

Written Testimony of Alex Roetter

Former Senior Vice President of Engineering, Twitter

Before the U.S. Senate Committee on Homeland Security and Governmental Affairs

On "Social Media's Impact on Homeland Security: Part I"

Wednesday, September 14, 2022

Good morning Mr. Chairman and Members of the Committee. Thank you for holding these hearings and for inviting me to provide testimony.

We live in a world where an unprecedented number of people consume news and information based on what is fed to them by a small number of massively influential social networks. A small group of people run these companies and have substantial power over shaping reality for billions of people. As others in front of this committee noted, viral content, misinformation, and disinformation can propagate via these platforms on a scale unseen in human history. The companies responsible for amplifying viral content, while focused solely on maximizing their profits, have either not taken responsibility for the effects this content has on society or paid lip service to them.

Regulators must understand these companies' incentives, culture, and internal processes to fully appreciate how resistant they will be to changing the status quo that has been so lucrative for them. Without increased transparency and a change of incentives, we should expect the same behavior and continued lack of meaningful, quantifiable progress on these problems.

In over 20 years of working in Silicon Valley as an engineer and an executive, I have seen first-hand how several companies prioritize features, launch products, and optimize their metrics. I was an early engineer on the Google syndicated ads product. Then, I started the ads engineering team at Twitter, growing it to \$2.5B in annual revenue, before taking over as SVP of engineering for all development efforts at the company.

Everything I will talk about today is an inevitable result of the culture of the companies and the incentives they face. While it may be easy to criticize specific leaders, that is the wrong target. As a thought experiment, we would not see a meaningful change even if we were to replace leadership at all these companies. What must change is the incentive system that makes the companies act the way they do.

Today I will explain the internal systems essential to understanding how things can or won't change, and the types of data we should seek. Finally, I will recommend how to obtain and analyze that information.

## Product Development

The product development life cycle works as follows:

1. Small teams of product managers, engineers, and designers brainstorm ways to meet and exceed specific internal goals that are measured with high precision. Metrics assigned to them come from various user growth and retention metrics (churn, engagement, time on sight, new user signups) and financial health metrics (revenue, ad price, click-through rate). Underlying these goals is a mission to drive revenue and maximize shareholder value and stock price. Typically, other metrics (user safety, etc.), if present, are a distant second in terms of importance.
2. To measure their effectiveness, these teams use an experimental system to launch their new changes to a small percentage of traffic (1% or 0.1% are typical values). I spearheaded the development of this system while at Twitter. A very similar approach is also in place at Google.<sup>1</sup> I believe other companies also operate related systems. This system logs a slew of data for every live experiment. Teams use this data to show per-experiment effects on various user and revenue metrics. Noticeably absent were any values tracking impacts on trust and safety metrics. For example, I never once saw any indication of how a given experiment affected any types of manually reviewed or actioned content, e.g., “Did a given experiment increase the reach of content later identified as hate speech?”
3. Executives will, typically once a week, as was the case at Twitter, hold an experimental review meeting and review all active experiments. After reviewing the data, they would shut down experiments, ramp them up to a more significant percentage of users, or fully launch them to 100%. Product and engineering ran these meetings. To the extent other functions like legal, finance, or trust and safety are present, they play a minor decision-making role compared to product, engineering, and sales leaders. Frequently they are not even present (and not viewed as mandatory attendees for a quorum).

## Company Culture

It is essential to understand the culture of these companies in order to appreciate how the product development life cycle is applied. In terms of culture, they are hierarchies, with the “builders,” namely, engineers, product managers, and designers, held in the highest regard. The strong norm in Silicon Valley is to not get in the way of “builders.” This stems in part from the fact that the founders of these companies come from “builder” backgrounds. Other functions (e.g., legal, finance, compliance, etc.) are viewed very skeptically, and the bias is to make sure “corporate bureaucracy” doesn’t slow down the building and releasing of new products or features. Perhaps no single illustration of this is more famous than Facebook’s old motto, “Move fast and break things.” A significant premium is placed on not slowing down the pace of development through reviews or bureaucracy.

---

<sup>1</sup> <https://static.googleusercontent.com/media/research.google.com/en//pubs/archive/36500.pdf>

## Review & Promotion System

The companies evaluate individuals via quarterly, semi-annually, or annual performance cycles. In this process, individuals typically write up self-reviews describing their accomplishments, reviews of their peers, and reviews up and down the management chain. Reviews heavily consider the individual's impact on their team and company. The way impacts are described is first and foremost in terms of what products or changes an employee built and what results they achieved. Benefits are measured most prominently by user engagement and financial metrics. Reviews determine the stack ranking of employees relative to one another, eligibility for cash or stock bonuses, and are inputs to the promotion process. Reviews are also considerations when companies make downsizing decisions in leaner times.

The promotion systems that I have observed highly emphasize impacts on key company metrics as well. Engineers highlight the effects that their work had on the company in as quantifiable a way as possible to promotion review committees. The most effective way to do this is to show an impact on financial or user engagement metrics. To the extent that you can get promoted for something on the trust and safety or privacy side of the house, a minority of builders get promoted this way, and the impact isn't valued as heavily as the core drivers of the company's growth. For example, the fastest way to get promoted as a front-end engineer on a core product feature team is to show that you made changes that drove more viral growth or increased revenue. There is an inherent tension between this progress and improving trust and safety. The former nearly always wins over the latter.

## Company Behavior

If one understands these companies' incentives, product development processes, and internal culture, it should be no surprise to see them respond the way they do to external pressures. While attempting to convince the public that they are making serious investments in these areas, we see them continuing to be unaccountable for any measurable results against the questions of interest to this committee. They are incented to delay any oversight while they continue to build some of the most valuable companies in the world.

To change things, we should not accept certain types of answers as sufficient, and we must demand more transparency. Finally, we should be able to observe if their incentives have truly changed by pressing for transparency on their internal processes and testing to see if they are acting as one would expect against a new set of incentives.

## Historical Data Sharing

Historically companies have released selective statements designed to convince people that they are taking the problem seriously in response to public or regulatory pressure. Examples of these statements from the companies include: "We are spending \$X on this problem," "We are employing Y people," or "We remove Z pieces of content every day/week/month." These

numbers lack context (i.e., a denominator) and are shared because the numbers seem large in absolute terms. They are designed to convince external parties that they are tackling the problem. These numbers are significant because they come from some of the largest companies in human history. Alphabet/Google and Meta/Facebook are in the top ten companies by market capitalization globally. Any single spending or investment figure will look large, but the numbers are less impressive in the context of how much they spend on other initiatives (new product initiatives, stock buybacks, stock-based compensation, etc.).

Similarly, how many pieces of content are actioned or removed tells you nothing about what fraction of content that represents or the reach of the content (which is the crucial part). Finally, simply saying they are investing a lot is not the point. The metrics that the public receives should be measuring outcomes (e.g., as proposed by an external group), not simply about the size of the investment. If companies share only how much they are investing, it is too easy to hide a complete lack of measurable results. Even a significant investment, when made in the context of the current incentive system, should not be expected to make a meaningful improvement.

## Transparency

Instead, regulators should demand transparency that gives a picture of how the system is behaving. What matters to understanding a network is not content creation but the distribution and engagement of individual pieces of content. It is impossible to understand the network without understanding the algorithms that govern the spread of content and drive interactions on that content. Understanding this requires more transparency than companies have historically shared.

If you want to understand what is happening more clearly, I recommend assembling an independent group of researchers and data scientists. Task them with coming up with the right questions to ask and a list of data they need to answer those questions. Then, fund a team of third-party individuals to constantly analyze this data, publish their findings openly, and recommend new questions to answer and the resulting datasets required.

The government already holds the private sector accountable in other technically demanding fields. Examples include building code inspections, NTSB crash investigations, FAA aircraft certification, and SEC financial crime enforcement. Consider also the role that third-party accounting firms play in generating audited financial statements of public companies. A functioning stock market is impossible without assurance that companies' financial statements are accurate. At the same time, companies need to keep proprietary information confidential. This is solved by third-party accounting firms that can audit and certify company financials. These firms successfully balance the public's need for accountability with the company's need to protect private and confidential information.

Due to their current incentives and how lucrative the status quo is to them, there will be pushback on any requests for more data. This pushback will come in the form of arguments

about user privacy, company confidentiality, the expense of these requests, and a myriad of other concerns. There are straightforward answers to these objections, and they can be overcome if someone is genuinely incented to overcome them.

The more complex the data requests, the more opportunities companies will have to obfuscate their answers or figure out how to make them unhelpful. There are many ways to do this, such as sharing numbers without sufficient context, sharing summary statistics such as averages when we look at highly skewed distributions where simple averages are not very informative, and many other obfuscating techniques. We should not underestimate the strength of the companies in figuring out how to avoid such requests for data. To understand their power, consider that just one team at one of these companies (the legal team at Google) has approximately the same number of lawyers as the total employee count of the entire Federal Trade Commission. This isn't even considering the amount of money they spend on external law firms.

So, as important as transparency is, it is necessary but not sufficient to tackle these problems. Other policy and legal experts have testified to various proposals that could change these companies' incentives, and I defer to those experts on that front. But until this is done, and until we change the fact that virality, attention, eyeballs, and clicks are the things the companies care about above all else, all the data sharing in the world will not address the problem.

## Incentives

One way to measure whether incentives are changing is to compare companies' behavior when they believe content quality does matter to their financial performance against cases when they do not. One simple litmus test would be how companies handle advertisers' content vs. non-advertising (organic) content. If propagated by advertising messages, hate speech or inciting violence would be terrible for the company and likely to harm their core business. Having offensive ads, or ones that violate their ad policies could quickly and materially harm their financial performance. This is why most advertising systems I am aware of place ad review (automated and sometimes manual) as a step that has to occur before the new content ever makes its way into the live system for consumption by real users. Contrast this to how organic content goes live. Organic content is allowed to go live and potentially go viral instantly. Only after the fact, when and if the team gets around to it, is there any review or take-down action. In many cases, societal harm has already been done, while any interim set of impressions or engagement is good for the company. Today, if a company changes its behavior to prioritize content review before new content goes live, that would harm its financial and user growth performance. Of course, this would not be true with the right incentives.

To give an example of how incentives can change what content is recommended, let me share the following example. It is public information that the Chinese government is an investor and board member in Bytedance, Tiktok's corporate parent. Of course, the Chinese government has a series of incentives other than profit maximization and user engagement. As a result, I have come to understand the following: the Tiktok algorithm pushes educational science, engineering,

and math content on Chinese youth while pushing a feed containing twerking videos, misinformation, and other destructive content to US children. In addition, they enforce daily time limits on usage for Chinese students.<sup>2</sup> Even worse, US companies' profit maximization motives mean that they are essentially doing the work of the Chinese government for them.

## Conclusion

Any suggestion for more data to share, or checks put in place before content can go live, will be met with strong objections. While this would certainly be an investment in computational power and software and have consequences that need to be considered, it is possible, given the resources at these companies' disposal, if they had the will to do so. Certainly, they have proved responsive to regulation by other governments (for example, in Europe). Many privacy and confidentiality concerns could also be addressed in a way that is not overly demanding. As seen with the advertisement content example, companies happily change their behaviors when it is in their best interest. That said, as long as their incentives are aligned to fight change, they will have endless reasons to oppose any such change.

Given what we know about companies' incentives, internal processes, and culture, we should not expect meaningful progress voluntarily. They will continue to benefit from the rampant spread of viral content online without feeling any significant downside from any harmful effects it may have on society. With the proper transparency, third-party oversight, and regulatory environment, I hope they will start earnestly tackling and making real measurable progress against these problems they have been so integral to exacerbating.

---

<sup>2</sup> <https://www.cnn.com/2021/09/20/tech/china-tiktok-douyin-usage-limit-intl-hnk/index.html>