
Make the Internet Moderate Again



By Nadav Klein , INSEAD

Extremist opinion looms larger when the moderate majority stays silent online. But there's a simple and cost-free way to balance out the discourse.

When the free market meets the free-for-all of internet opinion, the result should reflect a pure democracy – even if it outwardly resembles a chaotic mess. On the theory of the wisdom of crowds, the welter of customer posts, reviews and ratings should, in the long view, cohere into an accurate barometer of quality. Choosing between workplaces, mobile phones, automobiles – anything subjected to the cogitations of the internet hive mind – should be a simple matter of weighing online sentiment averages against financial considerations such as price or salary.

Unfortunately, the reality is not so simple. Far from handing everyone a microphone, the internet has developed into a system that grants bullhorns to some while reducing the voices of others to a veritable whisper. Part of this has to do with engagement-boosting algorithms that amplify extreme opinions. Intense emotion motivates far more clicks and shares than sober, fact-based analysis.

But much of it is about self-selection. Opinions between the extremes are not only less galvanising to the public, but they also generate less urgency on the part of the opinion holder to share their thoughts. Someone with a modestly positive or negative opinion is less likely to race to their computer to write a review than a disgruntled or delighted customer. The agitated few thereby out-influence the moderate majority. This should be a matter of concern for those who care about the general accuracy of online opinion. However, the direct relationship between vehemence of feeling and willingness to speak is a fact of human nature that can be overcome – with the right incentives.

Behind the Glassdoor

My **paper** in the *Journal of Experimental Psychology: Applied*^[1] explores this phenomenon in a particularly consequential context: online workplace reviews. Composed by employees under the protection of anonymity, these reviews are meant to capture an unvarnished view of the organisation from the inside. If the prevailing extremist pattern extends to this area of the internet, job seekers and industry observers (not to mention the companies themselves) would receive a deeply distorted picture dominated by the views of a few misfits and model employees. The entire spectrum in between – ranging from contented to disengaged – would be elided.

For the paper, we examined 188,623 Glassdoor reviews written by US employees. Most (76 percent) were purely voluntary; the remainder were written by users who received a “give-to-get” prompt requiring them to leave a review in order to continue viewing Glassdoor content. This promise to see additional Glassdoor content acted as an incentive to post a review.

As expected, the spontaneous reviews were significantly more polarised in their distribution, as well as more extreme in their assessments, than the incentivised reviews. One-star reviews were 1.4 percent more prevalent in the voluntary group; five-star reviews were 4.3 percent more prevalent. Without added incentives, extremists indeed exercised outsized influence on Glassdoor. The reason is simple: People with moderate opinions about their employers were less motivated to post reviews than people with extreme opinions. Adding the incentive of gaining access to additional Glassdoor information helped mitigate the motivational deficit of moderate reviewers.

Moreover, there were substantial differences between the two groups of reviews that could meaningfully affect decision making. For example, the

aggregated voluntary reviews (which, again, comprised more than three-quarters of our sample) ranked advertising and marketing companies higher on the whole than consulting, and the insurance industry above investment banking. But the incentivised reviews found the reverse. It is sobering to think that in today's volatile labour market, where switching industries is becoming commonplace, entire careers may hinge on the accuracy of this information.

Testing incentives

If incentives are the key to eliciting reviews from non-extremists, then it behooves us to know which types are most effective. That was the purpose of our next experiment. We recruited participants through Amazon's Mechanical Turk platform, promising payment in exchange for completing a survey about employer reviews. Participants were randomly split into "choice" and "forced" conditions – the former were offered the option to submit a review of their current employer with no financial penalty if they declined, the latter were told to write an employer review or else forfeit financial compensation for participating in the survey.

In addition, we randomly assigned some participants to various incentives for leaving a review. These included small and large financial incentives. There were also "prosocial incentives," emphasising how posting a review would help others, with either a neutral ("communicate important information to help others"), positive ("reveal the best employers to work for") or negative ("expose the worst employers") tenor. The large monetary incentive and the neutral prosocial incentive triggered the greatest increases in response rates, as well as less extreme review distributions.

We concluded that it takes more than a slight bonus to motivate non-extremists to perform the mental labour of writing a review. Incentivising moderates at scale with cash benefits would therefore probably not be cost-effective. Prosocial messaging, by contrast, costs little but the "warm glow" it provides can be emotionally valuable for many people. To draw out moderates, such messaging should understandably be neutral in tone rather than suggestive of either positive or negative sentiment.

Gold in the nuances

To be sure, our results may not sway companies that style themselves or their products best-in-class. Their natural inclination may be to encourage

positive extremism in online reviews, since they have very few detractors anyway. For them, soliciting moderates may produce a higher proportion of three- or four-star ratings to mottle their five-star record. All others, however, may grasp the benefit of providing more nuanced information to help consumers make educated trade-offs between cost and quality, or to assist job seekers in balancing desired compensation with other workplace priorities.

In truth, though, organisations on every level should acknowledge there's gold in the nuances. Often, companies solicit Glassdoor reviews only from employees they know will give fulsome praise. The gushing productions of model employees certainly make the organisation look good to outsiders, but they contain limited information for accurate matching between organisations and applications, as well as little constructive feedback for fuelling value-creating improvements. Online employer reviews are a matching mechanism – not a beauty contest. By attempting to game the system, firms may deprive themselves of priceless information hidden in the heads of their non-extremist employees. These rich deposits could be theirs for the asking, if they use the right prosocial language.

[1] Co-authored by Ioana Marinescu (University of Pennsylvania & NBER), Andrew Chamberlain (Glassdoor) and Morgan Smart (Facebook)

Find article at

<https://knowledge.insead.edu/leadership-organisations/make-internet-moderate-again>

About the author(s)

Nadav Klein is an Assistant Professor of Organisational Behaviour at INSEAD.

About the research

"Incentives can reduce bias in online employer reviews" published in *Journal of Experimental Psychology: Applied*.