



# 2019 TripAdvisor Review Transparency Report





# Table of Contents

---

Foreword	3
Executive Summary	5
Section One: A Global Platform for Content Submission	7
Section Two: How TripAdvisor Moderates Content	10
Pre-Posting Review Moderation	10
After-Posting Review Moderation	13
Volume and Speed of Assessments	14
Section Three: Protecting Users from Fake Reviews	15
What is a Fake Review?	15
Catching Fake Reviews	16
Volume of Fake Reviews	17
Types of Fake Reviews	18
Origin of Fake Reviews	19
Penalties and Enforcement	20
Penalising Paid Reviewers	21
What's Next?	23
Glossary	25
Methodology and Endnotes	28



# Foreword

---



**Stephen Kaufer,**  
**TripAdvisor President and CEO**

When TripAdvisor launched almost twenty years ago, none of us imagined just how large a community we'd build, or the impact TripAdvisor would have on the global travel economy. Today, hundreds of millions of unique visitors come to TripAdvisor each month – relying on more than 700 million reviews and opinions from people all over the world as they plan and book travel.

No matter where in the world someone is travelling, they can always be guided by the first-hand experiences shared by our global travel community.

Creating a platform for unbiased traveller reviews has also benefited the tourism industry. Now, both large and small hospitality businesses can reach a huge, global audience of potential guests – for free – and compete on the quality and value of the service they provide. For independent businesses with little to no marketing budget, this would have been otherwise impossible.

Over the years, I've heard countless stories from such businesses. One wonderful story in particular involves a small business in South Africa called [Siviwe Township Tours](#).

In 2007, a young man named Siviwe started offering tours of Langa, a part of Cape Town where he was born and raised. One day, two happy customers introduced Siviwe to TripAdvisor. Today, Siviwe Township Tours is one of the top-rated tours in Cape Town. It employs 15 tour guides who lead walking tours in Langa and surrounding townships. Siviwe also founded a successful community project, [HappyFeet](#), to provide a safe after-school environment for the children of his community.

Stories like this remind us of the incredible benefits that a platform like ours can bring to people all over the world. But, with that comes a great responsibility too. We have a duty to our users and to the industry to make sure that TripAdvisor is used fairly and honestly by all those who contribute content to it.

That is not easy. It requires a set of [clear guidelines](#) that every contributor must abide by. We have to effectively protect against fraudulent reviews that could mislead our customers and hurt businesses. And we have to make sure that legitimate reviews that meet our standards are posted as quickly as possible.

Getting that balance right across multiple languages requires the absolute latest in information technology, combined with a large team of highly trained moderators. Every hour of every day, we are working hard to get it right.

No one has a greater incentive than TripAdvisor to ensure the reliability of the content on our platform. It is the core of our business because if people don't find our content reliable and helpful, they won't keep coming back to our site.

However, too often we have seen third parties share inaccurate figures about the authenticity of review content and the scale of fraud on our platform.

These third parties do not have access to the key technical data necessary to determine whether or not a review is fraudulent. We do – and with this report, for the first time ever, we want to provide definitive insights into the details and data behind our extensive content moderation efforts.

We do this in the spirit of transparency. We know we're not perfect. But we're constantly working to stay one step ahead of the people who try to abuse our platform – and we believe that no other review platform does more to protect the integrity of their content than TripAdvisor.



# Executive Summary

---

Over the past 20 years, TripAdvisor has developed an industry-leading content moderation process to maintain the integrity of the more than 760 million reviews and opinions hosted on the platform.

This inaugural TripAdvisor Review Transparency Report provides insight into the moderation processes and key data for reviews submitted to the platform between January and December 2018.

The report follows the steps TripAdvisor takes to assess whether content submissions meet established guidelines. It also explains how TripAdvisor analyses every review before it is posted to the platform, as well as the extensive efforts undertaken to block, remove and minimise attempts to post fraudulent reviews.

## Key findings of the report include:

- In 2018, TripAdvisor received **155 million content submissions** from users. Of those, **66 million** were reviews.
- The average rating submitted by reviewers for businesses and locations listed on TripAdvisor was 4.22 out of 5.
- **100% of reviews submitted to TripAdvisor** pass through an advanced moderation process, which uses state-of-the-art analysis technology to identify potential problem comments.

- In 2018, 2.7 million of the reviews submitted were **further screened by a highly trained content moderation team**.
- 4.7% of all review submissions were **rejected** or **removed** by either the advanced analysis technology or manually by the content moderation team. There are a number of reasons why TripAdvisor rejects or removes reviews, ranging from guideline violations (such as the use of profanity) to instances of review fraud.
  - 3.4% of all submissions were rejected *before* they were posted to TripAdvisor, compared to 1.3% of all submissions which were removed *after* being posted.
- Looking specifically at fake reviews, **only a very small fraction of all review submissions – 2.1% – were determined to be fraudulent**, and the vast majority of those (73%) were blocked before they ever got posted.
  - Less than 0.6% of all reviews submitted (374,220 in total) made it onto TripAdvisor before being removed for fraud.
- **Fewer than 1%** of reviews were flagged by users or businesses for potentially violating TripAdvisor guidelines. Most of those community reports were reviewed within six hours of being submitted.
- TripAdvisor **penalises companies and individuals** who attempt to cheat the system with penalties that increase in severity for repeat offenders.

The Review Transparency Report also details how TripAdvisor plans to stay one step ahead of those who try to manipulate the platform by:

- Continuing to improve systems to identify fraud and penalise the perpetrators.
- Investing in training for human moderation teams to support the best decision making in a complex moderation landscape.
- Pursuing partnerships with law enforcement authorities to share information and support their efforts to tackle online fake reviews.
- Building on transparency efforts by sharing more insights into moderation processes and investigations on the online '[About TripAdvisor Reviews](#)' information resource.

*Note: All references in the report to data from the last year refer specifically to the period 1st January 2018 - 31st December 2018, unless otherwise specified.*



# Section One: A Global Platform For Content Submissions

---

## By the Numbers

66m

reviews submitted in 2018  
across every continent

53%

were for locations in  
Europe; 23% for those in  
North America

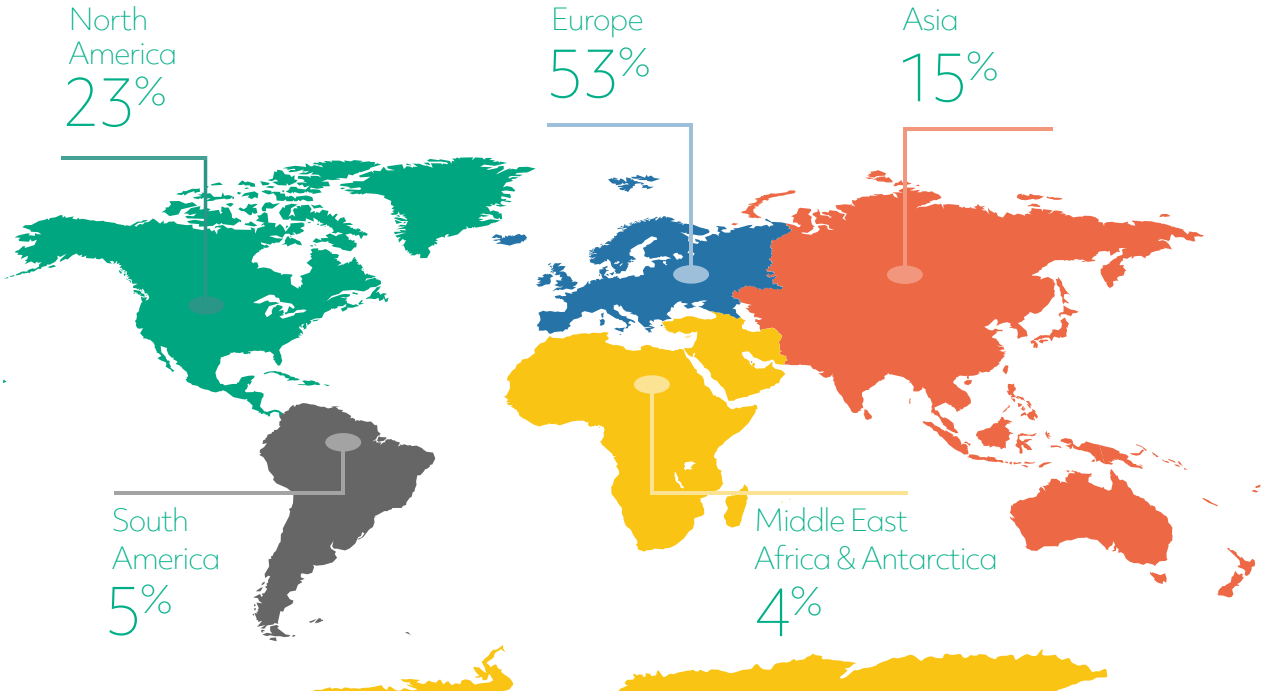
81%

had ratings of 4 or 5; just  
10% had ratings of 2 or 1


In 2018, TripAdvisor received approximately 155 million reviews and opinions from its members around the world. This includes a number of different forms of user-generated content (UGC), such as reviews, management responses and forum posts.

Looking specifically at reviews, TripAdvisor received more than 66 million submissions from its global member community in 2018. Every review submitted to TripAdvisor went through a rigorous moderation process, and 96.6% were posted onto the platform.

Review submissions were spread across every continent: 53% of review submissions were for establishments in Europe, 23% in North and Central America, 15% in Asia Pacific, 5% in South America and 4% in the Middle East, Africa and Antarctica.



In line with seasonal travel trends, the volume of review submissions spiked during the peak period between June and August 2018.






When submitting a review on TripAdvisor, users are required to submit a bubble rating on a scale from 5 (Excellent) to 1 (Terrible).

Last year, the average rating given by users on TripAdvisor was 4.22. In fact, more than eight out of ten reviews (81%) were submitted with a rating of 4 or 5 bubbles. One bubble ratings are rare, representing only 5.7% of reviews submitted in 2018.

When TripAdvisor asked reviewers about their motivations for writing reviews, 9 in 10 (87%) said they wrote a review because they wanted to share a good experience with other travellers(i).

Proportion of review ratings





## Section Two: How TripAdvisor Moderates Content

---

### By the Numbers

100%

of reviews submitted to TripAdvisor in 2018 were screened by an automated system before being posted

2.7m

submissions were further screened by content moderators

4.7%

of review submissions were rejected either before or after posting

TripAdvisor posts an average of 256 content contributions every minute, and operates in 28 languages globally. This presents complex challenges for review moderation, however TripAdvisor has robust systems and processes in place to screen content before it is posted and to moderate it afterwards.

### Pre-posting moderation

Before a review is posted to TripAdvisor, it is screened to assess whether it complies with the [platform's strict guidelines](#). These include making sure the content is relevant, non-commercial, based on first-hand experience and not biased by an affiliation with the business being reviewed.

These guidelines are vital to the success of TripAdvisor, both for travellers and businesses. Reviews that are most useful and relevant to consumers and businesses reflect a genuine, direct experience. Second-hand information or detail taken from other people's experiences or third-party sources is more prone to bias and misinformation. Biased reviews would make the platform unreliable as a source of credible information.

There are a number of reasons why a review might violate TripAdvisor's guidelines. Some violations may be unintentional, such as when a review is submitted for the wrong business, while others may represent a deliberate attempt to mislead consumers, i.e. fake reviews.

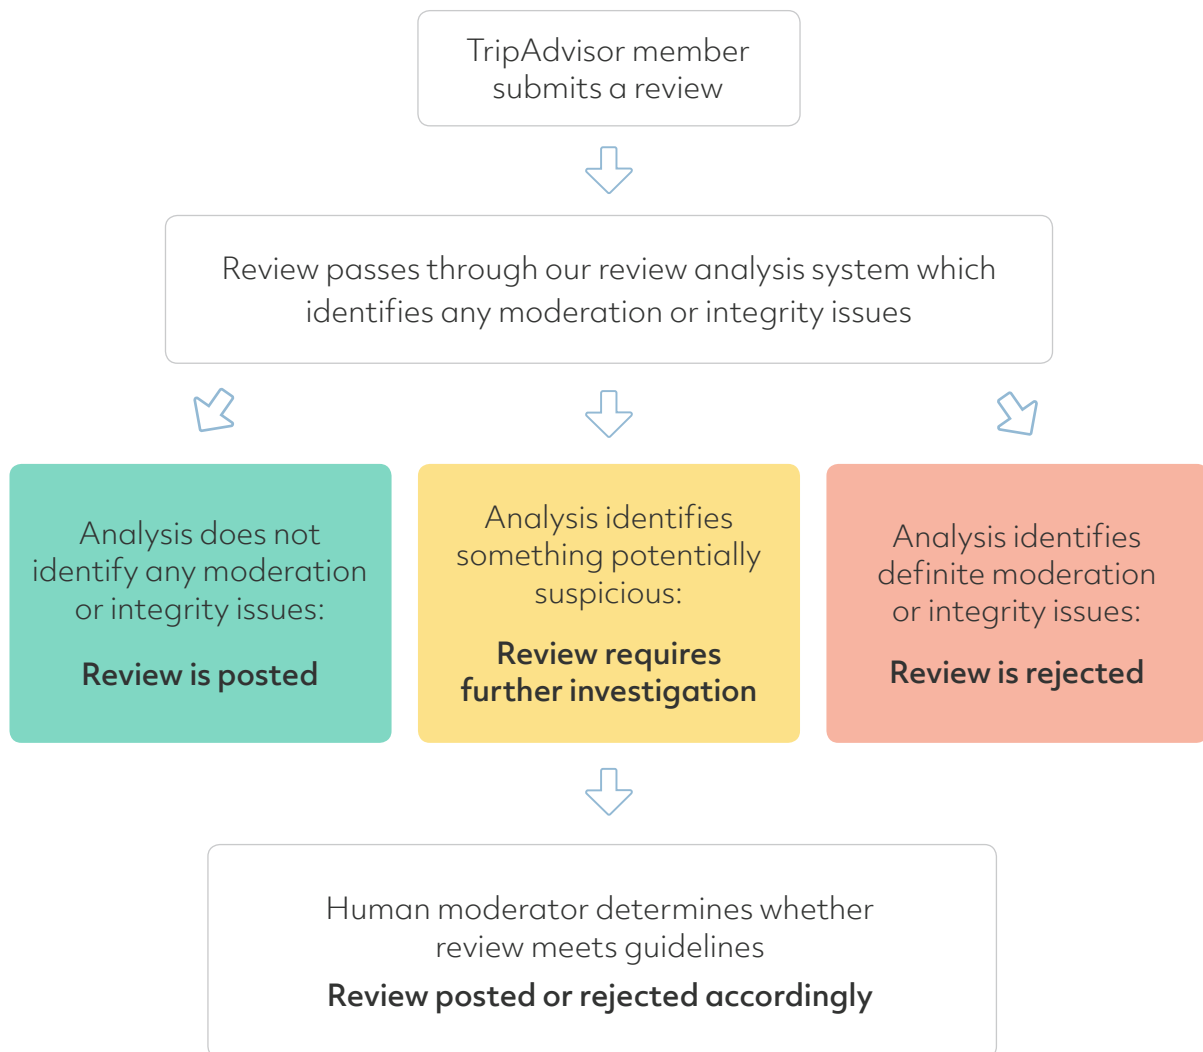
The first step in the pre-posting screening process is an automated **review analysis system**. This technology employs state-of-the-art data analytics to review hundreds of pieces of online information that can flag content that doesn't meet TripAdvisor's guidelines or that might be fraudulent. It uses network forensics and advanced fraud modelling to map electronic patterns that no human could ever identify.

If the system detects clear violations of TripAdvisor's guidelines, it will automatically reject the submission. Last year, the system rejected 1.4 million review submissions.

If the analysis system detects a potential problem, the review is sent to a content moderation team for further analysis. This dedicated team then assesses factors that the review analysis system cannot. That might include the context in which a review is written, or any cultural nuances in the wording. The analysis system or the moderation team might also, on occasion, reach out to a user to ask them to verify a detail of their review, such as whether it is for the correct location.

If the review analysis system doesn't detect any problems with a piece of content, it is posted on the relevant TripAdvisor page.

## How a review appears on TripAdvisor:



TripAdvisor employs hundreds of content analysts globally, working 24/7 to maintain the quality of the reviews on the platform, including analysts fluent in all 28 languages that the platform supports.

The combination of human and machine moderation is essential when processing content on such a large scale. By using advanced technology, TripAdvisor ensures that its moderation team is focused on potentially problematic content. This approach avoids a backlog of unpublished reviews while still achieving very high levels of detection and accuracy.

## After-posting moderation

TripAdvisor's objective is to catch suspicious content before it is posted to the platform. However, any system benefits from additional checks and balances, so in addition to pre-posting moderation, TripAdvisor also employs a robust process for after-posting community reporting.


Here's how that process works:

Anyone with a free registered TripAdvisor account can flag a review if they believe it violates the company's guidelines, either because of language used, relevance or suspicions about the authenticity of the reviewer. Registered business representatives can also report reviews in the Management Centre.

We refer to reviews that are flagged in this way as **reported reviews**.

Every reported review is acted on by TripAdvisor. This does not mean that a reported review will be removed. The moderation team will only remove a review if there is evidence that it violates the company's guidelines. Where there is no clear evidence, the review will remain live.

Decisions on reported reviews are made using identical criteria, regardless of whether the review was reported by a business owner or a traveller.



Fewer than  
1% of reviews  
were reported to us  
by the community

Last year, fewer than 1% of the reviews posted to TripAdvisor were flagged by the community. Of those reported reviews, 57% were found not to be problematic and 43% were removed after being assessed. This speaks to the vital importance of community reporting in helping to moderate the platform effectively, even though the number of reviews removed in this way represented just 0.3% of the total volume of review submissions.

## Total volume of content moderated

In total, 4.7% of all review submissions were rejected pre-posting or removed after-posting by either the analysis system or content moderators due to guideline violations (ii).

Guidelines violations occurred in a number of ways, ranging from simple errors (such as reviews found to have been submitted regarding an incorrect location) to more serious violations, such as reviews that were determined to be fraudulent.

## Volume of human moderation assessments


Last year, TripAdvisor content moderators examined a total of 2.7 million reviews, either pre-posting or after-posting. Of these, they rejected 62% for violating TripAdvisor guidelines.


This equates to an average of over 7,000 reviews manually assessed by the content moderation team per day – however, during the peak summer period, this figure increased to more than 10,000 reviews per day.

## Speed of assessments

When a review is reported by the community, the moderation team tries to resolve it as quickly as possible. Last year, the majority (79%) of moderation assessments on reported reviews were conducted within 6 hours of the report.

### Most assessments completed within 6 hours





## Section Three: Protecting Users from Fake Reviews

---

### By the Numbers

2.1%

of the 66 million reviews submitted were identified as fraudulent

73%

of those fake reviews were blocked by TripAdvisor's review fraud detection systems **before** being posted

<0.6%

of all review submissions were identified as fraudulent and removed **after** being posted

TripAdvisor devotes a significant proportion of its content moderation resources to catching fake reviews and penalising those responsible.

### What is a fake review?


TripAdvisor defines a fake review as any review written by someone who is trying to unfairly manipulate a business' average rating or traveller ranking, such as a staff member or a business' competitor. Reviews that give an account of a genuine customer's experience, even if elements of that account are disputed by the business in question, are not categorised as fake.

## Catching fake reviews


Over the past two decades, TripAdvisor has actively tracked billions of data points across hundreds of millions of reviews, allowing the company to detect clear patterns of what is typical review behaviour – and what is not typical. This helps catch fraud. (See illustration below).

To catch fraudulent reviews, TripAdvisor employs sophisticated review fraud detection technology that uses best practice techniques from the banking and credit card industry to map hundreds of discrete pieces of online information in order to help us understand more about every review's origin and circumstance, such as IP addresses and device specifications. Over the years, the system has identified the typical electronic behaviour of genuine reviewers. When the technology spots an electronic pattern that does not conform to this typical review model, it instantly flags the review as potentially suspicious. The review is prevented from being posted and sent to a member of the content moderation team for closer examination.

### Illustrative Example of Electronic Behaviour Patterns



Typical Reviewer  
Behaviour Pattern



Pattern of Suspicious  
Reviewer Behaviour

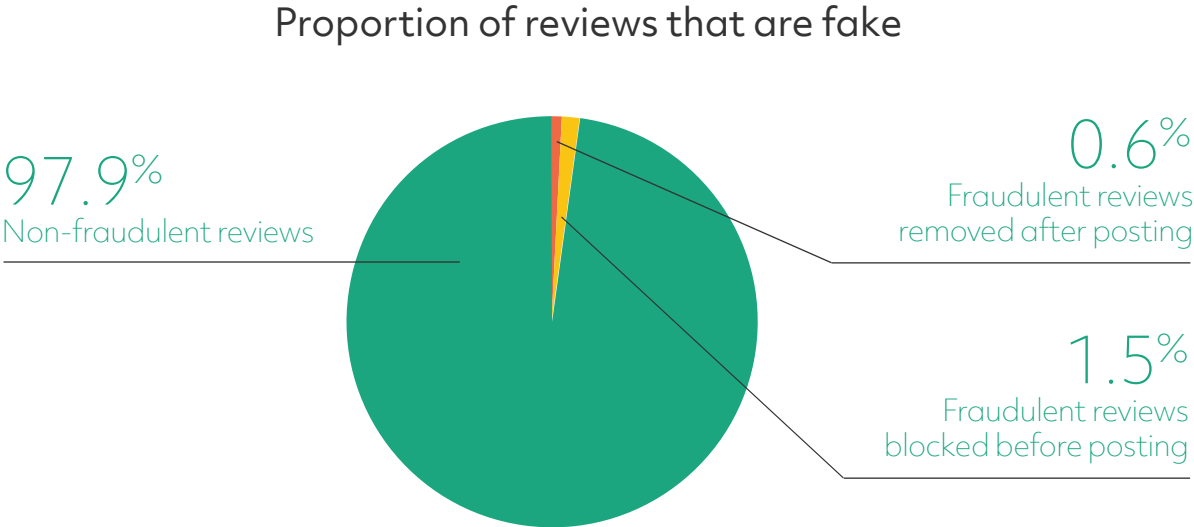


In addition to fraud detection technology, TripAdvisor has a dedicated team of expert fraud investigators who proactively target the websites, companies and individuals that try to sell fake reviews. These investigators focus on cutting off this type of fraud, known as paid review fraud, at the source.

### Volume of fake reviews

There is often inaccurate speculation about the volume of fake reviews on TripAdvisor. Here's the definitive story on what was detected in 2018:

- Of the 66 million reviews submitted to TripAdvisor in 2018, only 2.1% were determined to be fraudulent.
- The vast majority (73%) of these fraudulent reviews were identified and blocked by the fraud detection process **before** being posted – so they were never visible on TripAdvisor.
- That constitutes more than 1 million fraudulent reviews prevented from reaching the platform.
- Less than 0.6% of all reviews submitted (374,220 in total) made it onto TripAdvisor before being removed for fraud.
- 98.5% of the fake reviews caught were first identified by TripAdvisor, either by the review analysis system or content investigators.
- 1.5%<sup>(iii)</sup> of reviews identified as fake were removed as a result of reports submitted by the travel community. This represents 0.03% of the total volume of review submissions we received.

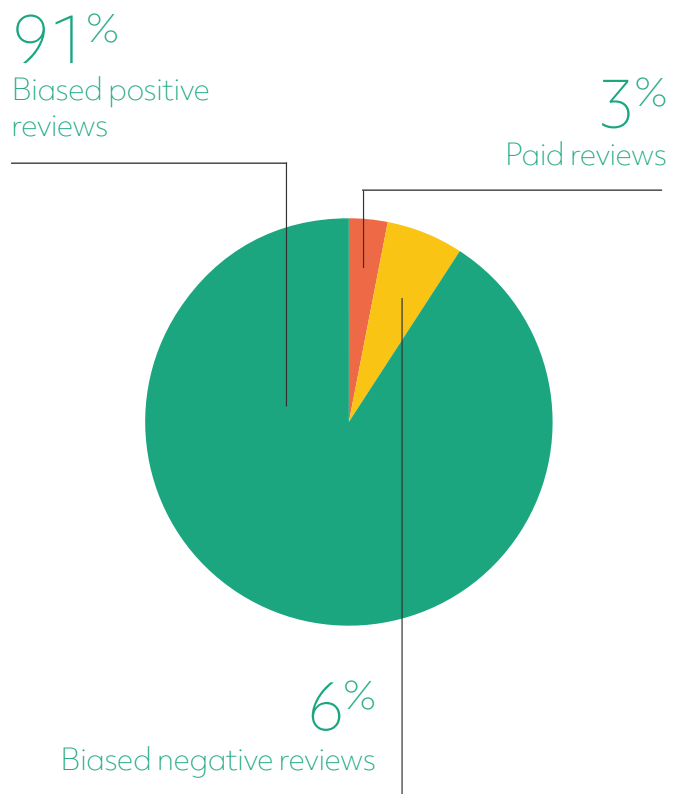


## Types of fake reviews

Fake reviews usually fall into one of three categories of business or property-based fraud:

- **Biased positive reviews** occur when someone who is connected to a business – such as an owner, employee or even a relative – submits a review in an attempt to improve the ranking or online reputation of that business. Reviews of this nature do not provide genuine, objective feedback that is useful for other travellers to make informed decisions. The vast majority of fake reviews submitted to TripAdvisor fall into the biased positive category.
- **Biased negative reviews** represent any attempt to submit a deliberately malicious review about a business in an effort to unfairly lower its ranking position or improperly discredit that business in some way. Most biased negative reviews come from one of two sources – either from someone connected to a rival establishment, or from a person who is trying to blackmail a business by threatening to submit a false negative review.
- **Paid reviews** occur when a business employs the services of an individual or a company in an effort to boost its ranking position or online reputation with fake positive reviews. Although paid review fraud represents the smallest proportion of fraud on the platform, it is a serious challenge for the industry and something TripAdvisor tackles aggressively.

### Proportion of Property-Based Fraud

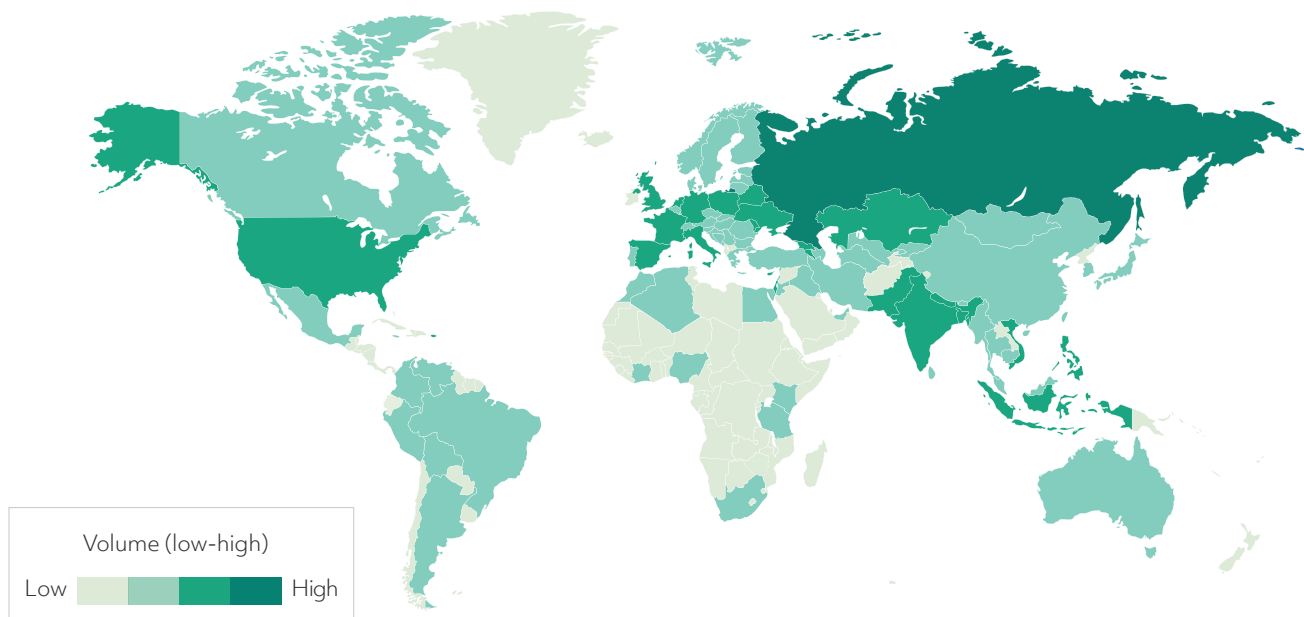


TripAdvisor's systems are also set up to identify and remove certain types of fake reviews that we refer to as member fraud. This can include instances where media commentary about a business or a high-profile incident at a property can spark a surge of reviews that do not relate to first-hand consumer experiences. In these cases, depending on the situation, a temporary lockdown of a business's TripAdvisor profile may be applied, which prohibits reviews from being posted to that business for a short period of time.

## Origin of fake reviews

Fake reviews originated from countries all over the world, however, when looking specifically at paid reviews, fraudulent activity is more prevalent in certain markets – most notably, in 2018, there was a higher percentage of fake reviews originating from Russia versus other countries.

### Source Location for Paid Reviews



In summer 2018, Russia hosted the FIFA World Cup. Major events like a World Cup attract a huge increase in visitors to the host country, and unfortunately there will always be fraudsters who try to exploit that situation. You can read more about how TripAdvisor's investigations team took proactive steps to protect against fraud during the World Cup [here](#).

## Penalties and enforcement

TripAdvisor believes that strong penalties are necessary to deter fake reviewers from targeting the platform.

Most fake review submissions can be connected back to a particular business that intended to benefit by fraudulently boosting their own ranking. However, it is important to note that most businesses listed on TripAdvisor play by the rules; fewer than 1% of businesses try to abuse the platform.

When it finds evidence of fraud, TripAdvisor applies a range of penalties:

**Content ban** – Individual members who persistently commit fraud face a content ban, which removes their ability to post reviews and content. As well as banning the user's account, the platform also takes measures to prevent the user from being able to create new accounts.

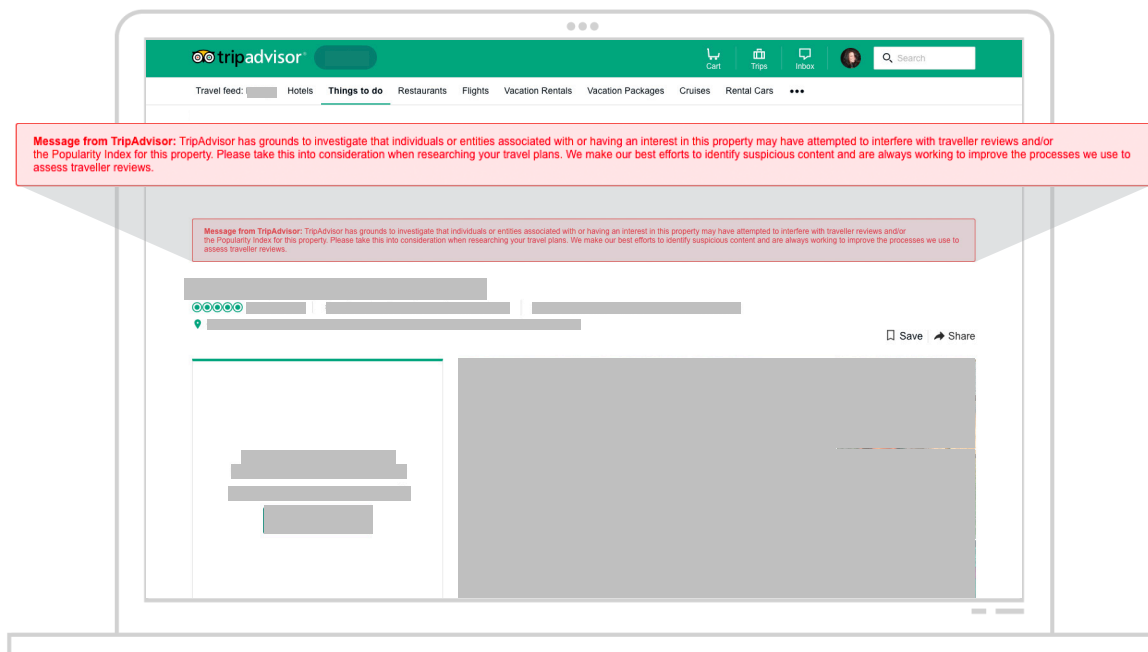
**Ranking penalty** – If a business is caught posting fraudulent reviews, the first step – after removing the fake content – is to impose a ranking penalty on the listing. A ranking penalty is a reduction of a property's position within TripAdvisor's popularity or traveller ranking for a period of time.

In 2018, **34,643** businesses out of the 8+ million locations listed on TripAdvisor received at least one ranking penalty. This figure includes businesses that did not submit fake reviews directly but encouraged or paid others to do so.

The majority of businesses receiving a ranking penalty quickly change their behaviour; in 2018, only 24% of businesses penalised in this way continued to attempt fraud after receiving a ranking penalty. However, in cases where fraud activity continues, further penalties are applied with increasing severity.

**A red penalty badge** – For the most persistent cases of fraud, TripAdvisor employs its strongest penalty. The red penalty badge is a highly visible notice on a business's TripAdvisor listing page that warns travellers about suspicious review activity spotted in relation to that business. It is accompanied by a significant enforced drop in the traveller ranking.

This severe penalty is only applied to the small minority of persistent offenders who have continued to attempt fraud despite being warned and penalised for suspicious activity previously.



TripAdvisor always notifies businesses before posting a red penalty badge (provided contact details can be identified for the business’s management) so they can respond to the allegations.

In 2018, TripAdvisor applied a red penalty badge to 351 businesses worldwide<sup>(iv)</sup>.

The effects of this penalty are significant, making it an effective deterrent against fraud. The enforced drop in the traveller ranking means the business becomes less visible to travellers, while the warning notification on the listing page makes travellers more cautious of the business. As this consequence is usually the opposite of what the fraud was designed to achieve, in most cases just the threat of a penalty badge is enough to stop fraud.

### Penalising paid reviewers

TripAdvisor also devotes substantial resources towards the fight against paid reviewers. Paid reviewers are individuals or companies that try to sell “user” reviews on either their own websites or on sites like eBay, Facebook and Fiverr.

TripAdvisor is prepared to take the strongest possible actions to stop paid reviewers, including legal proceedings and cooperation with law enforcement. For example, last

year the company supported the prosecution of a paid reviewer in Italy who received a nine-month prison sentence. Since 2015, TripAdvisor's efforts to fight paid reviews have led to the shutdown of more than 75 different paid review websites.

This aggressive investigative work has not gone unnoticed. There is evidence that fraudsters are increasingly aware of the risks that come with trying to target TripAdvisor, and are shifting their attention to other platforms that may have less sophisticated fraud detection measures.

### **Investigators in Action: A Case Study**

In 2015, TripAdvisor began gathering evidence against a paid reviewer in Italy after several Italian businesses forwarded emails flagging the fraudulent site.

The fraudster, who operated his illegal business under the name PromoSalento<sup>(v)</sup> was advertising fake review packages to hospitality businesses in Italy.

TripAdvisor investigators were able to identify details about the individual behind PromoSalento, including email addresses, residential addresses and personal connections. The team then applied advanced fraud detection techniques to identify which review submissions could be linked back to PromoSalento, and leveraged this technical analysis to find more than 1,000 attempts by PromoSalento to submit reviews to the TripAdvisor platform on behalf of hundreds of different properties. All of those reviews were either blocked or removed.

Armed with this data, TripAdvisor supported the prosecution of the fraudster as a civil claimant, by sharing evidence from the extensive in-house fraud investigations and providing support from Italian legal counsel.

As a result, in one of the first legal cases of its kind, in 2018 the Criminal Court of Lecce in Italy ruled that writing fake reviews using a false identity is criminal conduct under Italian criminal law. The fraudster who operated PromoSalento was sentenced to nine months in prison and ordered to pay approximately 8,000 euros in costs and damages.

It was a landmark ruling for the Internet and a milestone in the fight against paid review fraud. Writing fake reviews has always been fraud, but this was the first time someone was sent to prison as a result. TripAdvisor will continue to eagerly work with regulators and law enforcement to assist their efforts to prosecute other fraudsters.



## What's Next?

---

TripAdvisor is the global industry leader in review fraud detection. But complacency is not an option. We are constantly working to improve our ability to strike the right balance between being a valuable platform that lets travellers share experiences, while keeping the platform free of fraudulent reviews.

Fraudsters are always changing their tools and tactics, and we work tirelessly to stay one step ahead of them. Looking ahead, TripAdvisor's roadmap includes these enhancements:

- Continue to invest in training, processes and tools to improve decision making in an increasingly complex content moderation landscape.
- Improve the methods used to identify fraud and penalise the perpetrators – both from a technology and a human moderation perspective.
- Refine processes, tools, detection methods and policies around the new content types added to the platform in 2018 – links, videos and trips.
- Pursue and strengthen partnerships with law enforcement authorities to share information and support their efforts to tackle online fake reviews.
- Continue efforts to help travellers find and share safety information, introducing new features and policy enhancements that leverage TripAdvisor's scale and community to provide travellers with relevant and timely safety information.
- Share insights into the company's review moderation processes and investigations on the online ['About TripAdvisor Reviews'](#) information resource as part of TripAdvisor's transparency efforts.

## Stronger Together

TripAdvisor has a responsibility to protect the platform that we built for the benefit of everyone in the travel community – businesses and consumers alike. But there is always more we can do, on our own and with the help of our global community and the industry.

Our passionate community is a valuable partner in protecting the content on our platform, and we encourage any user who has a concern about a review to report it to us for further investigation. We act on every report we receive.

We also welcome more opportunities to work with hospitality businesses, consumer rights organisations, regulators and law enforcement agencies to identify and prosecute paid review fraudsters. By working together, we can be even more effective in shutting down these businesses. We encourage any business approached by someone offering reviews for cash to tell us about it. They can share any information on these companies or individuals by contacting our team directly at [paidreviews@tripadvisor.com](mailto:paidreviews@tripadvisor.com).





# Glossary

---

**After-posting moderation:** a process of determining whether content complies with the TripAdvisor guidelines which takes place after a piece of content has been posted to the platform.

**Assessment, or human assessment:** an assessment carried out by a member of the TripAdvisor content moderation team to determine whether or not a review complies with the content guidelines.

**Bubble rating, or rating:** a score from 5 (Excellent) to 1 (Terrible) that TripAdvisor users are asked to provide whenever they submit a review of a business or point of interest. A rating is intended to represent a reviewer's level of satisfaction with their customer experience.

**Business representative:** a business owner or member of staff who is registered as such on TripAdvisor.

**Content guidelines, or guidelines:** a set of rules that every piece of content must abide by, ensuring that user-generated content is useful and relevant to travellers.

**Content moderation:** the process by which TripAdvisor assesses whether reviews and other user-generated content comply with its content guidelines.

**Content moderation team, or moderation team:** a dedicated team at TripAdvisor tasked with ensuring reviews and other user-generated content comply with its content guidelines.

**Content removal:** the act of removing any content found to violate the TripAdvisor content guidelines. Note that once a review is removed, any ratings associated with that content no longer factor into the property's average rating or ranking position on the platform.

**Fake review:** any review that is written either by someone who is not an unbiased traveller (such as a business representative or competitor) or by someone who is deliberately trying to influence a property's average rating or ranking position even though they did not have a personal customer experience with the business they are reviewing.

**Listing:** a page on TripAdvisor that profiles a specific business or point of interest, and which displays user-generated content relating to that place, such as reviews and photos.

**Management Centre:** a dedicated section on TripAdvisor that registered business representatives can access.

**Management response:** a reply from a business representative submitted in response to a particular review. Management responses found to be in compliance with the content guidelines are posted below the review in question and are thus visible to all users of the platform.

**Posted/hosted:** the description given to any piece of content that is hosted live on TripAdvisor and therefore visible to other users.

**Pre-posting moderation:** a process of determining whether content complies with the TripAdvisor guidelines which occurs prior to that content being posted to the platform.

**Property:** a general term given to any business or point of interest listed on TripAdvisor.

**Ranking:** the Traveller Ranking of a property on TripAdvisor. The traveller (or popularity) ranking reflects a business's popularity based on the quality, quantity, recency and consistency of reviews on TripAdvisor.

**Ranking penalty:** the reduction of a property's position within the popularity or traveller ranking for a period of time.

**Red penalty badge:** a notice on a property's TripAdvisor listing page which warns travellers about suspicious review activity in relation to that business or point of interest.

**Rejected review:** a review that was found to be in breach of the TripAdvisor content guidelines at the pre-posting stage and at no stage was posted live to the platform.

**Removed review:** a review that is removed from TripAdvisor after being posted.

**Report:** an online form hosted on TripAdvisor which allows travellers or business representatives to notify TripAdvisor about a particular review if they believe that review may not comply with the platform's content guidelines.

**Reported review:** any review which a traveller or business representative has notified TripAdvisor about in relation to its potential non-compliance with the platform's content guidelines.

**Review analysis system, or analysis system:** the technology that TripAdvisor uses to track and analyse review submissions.

**Review fraud:** any deliberate attempt by an individual or group of individuals to post false, biased or misleading reviews.

**Review submission:** any review that has been submitted to the platform by a registered TripAdvisor user but is not yet posted or visible. Note that a submitted review is not immediately visible to other users on the platform, and must first go through a pre-posting moderation process to determine whether or not the review will be posted for others to see.

**User-generated content, or UGC:** any piece of content (such as a review, photo or forum post) that has been submitted to TripAdvisor by a registered user.



# Methodology and Endnotes

---

## How the data was gathered and analysed

All statistics are based on reviews submitted to TripAdvisor between (and inclusive of) 1st January 2018 and 31st December 2018, along with any further moderation actions that may have been taken on those reviews until 31st January 2019.

The metrics in the report focus on traveller review content only. Unless otherwise stated, the data does not cover other content types on the platform, such as management responses, forum posts, photos, videos and Q&As. For example, if a photo submitted alongside a review was queued for manual moderation, the review would not be counted as manually moderated in this report because the issue was only with the photo, not the review. However, if the review was also flagged for manual moderation by the technology system or reported by the community, it would be included in the data in this report.

The data collection methodology and its results have been vetted for accuracy and confirmed by multiple teams across TripAdvisor, including data analysts, review moderation managers and fraud investigators.

Every possible effort has been made to ensure the data in this report is accurate; however, it is not possible to completely eliminate every margin of error. By the same token, there is always the potential for human error in human content moderation, particularly on complex decisions where there is a degree of subjectivity – such as whether

a review includes offensive language, as some words and phrases may be offensive in some contexts and not in others.

There is also the possibility of a margin of error resulting from unforeseen technical errors, such as software bugs or tool outages, although given the scale of data included in the report and the rigor with which it has been analysed, any impact on the figures included herein would be minimal.

As a global report, the metrics are presented in aggregate across all 28 languages that TripAdvisor supports globally. With English being the largest language by review volume, the aggregate metrics are most representative of English-language reviews. Data related to smaller-volume individual languages might experience deviations from the aggregate averages, for example in areas such as processing times or removal rates.

Occasionally, a review may be reported more than once. For example, multiple users could dispute a single review or an owner could disagree with a decision to keep a review posted and may report the review a second time. For the purposes of this report, such a review is counted as 'reported' only once. However, the measure for turnaround times considers all reports, whether it is the first report of the review or a duplicate.

## Endnotes

*\*i Findings taken from The Power of Reviews project, which surveyed 23,292 TripAdvisor users across 12 markets (Australia, China, France, German-speaking, India, Italy, Japan, Singapore, Spain, Taiwan, United Kingdom, United States) between May and June 2019. The survey targeted users who had visited the site in the last 12 months.*

*\*ii Removal numbers refer to content that was submitted between 1st January 2018 - 31st December 2018, but removed or rejected between 1st January 2018 - 31st January 2019. This accounts for moderation of 2018 content that was submitted at the very end of the year, and includes removals of some reviews submitted prior to 2018. As moderation is an ongoing process, removal numbers continue to be updated; however, for the purposes of this report, the figures relate to removals in the 1st January 2018 - 31st January 2019 time period only.*

*\*iii 1.5% figure is differentiated from the 1% figure quoted on page 13, as that 1% reference is to the number of all reviews reported by the community, whereas this 1.5% figure is the number of fake reviews reported.*

*\*iv Figures refer to properties receiving fraud penalties 1st January 2018 - 31st December 2018. The fraudulent reviews leading to the penalties could have been submitted before 2018.*

*\*v PromoSalento was the brand name fraudulently used to sell fake reviews to hospitality businesses. It is not a genuine business and is no longer operating given the conviction of the person behind it. It has no connection to any unrelated businesses with the same name.*