubble policies, past decisions (to ensure consistency of approach and decision-making) and further research particularly in relation to new or emerging topics. The Content Safety Team's framework for content review is built on clear criteria and a repeatable and scalable workflow. This allows the team to make moderation decisions in an unbiased and consistent manner. Furthermore, policies and decision-making frameworks undergo continuous review and refinement to ensure their ongoing effectiveness.

Redbubble proactively monitors and screens for over 350 different content safety topics on its platform, spanning issues like incitement of violence, racism, disinformation and misinformation.

The following are examples of misinformation and disinformation topics that Redbubble screens for:

- Medical misinformation, such as anti-vaccine propaganda that may encourage the spread of communicable disease;
- Denials of real-world catastrophes, such as the Holocaust; and
- Political misinformation, such as false political conspiracy theories that are linked to real-world harm.

Redbubble strives to make use of unbiased news sources in determining whether user-generated content violates policies around harmful dis-/misinformation, including review of reports from a range of global health authorities that are based on peer-reviewed scientific research and from independent, non-partisan fact checking organizations.

To create proactive screening guidelines, the team uses the above sources to compile training content for content review teams and quality assurance of moderation decisions.

The Content Safety team recognizes the importance of keeping content policies up to date. Regularly reviewing existing policies is essential, particularly for topics that are no longer prominent or have decreased in public relevance due to global changes (e.g., the COVID-19 pandemic).

Dedicated Content Review and Moderation Capabilities

Redbubble has dedicated resources based in the United States, Europe and Australia to detect and moderate content available to Australian end-users. These resources are responsible for implementing content safety policies around disinformation and misinformation.

When Redbubble gains knowledge of user-generated content that violates its policies, examples of action include, but are not limited to:

- Moderation of listing
- Temporary suspension of account
- Account level warnings
- Restriction of account for self-review
- Account deletion

Reach of Content

Redbubble strives to limit the impact of harmful content by promptly and accurately identifying and removing user-generated content that violates its policies. Redbubble actively works to limit the number of Australian end-users who may be reached by harmful mis-/disinformation. Measures highlighted throughout this report have helped to contribute to fewer Australian end users being reached by violative content. The reach of harmful misinformation on the platform is quantified by “impressions,” which measures user engagement with violating content before it was moderated. Number of impressions is how

many times a listing appeared to a user in Australia on search results pages or in recommended carousels. "Clicks" indicate how many times an Australian user clicked on a listing.

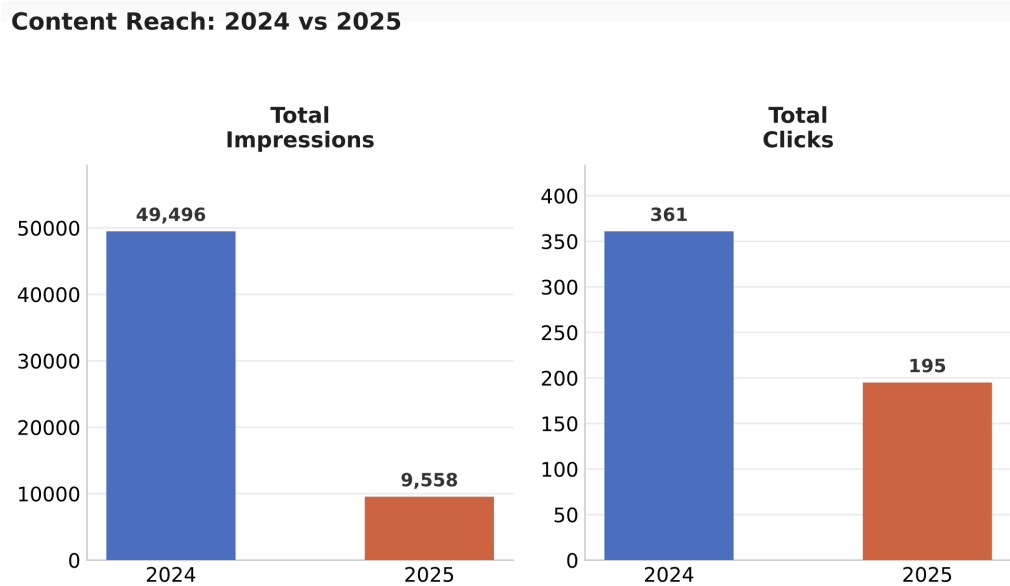


Figure 2: Shows the total number of impressions and clicks by Australian users recorded on content moderated for harmful misinformation in 2024 and 2025.

In 2025, overall reach of misinformation dropped sharply as shown in Figure 2. Total impressions fell from 49,496 in 2024 to 9,558 in 2025 — an 81% decrease. Works moderated dropped from 832 to 201 (76% fewer), and average time on site fell from 776 days to 139 days. “Works moderated” refers to the total number of individual product listings removed from the platform for violating Redbubble’s harmful misinformation policies, which differs from the quarterly moderation counts shown in Figures 1A and 1B. This decline is attributable to a combination of factors: continued improvements to proactive detection tooling, which enabled faster identification of violating content; routine policy and process updates that resulted in the removal of older listings during the year; and a reduction in the volume of certain high-impression topics compared to prior years. The drop in average time on site — from 776 days to 139 days — suggests violating content is being detected and removed before it can accumulate substantial reach.

Outcome 1b: Inform users about what content is targeted

The communication to users of what constitutes misinformation and disinformation is important in stopping its spread and impact. Australian end-users are informed of this

topic and the types of behaviours and content that are prohibited.

Redbubble does this by publishing content rules in its Community Guidelines (URL provided above) and other Help Centre articles. In the Redbubble User Agreement (URL provided above), users are required to adhere to the Community Guidelines and represent and warrant that the products they sell are free from harmful misinformation and disinformation. The User Agreement states that Redbubble reserves the right to remove any such content from the website and terminate user accounts.

The Community Guidelines detail Redbubble's rules regarding user behavior and content on its marketplace and are made public in the spirit of open communication with artists and their customers. The relevant prohibition under the Community Guidelines addresses misinformation and disinformation in the following terms:

“Harmful misinformation is not permitted. We define this as any misleading or false information that harms or significantly threatens public health and safety, or where the intent is to cause fear and suspicion about a topic that can cause real-world harm.”

The Community Guidelines are expressly noted to be adaptable and subject to refinement over time as the environment and circumstances change.

The Redbubble User Agreement provides further information to artists proposing to upload content to the marketplace. It provides that an artist uploading content represents and warrants that:

“The content does not contain material that is harmful, abusive, inflammatory or otherwise objectionable; and the content is not misleading and deceptive and does not offer or disseminate fraudulent schemes or promotions.”

Additional information for users is included in the Redbubble Help Centre's Policies & Guidelines page located at:

<https://help.redbubble.com/hc/en-us/categories/9681032172564-Policies-Guidelines>

The Content Safety team adds to this FAQ on a regular basis to provide helpful information in response to commonly asked questions regarding disinformation and misinformation and other content safety topics.

Active Australian end-users of Redbubble are informed when action is taken against their account or content they publish on the service is in violation of these policies that prohibit harmful dis-/misinformation. Users receive notification via e-mail for moderated listings and account-level actions.

Outcome 1c: Users can easily report offending content

Redbubble requests that all users of our marketplace flag behavior or content that contravenes the Community Guidelines through one of the reporting functions on our site. The User Agreement prompts users to report such content:

“Please help us by letting us know straight away about any inappropriate Content you see on the Marketplace. You can do this by clicking the "Inappropriate Content" link displayed on each Product listing page.”

Redbubble provides a simple and accessible reporting tool for users. On every product listing page created by sellers on the marketplace, a prominent “Report Content” link is provided.

This link directs users to a web form (*Figure 3*) where they can quickly and easily report content that they believe falls outside of Redbubble’s policies, including works that may contain disinformation or misinformation.

Inappropriate Content

Note: Your information will *not* be relayed to the author, you will remain anonymous.

If you'd like to report [redacted] by [redacted] as inappropriate, please complete this form and we'll make sure the review team takes a look.

Why are you reporting this content as inappropriate?
Please Select...

If the content does not meet these guidelines, please provide any additional comments or information and click "send".

SEND

How does this work?

When you report a concern a notification is sent to the Redbubble objections team. We review the content and follow up in cases where the content falls outside Redbubble's guidelines. Due to the volume of emails the team receives, we cannot respond to every query regarding these reports but please rest assured we do check every single report carefully and we'll be in touch if we need any further information. If you would like more insight into the guidelines we apply, further information is available in our [community guidelines](#). Thanks again!

Figure 3

Redbubble investigates all user reports to assess whether the content or uploading account should be subject to moderation or other actions. This reporting tool continues to be an important way to detect content that violates Redbubble policies, including disinformation and misinformation.

Outcome 1d: Information about reported content available

As outlined above, active Australian end-users have access to general information about Redbubble’s actions in response to reports about the service.

In 2025, Users submitted 140 reports using the reporting functionality described in Outcome 1c, which includes but is not limited to reports of misinformation or disinformation. This is highlighted in *Figure 4*.

Community Guideline Violation Reports, YoY




Figure 4: Shows total number of user reports from Australian users for potentially violating content over 2022-2025.

The reporting tool used by Australian users does not require reporters to specify a category of violation, meaning the 140 reports represent all community guideline concerns submitted during the period and cannot be disaggregated by type, including for misinformation or disinformation specifically.

Outcome 1e: Information about recommender engines

Product recommendations to users of the Redbubble marketplace are primarily based on keyword-matching algorithms that connect the text-based title, tags and description generated by the artists with user search and navigation behavior, such as prior keywords the user searched for and keyword similarities between content they clicked on and additional content that may interest them.

In addition, information related to the inclusion of user-generated content in feeds for off-site marketing platforms are available to users via the Redbubble Help Center, such as:

- [What is off-site marketing?](#)
- [How does off-site marketing of your products work?](#)
- [How do my off-site ads get created?](#)
- [Can I change my off-site marketing preferences?](#)

Artists may opt in or out of advertising their products offsite, and keyword blocklists are aimed at preventing products with titles or tags associated with harmful mis-/disinformation from appearing on offsite platforms.

Objective 2: Disrupt advertising and monetisation incentives for disinformation.

Redbubble reduces monetisation incentives from artist sales of products on the marketplace by swiftly detecting and removing accounts and content that violates Redbubble policies. These measures are described in more detail throughout this report. If a user violates the Redbubble User Agreement or Community Guidelines, their violative content will be removed and their accounts will be subject to account penalties, up to and including account termination. Networks of connected accounts will also be terminated.

Redbubble also disrupts incentives by using keyword blocking tools that prevent content tagged with terms related to misinformation and disinformation from appearing on offsite marketing platforms where artists promote their products and generate sales.

Redbubble's User Agreement (<https://www.redbubble.com/agreement>) and Community and Content Guidelines (<https://help.redbubble.com/hc/en-us/articles/202270929>), referenced in full under Outcome 1a above, set out the specific prohibitions on content that promotes disinformation or misinformation and the consequences for violation.

Objective 3: Work to ensure the integrity and security of services and products delivered by digital platforms.

At the account-level, Redbubble uses account abuse detection software that combines machine learning with a network of global data to detect users who have violated Redbubble policies. This software uses data points that are customised to the Redbubble marketplace and allow room for adjustments based on trends and current events. This is an effective tool for maintaining content integrity on Redbubble and detecting users who upload content that perpetuates disinformation or misinformation. For example, this tool has helped detect users who use bots to create networks of multiple accounts and attempt to upload large amounts of images intending to sell products related to trending topics that may violate Redbubble policies. In 2025, over 495,000 accounts were disabled by these measures – a 65% increase compared to the approximately 300,000 accounts disabled in 2024.

Objective 4: Empower consumers to make better informed choices of digital content.

The measures and tools discussed throughout this report summarise the ways that Redbubble detects and removes misinformation and disinformation, which mitigates the risk of harmful content and accounts that Australian end users could be exposed to on its marketplace.

Redbubble also puts artists front and center in its public communications, such as web copy and promotional materials, to ensure that consumers understand that images have been uploaded by independent artists and that the products offered on the marketplace are designed and sold by artists.

The Redbubble User Agreement makes it clear to users that artists are responsible for the content they upload, the products they offer for sale in their shops, and the titles, tags and descriptions they write to describe their products.

Email marketing materials contain the following statement: “All products on the Redbubble marketplace are designed and sold by independent artists”, and every product listing page created by artists contains the words “designed and sold by [artist’s username]”.

Objective 5: Improve public awareness of the source of political advertising carried on digital platforms.

Redbubble has opted out of Objective 5, because it is not considered to apply to the Redbubble marketplace, and Redbubble does not sell ad space to parties conducting political advertising.

Objective 6: Strengthen public understanding of Disinformation and Misinformation through support of strategic research.

Redbubble is open to supporting independent research that has the purpose of improving public understanding of disinformation and misinformation. At this time, Redbubble does not provide financial support to third-party strategic research in this area.

Objective 7: Signatories will publicise the measures they take to combat Disinformation.

Redbubble continues to publish these transparency reports, which are accessible to the public on the DIGI website at <https://digi.org.au/disinformation-code/>.

