

**Which?**

 the  
behaviouralist

SUMMARY REPORT MAY 2020

---

# The real impact of fake reviews

a behavioural experiment on how fake reviews  
influence consumer choices

# Introduction

One of the most striking ways in which the internet has transformed how consumers make purchases is the unprecedented access it has given them to the experiences and opinions of other customers. Consumers rely heavily on online customer reviews to help choose between a range of possible options and it has been widely found that reviews can have a strong effect on sales.

However, consumer trust in these reviews may often be misplaced as there is growing evidence that reviews are being faked to manipulate consumer decisions. Which? has investigated the spread of false and misleading information through online reviews for several years. Our research and investigations have identified myriad instances of products and services with suspicious reviews posted on popular sites like Amazon, Tripadvisor and eBay. This included instances of reviews being traded on social media, sellers offering customers free items or cash incentives for positive reviews and the use of paid 'review farms' or bots to flood products with five star reviews.

The potential for these activities to create consumer harm is clear – consumers exposed to misleading information are more likely to make bad decisions. However, the research around how fake reviews actually influence consumer behaviour is limited. This leaves a critical gap in the understanding of the problem for online platforms and regulators looking to solve the issues of widespread manipulation of customer reviews. Without a clear understanding of the harm, remedies are liable to be insufficient or interventions poorly designed.

To address this, Which? worked with the research consultancy The Behaviouralist, to produce experimental evidence of how consumer behaviour changes in the presence of fake reviews. The results were striking, and provided clear evidence that fake reviews cause consumer harm. We found **fake reviews were highly effective at misleading consumers, causing them to choose poor quality products instead of better alternatives**. All of our fake review scenarios had harmful effects on consumer behaviour and, in the worst scenario, demand for products with fake reviews increased by more than 135%. Furthermore, these effects were similar even when we left clear signs of manipulation in the review text

This research provides robust evidence of the harm from fake reviews and highlights the clear need for action to stop fake reviews. The results demonstrate that the harm arising from fake reviews could be vast, with many consumers at risk of being misled into spending their money on substandard products.

Given the strong evidence of harm identified in our research, the CMA must urgently investigate how fake reviews are being used to manipulate online shoppers, taking the strongest possible action against sites that fail to tackle this problem.

This summary report outlines the experimental approach we used and the key results, but for full details of the methodology and results please see the full report at [which.co.uk/realfakereviews](http://which.co.uk/realfakereviews).

# Approach: Experimental Design

We designed our research to answer five main research questions:

1. Can fake reviews influence what products consumers buy and to what extent?
2. Can fake reviews make people more likely to choose poor-value products?
3. How sophisticated does review fakery need to be to influence purchase decisions?
4. Can platform endorsements whose award may be linked to fake reviews, lead to greater harm?
5. Can consumer education mitigate the influence of fake or suspicious reviews?

To answer these questions, we used an experiment in which **9,988 consumers** completed a hypothetical shopping task through an online survey. The participants were not informed that the task was related to fake reviews.

Participants were asked to choose between three product categories: headphones, dash-cams and cordless vacuum cleaners. We included those products because these are products for which fake reviews have been found in Which? investigations. Once participants decided the product category, they were then asked to choose their favoured option among a choice of five different products presented on screens designed to look like the Amazon website. They could review information on a search page, designed to look as though they had just searched for a product category (e.g. headphones) on the Amazon site, and then choose to see more information on each of the five products listed on that search page, including seven written reviews of each product.

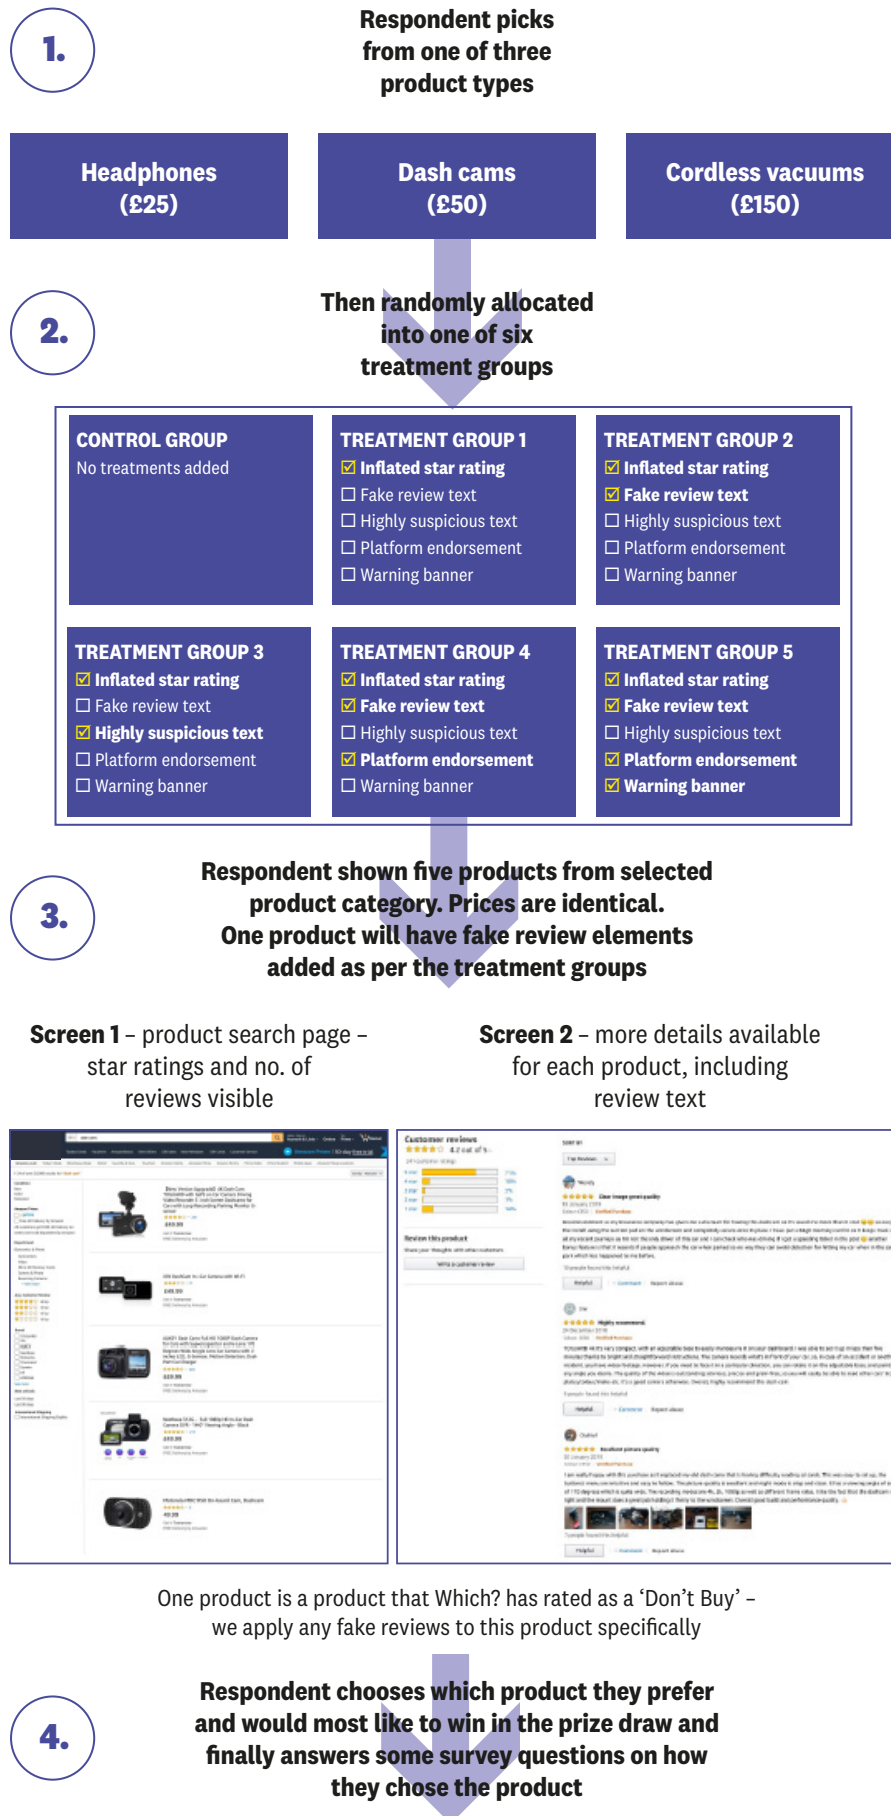
The pages and content that were shown to participants as part of the survey were not real Amazon pages or content. Amazon states that it has policies that prohibit abuse of its community features, including reviews, and that it suspends, bans and takes legal action against those who violate these policies. Amazon also states it has systems and tools in place aiming to stop abusive reviews before they are ever published, and allowing customers concerned about the authenticity of reviews left on a product to use the "report abuse" link, available on each review.

At the start of the experiment, all participants were informed they would have a chance of winning whichever product they chose, which was important as it provides the participants with an incentive to make choices as they would if they were actually making the purchase in real life.

The consumers were randomly allocated into a control group and five treatment groups (see Figure 1), in which they were exposed to different levels of fake reviews that were added to one of the five products they could choose. The fake review content on the pages participants saw were not real but incorporated features that Which? has seen in real life. They were not taken from actual Amazon pages and in real life these particular fake reviews may have been prevented from appearing on the site.. In one group (the control group), no fake reviews were added, giving us a baseline against which to measure how introducing different types of fake reviews influenced participants' choices. For one treatment group the fake review product also had a platform endorsement, while the final treatment group saw a warning banner shown at the top of the screen educating consumers about fake reviews.

To get a measure of how susceptible consumers could be to making bad choices as a result of fake reviews, participants were asked to choose between five products of variable quality as tested by Which? experts. This included one Which? Best Buy product and one Don't Buy, with a further three 'filler' products in between. We always added the fake reviews to the Don't Buy in order to test whether fake reviews could lead the participants into making a demonstrably poor choice and create consumer harm.

Figure 1 – Experimental design

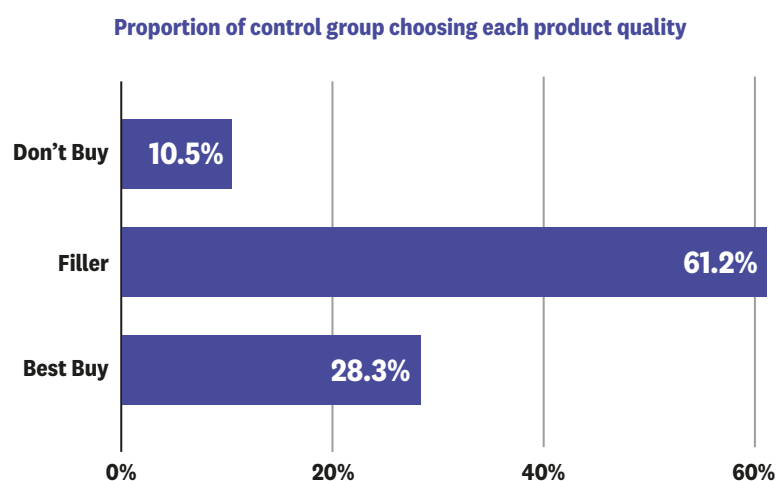


# Results

Below, we present the key results averaged across all three of our product categories with reference to the control and treatment groups in Figure 1 above. All the results presented in this summary report are statistically significant and similar across the three product categories. For results for individual product categories please see the full report.

## Consumer choices without fake reviews (Control Group)

**In the control group, with no fake reviews, 10.5% chose a Don't Buy product**, with 28.3% selecting a Best Buy and the remaining 61% a filler product. This suggests a retailing and review system that works reasonably well in guiding consumers towards good choices and away from poor choices, given that only around one in ten consumers chose the poorest-quality product available. In addition, the Best Buy product was the most chosen product while the three filler products were chosen 60% of the time.



## Inflated Star Ratings (Treatment Group 1)

When companies use fake reviews, they often target an improvement in the average star rating of their product. To achieve this they may buy or incentivise positive reviews, or employ so-called 'reputation management' services or bots to post a large quantity of five star reviews on their products.

In our experimental setting, we mirror this type of behaviour by presenting some participants with identical information to those in the control group but with the Don't Buy receiving a large increase in the number of five star ratings. We also moved the product to the top of the search listings to reflect the impact that average ratings can have on search rankings. Any corresponding increase in the proportion choosing the Don't Buy product can then be attributed to this rating and ranking effect.

We found that by making these adjustments, **the proportion of participants choosing a Don't Buy product rose to 16.3%, in comparison to 10.5% in the control group**, meaning an increase of 55% in the number of consumers choosing the low quality product.

## Inflated star rating treatment



**This result demonstrates that simply manipulating reviews to include a larger number of five star ratings can be highly effective at driving consumers towards a low quality product, even when the text of the reviews left for products remains relatively negative.** This is consistent with the literature on reviews more generally, where research has found consumers use star ratings as a heuristic to judge quality and consensus of other consumers' opinions.

### Inflated Star Ratings and Fake Review Text (Treatment Groups 2 and 3)

More sophisticated methods of review fakery used by companies to mislead consumers include the production of reviews containing highly positive descriptions. If done convincingly such methods may be more effective at influencing consumer opinion, but are also prone to leaving consumers hints that the reviews have been manipulated.

To test what effect the manipulation of the text of reviews might have on consumers, we presented some participants with faked text in the seven reviews beneath each product in addition to the inflated star ratings discussed above. To reflect real examples accurately, we made the text of the reviews more positive while also leaving some indicators of suspicious activity that could be spotted by a well-informed consumer, including exaggerated language, repetitive phrases and fewer verified purchases.

When this fake review text was added along with the inflated star ratings, **we found that 23.1% of people chose the Don't Buy product, which is more than double the proportion in the control group.** Furthermore it also led to a large increase (+6.8pp) over the simple manipulation of the star ratings.

We also found that even more egregiously fake reviews could have a similar effect on consumer choices. Some participants in the experiment were presented with reviews containing even more signs that the reviews had been manipulated, including admissions that reviewers had been incentivised to leave positive reviews and signs that positive reviews of completely different products had been hijacked. **Despite including these clear signs of manipulation, 21.6% of the participants still chose the Don't Buy**, only slightly less than the less obviously fake reviews and still more than twice as many as the control group.

### Inflated Star Ratings, Fake Review Text and Platform Endorsement (Treatment Group 4)


Some online review platforms give endorsement labels to products or services receiving particularly good customer feedback, providing a potential route for those manipulating reviews to extend the influence of that manipulation. In our experimental setting, we explored the impact of these endorsements on consumer choice through the inclusion of an “Amazon’s Choice”-style treatment. While such endorsements are not an element of fake reviews as such, some of Which?’s recent investigative research has found instances where fake reviews had contributed to products receiving platform endorsement labels.<sup>1</sup>

Our experiment found that **when fake reviews were coupled with the platform endorsement label, an even higher proportion of consumers chose the Don’t Buy product, amplifying the potential consumer harm from the fake reviews.**

Adding the platform endorsement label further increased the proportion of people choosing the Don’t Buy product with an inflated star rating and reviews containing faked text. **This means that around one quarter (24.8%) of our participants in the experiment were choosing the Don’t Buy product in this treatment group in comparison to only around one in ten (10.5%) in the control group – an increase of more than 135%.** These results demonstrate that fake reviews could have a huge impact on purchasing behaviour, particularly when coupled with endorsements linked to the reviews.

### Warning banner (Treatment Group 5)

Finally, we tested whether an information remedy might be able to reduce the harm from fake reviews. We tested this using a simple intervention in the form of a warning banner at the top of all of the search and product information screens for every product, not just the Don’t Buy. The banner warned about the possible presence of fake reviews and had tips consumers could follow to avoid being misled.



**Please be aware that the product pages may contain untrue or misleading customer reviews.**

**Some tips to avoid fake reviews:**

- Inspect the comments (don’t rely on star ratings alone).
- Watch out for suspicious language (such as lack of punctuation, capitals or odd formatting).
- Be suspicious of products with an unusually large amount of reviews.
- Check the dates of the reviews (if a lot of the reviews were posted the same day, be more cautious).
- Check the one-star reviews, and in particular if they seem to directly conflict with the praise in the five-star reviews.

To learn more about spotting fake reviews, please visit [This Page](#).

We found that the banner had a significant impact in reducing the proportion of consumers choosing the poor product, with around one in five consumers (19%) choosing the Don’t Buy product when the banner was introduced, down from around one in four (24.8%) when the fake reviews were present. This, however, still represents an 81% increase compared to the control group, demonstrating that **while the intervention could reduce the harm arising from fake reviews it is certainly not sufficient to eliminate the harm.**

The results presented in this report are consistent across demographic groups. They are also consistent across product types, indicating that consumers are likely to be deceived at a range of price points. Indeed, **in our experiment consumers were most likely to buy the Don’t Buy product with fake reviews in the most expensive product category, where consumer harm is likely to be greatest.** Finally, we found that fake reviews attracted choices away from all other products, but disproportionately so from the best (Which? Best Buy) alternative.



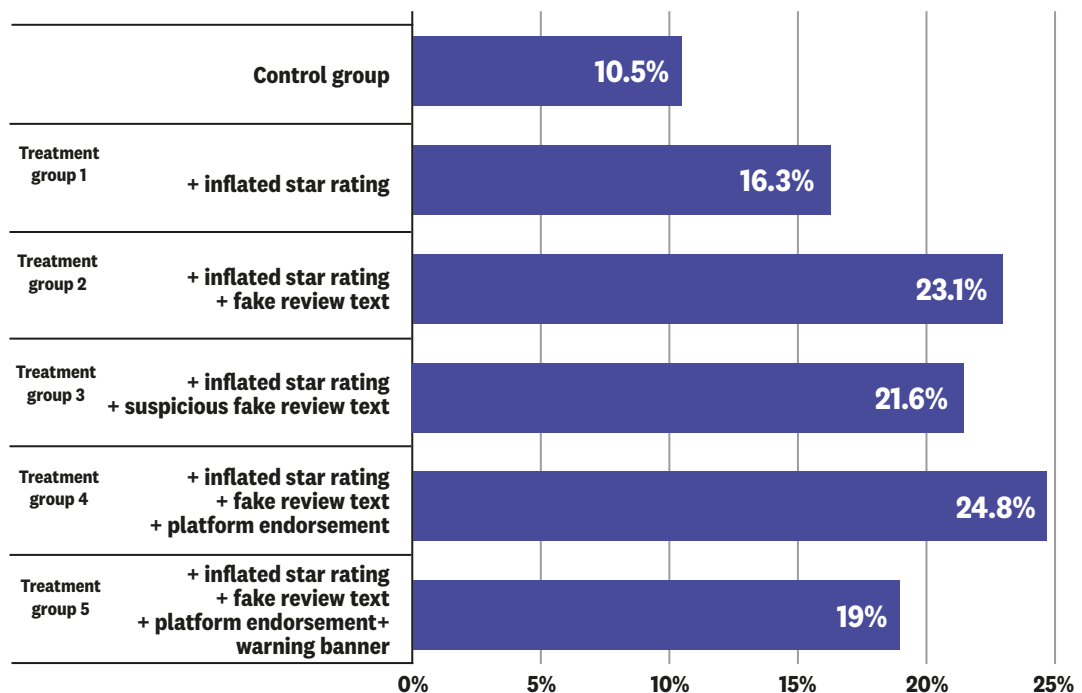
# Summary and conclusions

**Our results are unambiguous and provide clear indications of the harm that fake reviews can cause for consumers as well as initial insights into possible remedies to reduce this harm.** Our results are an important addition to the existing research on fake reviews and reviews more generally. Other researchers have identified the presence of fake reviews, but findings on how they affect consumers are few and inconclusive.

In summary, we found that (see Figure 2 for detail):

- adding fake reviews to the Don't Buy products in our experiment led to more than double the proportion of participants choosing them as their favoured option.
- both inflating star ratings and including fake text in the product reviews led to more consumers choosing the product, even when clear signals that reviews were being manipulated were included within the review text.
- when the platform endorsement was added to the Don't Buy product, the proportion of consumers choosing it increased even further.
- including a simple and untargeted banner at the top of the shopping screens reduced the harm associated with fake reviews, but a significantly higher proportion of consumers still chose the poor product than when there were no fake reviews.

**Figure 2 - proportion of respondents choosing the Don't Buy in each treatment group**



This research serves as an important evidence base for regulators and all review platforms to acknowledge and find effective ways to reduce the impact of fake reviews on their customers. It demonstrates how large the effect on consumer behaviour can be and as such it is paramount

that regulators and review platforms, such as Amazon, eBay and TripAdvisor take responsibility to make improvements to systems to remove fake reviews and stop consumers being misled. While the set-up of this experiment was designed to look similar to Amazon, as the UK's largest online shopping platform, we expect these results would be similar across review platforms where consumers use online feedback from other consumers to inform their purchasing decisions.

Which? remains concerned that review platforms are not doing enough to tackle the system wide problem of fake reviews. Amazon has policies in place to prevent review manipulation, including allowing users to report abuse. However our investigations continue to find examples of suspicious reviews across major websites. Therefore we welcome the regulator's investigation into whether major websites are doing enough to ensure that the measures in place are effective in ensuring that their customers are not being misled. People should be able to trust the reviews they read on these sites, without risking unwittingly purchasing a poor quality or even unsafe product. Platforms that host reviews should, as a minimum, uphold the following principles:

- Sites should have effective and adaptive policies and processes in place for preventing and removing fake reviews.
- They should take swift and effective enforcement action against those that breach these policies, including rogue sellers, businesses that incentivise illegitimate reviews and those offering 'black hat' services that offer to manipulate platforms' policies and processes in order to inflate product ratings and reviews.
- Sites should only accept reviews from customers who have made genuine transactions. Equally, platforms should be transparent about reviews that have been incentivised and such reviews should reflect genuine consumer experiences.
- Only reviews for genuine transactions should be used in determining rankings and endorsements awarded by sites. Sites should be clear about how such rankings are determined.
- Review hosting sites should be transparent about their business models, including the services offered to firms that pay to use their sites, and how they display reviews.

Given the strong evidence of harm identified in our research, we welcome the CMA's investigation into how fake reviews are being used to manipulate online shoppers on major websites, and expect the regulator to take the strongest possible action against sites that fail to tackle this problem.

**Which?**

Which?, 2 Marylebone Road,  
London NW1 4DF  
Phone +44 (0)20 7770 7000  
Fax +44 (0)20 7770 7600