

The Unseen Teen – Lenhart and Owens – 2021

Summary

While the user may be the constant focus of social media and gaming in one sense (engagement), there is less attention put toward the wellbeing of young users. This report asks, 'how do social media platform companies think about and design for the wellbeing of young people?'. It is a longitudinal, qualitative research project, interviewing various tech workers.

The key outcomes are:

- Many state that companies treat adolescent use of the platform as an afterthought. Given that they are a large user group, this is said to be a lost opportunity to aid in particular developmental needs.
- Defining wellbeing and harms from social media is difficult and lacks consensus. Many focus on screen time rather than more embodied forms of health. There needs to be greater focus from companies and regulators on reducing potential harms for specific minority subgroups.
- Companies design for an imagined 'average' user, missing the full range of user experiences. Designing for the average user can mean it fits no one, while designing around smaller subgroups can benefit everyone.
- Many companies use strategic ignorance to abstain from responsibility for the negative impacts of their platforms. This responsabilises subgroups for their own management of wellbeing.
- Company structures, cultures, or incentives don't promote a focus on user wellbeing. There tends to be a focus on exponential growth, with wellbeing only becoming a focus once harms to adolescents have already become public.
- Actors within and outside of tech companies can improve adolescent wellbeing on social media platforms. Within companies this can include: creating policies for adolescent users to empower them to learn and rehabilitate amongst other young people (via age gating); developing with real people in mind and a diverse team; integrate wellbeing expertise into every point of the design process; hiring a diverse workforce. For regulators and civil society this includes: applying pressure to these companies to address the issue of wellbeing; well thought out regulation without narrow definitions of harm.

Introduction

While the initial phases of social media and the internet were largely very positive, in more recent times we've come to ask what the health effects (mental and physical) are of all of this time online? And what responsibility does the platform have to their users? This is often particularly so for young people, who have different needs than adults, due to their status as 'vulnerable'. This is borne out in much of the legal policy around social media. The focus of research, equally, tends to be on adolescents and other users, rather than the companies who create the products. This project examines sites of production, as digital wellbeing is not just down to user behaviour. 25 former and current tech company workers were interviewed here, from both small and large companies. All companies were either a social media platform or a social gaming space. Interviewees spanned across the company, from legal to engineers, etc.

Adolescents deserve special attention in social media product design

Adolescents are in the process of becoming adults, learning and testing norms about how to interact with others and form their identities. They are also legally classed as minors, necessitating more thought into the protection of their online selves. However, designing for adolescents can be beneficial for everyone, essentially based on the principles of universal design, in which edge populations are prioritised to ensure everyone can participate. This is not to say that adolescents are the only subgroup, obviously, and this research is likely applicable in many ways to broader design processes. Adolescent here is used to mean someone between 13 and 17 (as is teen and youth); child is under 13; young people or minor is under 18 generally.

Digital wellbeing is hard to define and won't mean the same thing to every adolescent or population

One major issue here is trying to define such a contested concept as health/healthy use of technology (and so unhealthy use as well). The likelihood is that the answer is contextual and depends on *how* the platforms are being used, not just that they *are* being used. Equally though, defining wellbeing in general is very difficult. Even meta-analyses of wellbeing and digital media struggle to find any consensus (Orben, 2020). Many researchers focus on a range of measurements as a result e.g. mental health outcomes; social support and connections; life satisfaction, etc.

Digital definitions also include issues like equity and justice, based on the research subject here. Some interviewees felt it was about users being able to express their true selves on a platform or felt their time on a platform was beneficial/fulfilling. For adolescents in particular, the idea of encountering new people and ideas was also seen as a means of digital wellbeing. Some question the idea of there being 'digital wellbeing' as a separate concept from just 'wellbeing'.

Overall though, no consensus was found on what counts as healthy or who can define healthy, particularly in regard to technology use. There was a recognition that these concepts are highly subjective. For this research, healthy technology is seen as technology that improves (or doesn't diminish) mental, social, and physical health of all of its users. Digital wellbeing is defined as a sense that interaction with an online space can be a positive force for a user's mental, social, and physical life, as well as offering the user agency to manage their use of the platform. This was not the understanding of all participants but rather the researchers. Instead, companies often use the conceptual contesting as a means

to accept a narrow, metric-driven understanding of wellbeing which focuses on time and attention to a platform, as these things are easily measurable.

Current approaches to healthy technology for teens tend to be reductive and narrowly-focused on screen time

One of the main ways interviewees talked about healthy technology was via easily measurable concepts which act as proxies for digital wellbeing. This is the usual ideas of screen time, addiction, attention economies, etc. which comes from the Center for Humane Technology. Many participants talked about this and compared it with in-person activity. However, this is obviously a simplistic rendering of what it means to use technology in a healthy manner. This was recognised by some participants, who recognised not all screen time is created equally e.g. some people may benefit from longer screen time if they have supportive online communities. Basically, the traditional, quantitative measures are not enough to account for the impacts of technology (both positive and negative). Concepts of digital health and technology use have to be multi-faceted and based on engagement.

The risks and harms of social platform use are complicated, but necessary for tech companies to address

How can we better conceptualise the risks and harms of social media use then? Social media clearly does connect and inform us but it can also distract us, and compromise our mental and physical health. Again, young people are focused on here because they are in an ongoing developmental process so can be subject to the most harm. Within research, there is little consensus around the degree and nature of harm, even with behaviours suspected to be problematic (Livingstone, 2013).

There is research to show that different adolescents experience things differently on the same platforms, with the same features, and experiences. There tends to be small groups upon which negative factors for wellbeing have strong effects. We should likely be looking specifically at groups and individuals who may be at risk then, rather than user averages. This requires companies doing so as well, to develop efficient harm-anticipation systems. This includes a better understanding of how teenagers use platforms, which is often different than how developers intended. Companies cannot predict these things but they still have a moral responsibility for the outcomes of their platforms. Best practices to deal with wellbeing issues are outlined in a later section.

Methods

‘The data presented in this report come from 25 semi-structured, in-depth interviews with current and former employees of social platform companies. Participants worked at a range of small, medium, and large social media and social gaming platforms... We sought to include participants from a diverse range of genders and ethnicities. While we were able to speak with a large number of women, we were less successful in recruiting people of color, especially Black participants’.

Key Findings

The research basically wanted to see how workers from the platforms integrated concern for adolescent wellbeing into design, features, and products. However, the most notable response here was that these concepts are only brought into the research and design process at a very late stage (if they are considered at all). This was common enough to raise the question of *why* adolescent wellbeing was often ignored, despite user health being publicly prioritised by many Silicon Valley companies?

This key findings section will examine how adolescent (and general) users are overlooked by social media and social gaming platforms. Then mechanisms influencing the overlooking of adolescent wellbeing are investigated. The first is that companies tend to design for the 'average user', meaning young adult, white males. Other issues such as business models, lionisation of metrics, and lack of diversity in tech teams all contribute to the emphasis on the 'average' user. Secondly, companies use strategic ignorance to avoid difficult topics around adolescents. By not recognising adolescent use of their products, companies can abdicate responsibility for them, while burying information under piles of data. Thirdly, workers who do consider adolescence and user wellbeing are made out to be 'blockers'. Finally, how the platforms and their leaders understand their own mission (including profitability) changes how adolescent wellbeing is prioritised.

Adolescents are an afterthought at most social media companies

The interviews suggest that adolescents are rarely considered in the product design process, with this choice often being conscious. Some participants stated they avoid designing for *any* specific social group. This is the case with platforms even when adolescents make up a large proportion of the user base, with companies not seeking out adolescent users, thus not designing for them. Alongside this, companies do not collect data on adolescents specifically, making them harder to see within their systems. Despite this, we know that large numbers of adolescents use social platforms and we cannot pretend to be ignorant of this.

Laws like COPPA have driven some companies away from providing child-centred products or to only narrowly manage compliance. Essentially, concern for adolescents is taken out of any design aspect and moved solely to a legal role. Similarly, the difficulty in recruiting those under 19 for product research is pointed to as a difficulty in designing for these uses. Other companies seem to avoid designing for children, having seen large companies like YouTube already struggle with this. Despite all of this, there are real consequences of adolescents using platforms (especially when they are not considered in the design process). Teenagers will use these platforms regardless. Even well intentioned designers are seen as creating features with negative mental health impacts for adolescents due to not considering them in the design process.

Companies design for their imagined average user, which leaves out youth and others

Designing for the average user can leave out any subgroups who do not fit this mould. Here, it was seen that workers mostly thought about early adopters or tech savvy users like themselves (generally adult, high income, white males). This underpins many of the flaws in the product design process. Some participants did say they seek out experiences from groups who aren't traditionally early adopters to combat this. Conversely, at another company, the average user was actively cultivated, in the hopes of treating everyone the

same. So it is not a simple case but it seems that it depends on the individual company and their culture. Three main reasons were identified for how the average user becomes constructed: business incentives, metrification, and a lack of diversity in the workforce.

Business models and incentives contribute to ignoring important subgroups

Business models obviously shape decision making within companies. For example, accepting VC funding means you need to provide growth and ROI. This is felt to an ever greater extent after an IPO. If a company uses an ad-based revenue system, they are based upon having as much attention from users for as long as possible. Metrics like daily active users (DAU) and monthly active users (MAU) can be measured and sold to advertisers who pay proportionately. These growth metrics were accepted by product managers and designers, who recognised their ultimate goal was to build revenue, rather than a product. This still meant focusing on the largest and most financially desirable groups, missing nuances and harms faced by smaller groups, like adolescents.

The centrality of growth metrics promotes viewing all users as the same

Growth metrics treat users as a monolith, rather than recognising specific subgroups with specific needs within the userbase. Research with adolescents has shown differential susceptibility to harm (as well as positive outcomes) from social platform use. Treating users as a monolith makes it harder to see them as individuals for platform employees. A/B testing is highlighted here as a specific means of treating users as aggregates and not considering what can be negative for smaller groups. Even when companies do work on improving user wellbeing, these efforts are made secondary by larger, more quantifiable goals (generally things which are more directly monetizable). Again, the average user is designed for and this tends to be reflective of the workforces in tech.

Lack of diversity in tech workforces limits who is imagined as a user

Workers view the needs of the community based on their own needs, which can be a problem with a lack of diversity existing within tech companies. Men are still overrepresented, as are white and Asian people. Black people are particularly underrepresented. Many workers also lack children or only have very young children. This research (and others) indicates ‘that if tech companies continue to primarily employ young, white men to design and engineer their products, those products will be disproportionately designed for the needs and experiences of that population’.

Many negative health impacts stem from what companies choose not to know about their users

It is also important to examine which types of data companies ignore. What data *not* to collect is as much of a choice as which data *to* collect. This is strategic ignorance, with companies choosing not to know things about their users to avoid responsibility for potential issues and a need to address these issues (ethically and legally).

Three main methods of strategic ignorance in company structures are laid out here: firstly, not collecting data. This allows companies to feign ignorance of an issue or argue that the problem does not exist at all. This may be not collecting age data about users, which allows companies to ignore issues or adolescent wellbeing. Many companies are able to

document adolescent use of social media though, showing that it doesn't have to violate privacy protections; secondly, collecting too much content. This makes relevant data difficult to find or understand. If a company has a large dataset, they can 'miss' content that highlights issues of harm to adolescents. However, companies encourage a high pace and scale of content posting on social platforms so this amount of data collection is a choice they made. It is suggested here that more content moderators are needed, better automatic flagging, removing of content, etc.; finally, unclear lines of responsibility for knowledge, distributing ignorance across an organisation. With enough division of labour, no one person or team can be seen as ultimately responsible for a problem. The blinkering of teams within a company makes it hard to see a bigger picture of what the team is working on and so mitigates ethical and health related issues potentially in the product. If issues are identified, they are only addressed to act within the bounds of what is legally permissible, especially for adolescents.

All non-majority groups face harm from this strategic ignorance. Essentially, a lack of knowledge can lead to discrimination as well as negative impacts on health. Cultural nuances and slang are not recognised too, resulting in subgroups being flagged (black communities being 1.5 times more likely to have content flagged than any other group, for example). These subgroups are also subject to higher rates of hate speech (which is, presumably, not being moderated as much). Similarly, there is a large bias towards American users (although the quote here specifies California), even if a platform has a large international userbase. And adolescent girls are not as visible to social media companies either, again reflecting the homogeneous makeup of much of tech. The most notable absence here was lower income users and users without access to fast internet or new technology. This is important now with distance learning but equally important when discussing wellbeing because suggestions of disengagement from technology often assume the upper middle class problem of abundance.

Some participants suggested that when health of subgroups was considered, it was a PR move and aimed at reducing legal liability or regulation. Again, users are only a means to an end, they are not prioritised in the actual design process.

Tech workers who explicitly focus on user wellbeing often have little power to change product design and processes

Despite all of the above, many companies still have teams designated to think about wellbeing of users, sometimes specifically for adolescents and children. However, they are often only able to flag issues rather than change any product direction. They are siloed (tending to fall within the Trust and Safety team) or their recommendations are ignored. They do not want to be seen as blockers of innovation or revenue; instead they seek to illuminate problems which may or may not be taken onboard. In the tension between internal company interests, it seems wellbeing staff rarely win. In companies where wellbeing staff are integrated into design teams, however, they are often seen as a collaborator rather than a blocker. But this is a minority, as stated.

Platform self-concept influences the choices organisations make about how to build and manage their online spaces

What companies think their platform is for and what value it provides to the world shapes the choices made within the company. Some sites see themselves as a space for fun or pleasure, being more willing to remove content as a result. Companies which see themselves as a digital commons tend to privilege free speech and have it at the centre of their rule making.

Interactive gaming platforms are seen as thinking more seriously about community interactions to diminish user harm. This may be due to the interest in building an immersive world, which requires a positive and inclusive atmosphere. Social platforms who focus on adolescents directly see themselves as responsible for building communities with digital citizenship as well. These companies tended to take a much more pragmatic viewpoint of how and when teenagers use technology today. Companies which primarily prioritised shareholder desires found it very difficult to develop products which also prioritise users. This idea of the self-concept shows the platforms not to be neutral spaces but having specific ideas about themselves and their goals.

Leadership sets the tone for organisational culture, values, and decision making

Obviously the leadership of a company is influential in how others in the company make decisions. Those at the top set the tone, corporate culture, and how much a company focuses on growth and revenue compared to other metrics. The normal focus on key performance indicators (KPIs) takes away from the human aspect of the product, according to one participant. Leadership often does not really trust their userbase either, so issues like adolescent wellbeing are taken even less seriously.

From the beginning as a start-up, few founders spend time worrying about what could go wrong. There are signs this could be changing, with more focus on content moderation and social responsibility than in the past. Some participants felt that a newfound openness to change (and to being wrong) was a positive in a CEO, but could also suggest that up until now they had not considered any negatives of their platform. Overall though, participants were clear that it was difficult to build support for a focus on wellbeing from the ground up, that it has to come from the top and be a part of the corporate culture.

Advertisement-based business models may be fundamentally in tension with adolescent harm and wellbeing

Participants identified ways in which company cultures, policies, and processes de-prioritised adolescent health and wellbeing but also showed how business models put companies fundamentally at odds with user wellbeing. If a company is based around user engagement and attention, it can be very difficult to change any processes without risking financial disaster. This (alongside pressure from investors) pushes features which are monetizable, rather than beneficial for the user. This was highlighted for games development too, due to their tight budgets and timeframes.

Tech workers saw this as a structural tension, difficult to change without larger, systemic changes within capitalism. Issues like inequity of access were seen by workers as issues which need broader social change, rather than platform (or sector) specific change.

Basically, not everyone can afford a subscription for a product so ad-based revenue becomes essential. Others did believe that prioritising user wellbeing would actually benefit the bottom line for companies. This is seen as having improved over time as well, with Trust and Safety being more involved in the design process now compared with a decade ago.

Regardless of the changes needed to capitalism, social platforms can still choose to prioritise adolescent wellbeing, as a compliment to their long term growth and monetisation. But adolescents remain an oversight for companies, within the design process. Three main reasons are reiterated here: the focus on the imagined user (tech savvy, young, white, male); strategic ignorance about use of the platform; and company structures/business models making it difficult to prioritise wellbeing. Two questions remain here: what can be done by outside regulators to push for change in social platforms? And what should companies (and workers within companies) do to change their organisations for the better?

Levers of change and best practices

Levers of change for regulators and civil society

One of the goals of this research is to understand the development of social games and platforms, including the impact of the development processes and practices on adolescent health. This includes what drives change within companies to think more about digital wellbeing beyond that of the average consumer.

Outside pressure works

Participants stated that negative media attention, civil society, and regulators push organisations to change their behaviour; there is a fear of negative PR.

Negative publicity – and tragedy – can force company changes

There were specific moments when adolescent experiences came to the fore for the companies involved, usually when there was a publicised tragedy. Only hard evidence of harm to a child can convince these platforms to change priorities sometimes. The argument here is that it should not take such an event in the first place for platforms to take adolescent health seriously.

Other than this, the threat of negative media coverage was a major driver for change. Reputation impacts their revenue so they have to care about it. This means certain journalists can have a lot of power for change. Therefore, there is an importance placed upon the press to take nuanced stances in their reporting of digital wellbeing. This includes going beyond surface level takes about screen time, for example.

Public failures at other companies also spur action and coordination

When one company takes a big public hit, workers and civil society can use this failure to drive change more broadly. One example here is inBloom, a personalised learning technology intended to mine educational data to improve outcomes. It was met with much public and parental resistance about the privacy and use of children's educational data, resulting in collapse. However, it can often take more than one company suffering a failure to really change corporate accountability practices around wellbeing. Newer start-ups are more likely to take note of contemporary failures due to their inherent precarity. Similarly, huge crises within society can push change e.g. the Christchurch shooting. Despite this,

non-profits and advocacy groups should continue to push for change without needing a failure.

Pressure from civil society can shift company decision-making

Companies have become more pro-active about seeing features which could get negative coverage, part of which includes getting partnerships with advocacy groups and the like as a form of external validation for their platform. These external advocates can therefore build off of companies' concern, producing reports, rankings, etc which can shape internal and external debates around platform trust and safety. One participant clearly says that platforms need very specific and actionable feedback, which these advocacy groups are able to provide, due to their expert status. Internally, employees who seek change in their company are able to point to the findings from these advocacy groups, which will be taken seriously as it is seen as hard data.

Regulations can help but can also have unintended effects

Participants often talked about regulation (and the threat of it). Companies take this seriously and seek to be prepared for it and lobby for their own interests. Legal judgements can drive changes within companies as an imperative as well, demanding certain oversight by companies be taken.

Focusing just on regulation and legality can ignore broader discussions of toxicity though. This point essentially says that platforms should be concerned with more than just generating clicks and should actively moderate (and censor) some subjects like conspiracy theories.

Similarly, if legislation is written or executed poorly, the outcomes can be contrary to the intention. COPPA is mentioned here specifically, as 'an abject failure' by one participant. Legislation can drive companies away from a space too if it is deemed too risky, creating 'reg tech' spaces (with high cost of entry and risky ROI) e.g. FinTech, EdTech, and spaces aimed at adolescents and children. The report therefore concludes here that we should learn from failures in regulation and anticipate future negative outcomes of any policies.

Best practices for companies and workers

What should companies do to ensure their technology minimises harm and promotes user wellbeing? Due to health still being a very vague term, there would need to be a considered approach by individual companies, rather than a one-size-fits-all approach. Companies need to be agile in responding to new situations as they arise. This includes proactively thinking about adolescent and user wellbeing in product development as an ethical decision but also as a value-added decision for users, resulting in long term profit and lower PR risk.

Create targeted approaches and policies for adolescent users

Focus on empowering, rather than protecting adolescent users. This includes consulting with young people to understand their perspective on the platform. Important to this is moving away from parental consent models for adolescent technology use. The technological underpinnings of platforms make getting informed parental consent unfeasible and are easy to circumvent anyway; it is basically a form of feigned ignorance on the part of the platform. Instead, companies should work to make their whole platform or system safer for younger users.

Allow for adolescent learning and rehabilitation.

We treat children differently than adults because we recognise they are developing and are in the process of becoming more accountable over time. Therefore, they should be allowed to make mistakes. Companies should build systems that allow for adolescents to make mistakes while keeping them accountable for their behaviour. Actually teaching children about what they may have done wrong and how to be better might not be in the purview of companies though; it may be better suited to outside groups like parents or educators.

Intentionally design and promote the desired environment and culture on a platform

Platform cultures and communities develop from design and marketing choices, not always intentionally. Again, social gaming platforms seemed to consider this more in terms of product design to promote healthy environments for adolescents. This is important because social platforms are not neutral technologies, meaning companies can (and should) work to build safe and healthy environments. This can incentivise people to return to the platforms more often and spend more time, generating more revenue.

Remind product teams of the humans using the platforms

Consider a broad range of user subgroups when developing new features and products. Tech workers should not assume all users are the same or like themselves; they should collect data to determine which groups need more focus. This will be context dependent on each platform. Workers should speak up if they think a feature is inappropriate or causes harm.

Integrate expertise about user wellbeing into product teams and across all roles

Ensure that employees with expertise in user health and wellbeing are integrated into product teams. They will serve as internal experts and consider various viewpoints. This research suggests doing so produces better outcomes and fewer failed launches. The aim here is to build product development infrastructure which inherently considers wellbeing across user types. This should then centre wellbeing from the inception, rather than being tacked on at the end. Currently, many companies only seem interested in wellbeing in an abstract manner, rather than empowering workers to actually change product design to reflect a concern around wellbeing.

Hire, retain, and empower a diverse workforce

This particularly applies to product development, design, and technical roles. This research and others suggest more varied teams produce better results for users and profits. There should be internal infrastructures which allow employees to voice their concerns as well, to not be made to feel like blockers, or to feel precarious (particularly if they are not white, middle-class, male, etc). Again, incorporating wellbeing from the design process at the beginning would make this easier, as it would not necessitate many changes being made when the product is about to go gold. This needs to go alongside an acknowledgement of systemic racism in the tech industry and beyond. This may include prioritising minority communities in product design and research.

Unions are suggested as a positive step here too. Workers increasingly find themselves at odds with the decision makers within their companies. This acknowledges that companies are hostile to this and the Alphabet Workers Union required many walkouts and protests but encourages companies to nurture unions nonetheless?

Require training in ethics and humanities in all staff and new hires

‘This training should deeply engage with a wide range of thinking, perspectives and issues, and connect to an organization-wide commitment and belief in the importance of diversity and ethical work’. This is intended to make companies more effective at recognising structural biases. It also suggests more humanistic and social science course work across technical fields.

Consult experts, civil society, and diverse groups of users to anticipate effects on users prior to launch

‘Assume responsibility for the disparate impact of products on different subgroups of users, regardless of the intent behind them. Work to anticipate and avoid potential negative effects by consulting experts and members of those user subgroups’. This includes making real amends for any actions which do go badly. There seems to be a growing recognition amongst social platform companies that they cannot continue to ignore harms to their users caused by/on their platforms.

Collaborate with industry partners to mitigate harm across platforms

People within organisations should work together on similar issues they face, including sharing information related to bad actors and wellbeing. One suggestion is a repository of ‘bad words’ which should be moderated or banned. Sharing with others is already done through informal relationships so professionalising this through third party groups could provide a place for public and private individuals to work together more effectively. Working together should include civil society as well. Regulators and advocacy groups can put pressure on companies or threaten to regulate. Many tech employees want to improve their platforms so would benefit from outside thinking as well. This would be most effective in an ongoing manner because issues affecting adolescents are not static.

‘This research and report was supported by Pivotal Ventures, the investment and incubation company created by Melinda Gates’. This makes complete sense and is indicative of the whole thing.