EXHIBIT 79

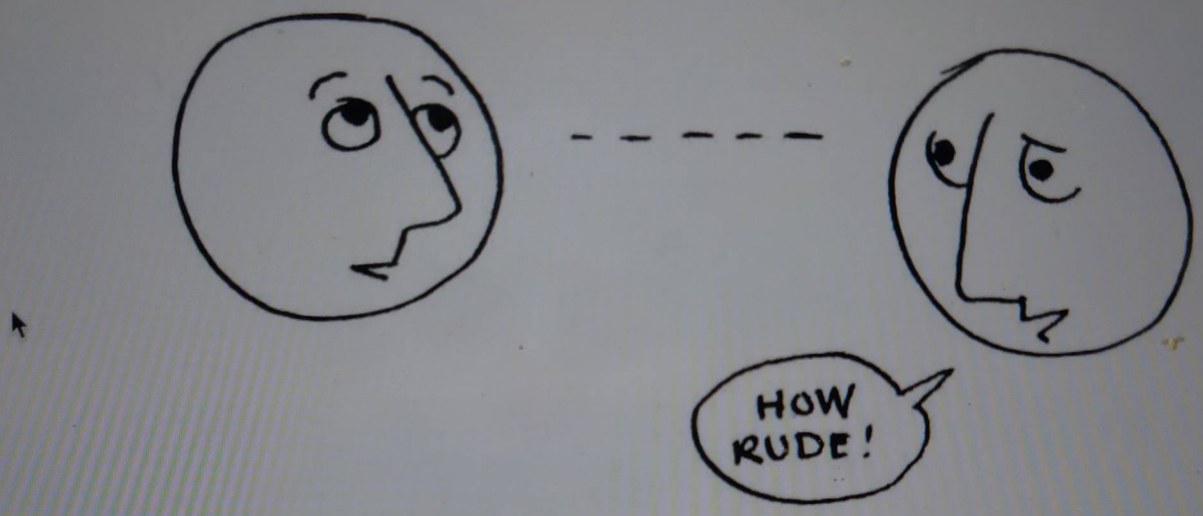

When User-Engagement ≠ User-Value

TOM CUNNINGHAM · THURSDAY, DECEMBER 10, 2020 · READING TIME: 11 MINUTES

(originally written in 2019, published in Dec 2020).

Three broad reasons why increasing user-engagement might not increase user-value:

- **First, when the product *stretches* value out** – i.e., it spreads value thinly, to maximize engagement.
- **Second, when the product is misleading** – i.e., we make promises about value without fulfilling them.
- **Third, when users' choices don't represent their own be** when they have self-control problems

- **Third, when users' choices don't represent their own best interests** – e.g. when they have self-control problems.

I think that all of these are important issues for our products.

Q: How to make product decisions that maximize user-value instead of user-engagement? I don't have a silver bullet, & just have a few observations.

- **Choosing engagement metrics carefully can help** – e.g., maximizing sessions, instead of time-spent, has some positives for the "stretching" problem, although it also may have negatives with respect to self-control issues.
- **Monitoring sentiment metrics can help** – but they're often too noisy to help.
- **A strong quality culture probably helps.**

Background

1. **There are a lot of critiques of social media that we maximize engagement at the expense of value.**.
2. **We are providing value.** People spend a huge share of their time on Facebook apps – something like 1% of worldwide time-spent – it would be surprising if we weren't delivering some positive value. When you ask people to put a monetary value on Facebook it's high: the average American says it would take $50-100 before being willing to give up Facebook for a month.
3. **There are reasons to be concerned.** Most importantly the relationship between usage of Facebook and psychological wellbeing is pretty murky: higher FB use is associated with worse wellbeing, in the cross-section (though not in the time-series), and in a recent experiment they found that a 1-month break from

3. **There are reasons to be concerned.** Most importantly the relationship between usage of Facebook and psychological wellbeing is pretty murky: higher FB use is associated with worse wellbeing, in the cross-section (though not in the time-series), and in a recent experiment they found that a 1-month break from using Facebook caused a small but significant *increase* in self-reported psychological wellbeing.

4. **Why would engagement and value diverge?** Basic economic reasoning would imply that people engage with a product to the degree that they get value from it, so engagement and value should be closely connected. I think this is broadly true – but there are a few *wedges* where value and engagement can diverge.

(1) Stretching Out Value (AKA Spence distortion)

Spreading the jam thin. Suppose you have a fixed amount of value to offer but you're trying to increase engagement, e.g. time-spent, then you can do this by spreading the value out: cut the value up into chunks, and put some filler in between.

A couple of economic papers in the 1970s by Michael Spence formalized this point: if you're a firm making choices about the design of your product then your incentive is to maximize the *quantity* of demand – i.e., the number of people who want it, or the time-spent using your product – rather than the *quality* of demand, i.e. the total value delivered. (More precisely: your incentives are around increasing the *marginal* value, not the *total* value).

Spence applied this reasoning to television, comparing cable TV to broadcast TV. Cable TV is funded by a fixed fee, so they have an incentive to maximize overall value. Broadcast TV is funded by advertising, so they want to maximize time-spent, AKA engagement. For this reason there's an incentive for broadcast TV to stretch out their value in a couple of ways: (1) if they can make a show more broadly appealing

Spence applied this reasoning to television, comparing cable TV to broadcast TV. Cable TV is funded by a fixed fee, so they have an incentive to maximize overall value. Broadcast TV is funded by advertising, so they want to maximize time-spent, AKA engagement. For this reason there's an incentive for broadcast TV to stretch out their value in a couple of ways: (1) if they can make a show more broadly appealing, at the expense of value to the core fans, that is always in their interest; (2) if they can artificially increase viewing time on a show they'll do that – by having constant recaps, and padding, and cliffhangers.

Stretching value occurs in our products in a few ways.

1. **We limit the amount of information in notifications because we want people to come onto the site.** A few years ago we stopped sending out emails telling you what happened – e.g. telling you what your friend did – instead we just say "someone comment on your post," in the hope that you'll click through. This is a clear value-engagement tradeoff. We see similar effects on push notifications about friend activity: when we send fewer notifications, that *increases* engagement because people come to the site to check what's happening.
2. **Lots of content on News Feed withholds information.** With titles like "you'll never guess...", "one amazing trick...," they try to get you to click through, even though they could've easily told you the information up-front. We currently have a ranking penalty for egregrious cases, but you can still see the effects of this incentive everywhere, even high-quality newspapers craft their headlines to leave out some tantalizing detail.
3. **Maximizing for time-spent incentivizes bad ranking.** If we evaluated ranking by the effects on time-spent then it would *not* be optimal to rank just by expected value – instead we would be better off by adding some randomness or filler, so that the value is stretched out. It's hard to tell how important this is for ranking in practice, it's a subtle thing to measure and ranking is a black-box, but I'd guess it has a significant influence.
4. **More friction for high-value things.** Suppose that someone

3. **Maximizing for time-spent incentivizes bad ranking.** If we evaluated ranking by the effects on time-spent then it would *not* be optimal to rank just by expected value – instead we would be better off by adding some randomness or filler, so that the value is stretched out. It's hard to tell how important this is for ranking in practice, it's a subtle thing to measure and ranking is a black-box, but I'd guess it has a significant influence.
4. **More friction for high-value things.** Suppose that someone really wants to do something – e.g. there's a video that they really want to watch – then any feature or experiment can add friction without having much effect on engagement – e.g. we can add more ads, interstitials, promotions. Again, these types of things will hurt consumer value without having a proportionate impact on engagement.

(2) Misleading Users

A second reason for divergence between engagement & value is misleading experiences. If you give the impression to users that they're going to get some value by engaging with your product, but you don't actually give users that value, that will introduce a discrepancy between engagement & value.

This happens with Feed links all the time. Headlines will say misleading things to get you to click, subtly exaggerate. We have a classifier and demotion for exaggerating headlines, but it's a very crude solution.

If we blindly followed AB tests to maximize engagement, we'd lie to users all the time. The simplest way to maximize engagement is just to lie to users, e.g. by saying "click this link for guaranteed long battery life", or "accept our terms and we promise not to track you." If an AB test has a surprisingly positive effect on engagement you should double-check that the feature isn't somehow inadvertantly misleading users.

Why doesn't this process self-correct? You'd expect that consu

Why doesn't this process self-correct? You'd expect that consumers will learn when firms mislead or exaggerate, and adjust their expectations, giving firms an incentive to tell the truth. This doesn't work perfectly in practice for a couple of reasons: (1) the cost of disappointed expectations is shared – when users have a hard time keeping track of who's responsible for the content they're seeing they'll blame *all* publishers, and *all* app developers, for disappointed expectations. This weakens the incentive to tell the truth, because the costs are borne by other companies. (2) many features are evaluated just on *short-run*engagement – if we mislead users it'll surely cause a lack of trust down the line, but that can be slow and cumulative, and it's very hard to detect with just short-run experiments.

(3) Self-Control Problems

Finally there are sometimes discrepancies between what people choose and what gives them the most value. The classical economic point of view is that a person's choice perfectly represents their preference – they're the same thing. However there are a bunch of cases where that seems overly simplistic – e.g. when people have self-control problems, or cognitive problems in processing information.

There are reasons to worry about self-control and use of our products. I think the biggest problems people have are: (1) compulsive use, to fill gaps in their time, and they end up spending more time than they'd like, instead of spending time with their families, or studying; (2) unhealthily fixating on other people – most prototypically, compulsively looking up the profile of your ex-partner.

A quick rundown of evidence:

1. **Higher Facebook use is correlated with worse psychological states** (wellbeing, loneliness, etc.), although *changes* in Facebook use are uncorrelated with changes in wellbeing measures. [8,5]

1. **Higher Facebook use is correlated with worse psychological states** (wellbeing, loneliness, etc.), although *changes* in Facebook use are uncorrelated with changes in wellbeing measures. [8,5]
2. **An experiment found that a 1-month break from Facebook improved self-reported wellbeing.** [6]
3. **A large fraction of users struggle with their Facebook/Instagram use.** Among US teens "84% had taken a break from social media, and 74% have taken a break from Instagram." [2] "More than a third of the survey respondents feel they have spent more time than they would have liked on Instagram. About one in five respondents think they have a problem controlling their time." [5]
4. **A significant minority report serious difficulties.** 3.1% of Facebook users in the U.S. report "serious problems with sleep, work, or relationships that they attribute to Facebook AND concerns or preoccupations about how they use Facebook." [1]

What do Do?

Despite lots of work I think it's still an open problem, of how to make sure that maximizing user-engagement doesn't harm user-value. I'm planning to spend some more time thinking about this, in the meantime here are a few observations:

1. **Choosing engagement metrics carefully can help.** – over the last couple of years News Feed slowly switched from maximizing time-spent towards maximizing sessions. This should be good from the point of view of the Spence distortion: the frequency of sessions should be roughly proportional to the total expected value, so there's less incentive to *stretch* value out to maximize time-spent. However this could be bad from the perspective of self-control problems: we actually find that the frequency of sessions is a strong predictor of problematic use, and our own internal researchers recently recommended we should "help people consolidate their use of Fa

1. **Choosing engagement metrics carefully can help.** – over the last couple of years News Feed slowly switched from maximizing time-spent towards maximizing sessions. This should be good from the point of view of the Spence distortion: the frequency of sessions should be roughly proportional to the total expected value, so there's less incentive to *stretch* value out to maximize time-spent. However this could be bad from the perspective of self-control problems: we actually find that the frequency of sessions is a strong predictor of problematic use, and our own internal researchers recently recommended we should "help people consolidate their use of Facebook into fewer sessions." [1]
2. **Monitoring sentiment metrics can help.** – e.g. CAU, WYT, etc., but my impression is that they're often too noisy to give guidance to everyday product decisions.
3. **A strong quality culture probably helps.** – I sometimes wonder what the alternative is to making decisions based on AB-tests of engagement metrics, and I scratch my head, then I remember that just over the road WhatsApp built a 1Bn user product without running any tests, they made all their product decisions just based on their perceptions of user quality. I don't feel qualified to judge on whether the outcome is better, but it's certainly worth remembering that there is an alternative. (Of course WhatsApp cares about growth and engagement, but the fact they didnt' run experiments until recently implies that they don't think short-run AB-test measures were very valuable measures of long-run value).

References

[1] Cheng, Burke, Davis, "Understanding Perceptions of Problematic Facebook Use", https://www.dropbox.com/s/bph75wxf1u27hnr/problematic_use_2019-01-07_final.pdf?dl=0

Notes

+ Write Note

References

[1] Cheng, Burke, Davis, "Understanding Perceptions of Problematic Facebook Use", https://www.dropbox.com/s/bph75wxf1u27hnr/problematic_use_2019-01-07_final.pdf?dl=0

- See also summary at https://our.internmc.facebook.com/intern/research/p/6gcvu18gbp/
- 3.1% of Facebook users in the U.S. experience "Serious problems with sleep, work, or relationships that they attribute to Facebook AND concerns or preoccupations about how they use Facebook (e.g., a fear of missing out (FOMO) or lack of control)"
- "people who reported problematic use spent signifcantly more time on Facebook than people who did not ... they spent 21.6% more time on the site ... had 13.5% more distinct sessions."
- We should "Help people consolidate their use of Facebook into fewer sessions, with more control over notifications."

[2] Darwin Mastin, 2018, "Taking a Break from Tech" – https://fb.workplace.com/groups/ResearchInsights/permalink/1824949684220305/

- "84% [of US teens] had taken a break from digital technology in the last year"
- "84% had taken a break from social media, and 74% have taken a break from Instagram."

[3] Ryan et al., "The uses and abuses of Facebook: A review of Facebook addiction" - https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4189307/

- Review finds 9 studies of FB addiction, but methodologically inco

Chats 6

MacBook Pro

Notes

+ Write Note

[3] Ryan et al., "The uses and abuses of Facebook: A review of Facebook addiction" - https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4189307/

- Review finds 9 studies of FB addiction, but methodologically inconsistent.

[5] Ruogu Kang (2018) "User Perception of Time-Spent on Instagram"

- https://fb.workplace.com/groups/432156973909421/permalink/482477452210706/
- "survey in May 2018 with mid to high time spent users on Instagram,"
- "More than a third of the survey respondents feel they have spent more time than they would have liked on Instagram. About one in five respondents think they have a problem controlling their time."
- 42% of the survey respondents have taken a break from Instagram. The top reason for taking a break is "I had something at work or school that I needed to focus on." More high TS users report that the break has made them feel worse than before than low TS users.
- People who spend longer time think their time spent is more meaningful than those who spend shorter time. Those who spend longer than 3 hrs/day have lower positive affect, lower life satisfaction, and higher negative affect than people who spend shorter time on Instagram.

[6] Allcott et al. "The Welfare Effects of Social Media"

- https://web.stanford.edu/~gentzkow/research/facebook.pdf
- "Deactivation caused small but significant improvements in well-being, and in particular in self-reported happiness, life satisfaction, depression, and anxiety ... Our overall index of subjective well-being improved by 0.09 standard deviations. As a point of comparison, this is about 25-40 percent of the effect of psychol

Notes

[6] Allcott et al. "The Welfare Effects of Social Media"

- https://web.stanford.edu/~gentzkow/research/facebook.pdf
- "Deactivation caused small but significant improvements in well-being, and in particular in self-reported happiness, life satisfaction, depression, and anxiety ... Our overall index of subjective well-being improved by 0.09 standard deviations. As a point of comparison, this is about 25-40 percent of the effect of psychological interventions including self-help therapy, group training, and individual therapy, as reported in a meta-analysis by Bolier et al. (2013)."
- (after deactivation, persistently lower use)

[7] Burke & Kraut (2016) "The Relationship Between Facebook Use and Well–Being Depends on Communication Type and Tie Strength"

- https://onlinelibrary.wiley.com/doi/full/10.1111/jcc4.12162
- only looks at *changes* in wellbeing over a month, doesn't look at

[8] Ritter et al. (2018) "A Longitudinal Analysis of Facebook Use and Well-Being"

- "people who use Facebook more have lower well-being compared to people who use Facebook less"
- No relationship between time-spent and *change* in well-being

[9] Charles Midwinter, "More time spent per session is associated with a lower risk of problematic use"

- Findings:
  is_problematic ~ sess_per_day (+.01) + controls
  is_problematic ~ session_per_day (+.007) + ts_per_session (-.0

Notes

[9] Charles Midwinter, "More time spent per session is associated with a lower risk of problematic use"

- Findings:
  is_problematic ~ sess_per_day (+.01) + controls
  is_problematic ~ session_per_day (+.007) + ts_per_session (-.003) + controls
  that, when you control for both, session-frequency positively associated with problematic use, and session-length negatively associated.
- https://fb.workplace.com/groups/558331381190793/permalink/903051943385400/
- notebook @ https://our.internmc.facebook.com/intern/anp/view/?id=132931

[10] Charles Midwinter (2019) "The Good, the Bad, and the Murky: Initial Findings on Session Quality"

- https://fb.workplace.com/notes/charles-midwinter/the-good-the-bad-and-the-murky-initial-findings-on-session-quality/484659979043924/

73 14 Comments 10 Shares

Like Comment Share Unsave

Tom Cunningham cc Nihar Shah
Like · Reply · 1y 1

Tom Cunningham Peter Tu
Like · Reply · 1y 1

Write a reply...

Michael Salem FWIW I suspect we have a lot of #1 going on (stretched out value) via d

Chats

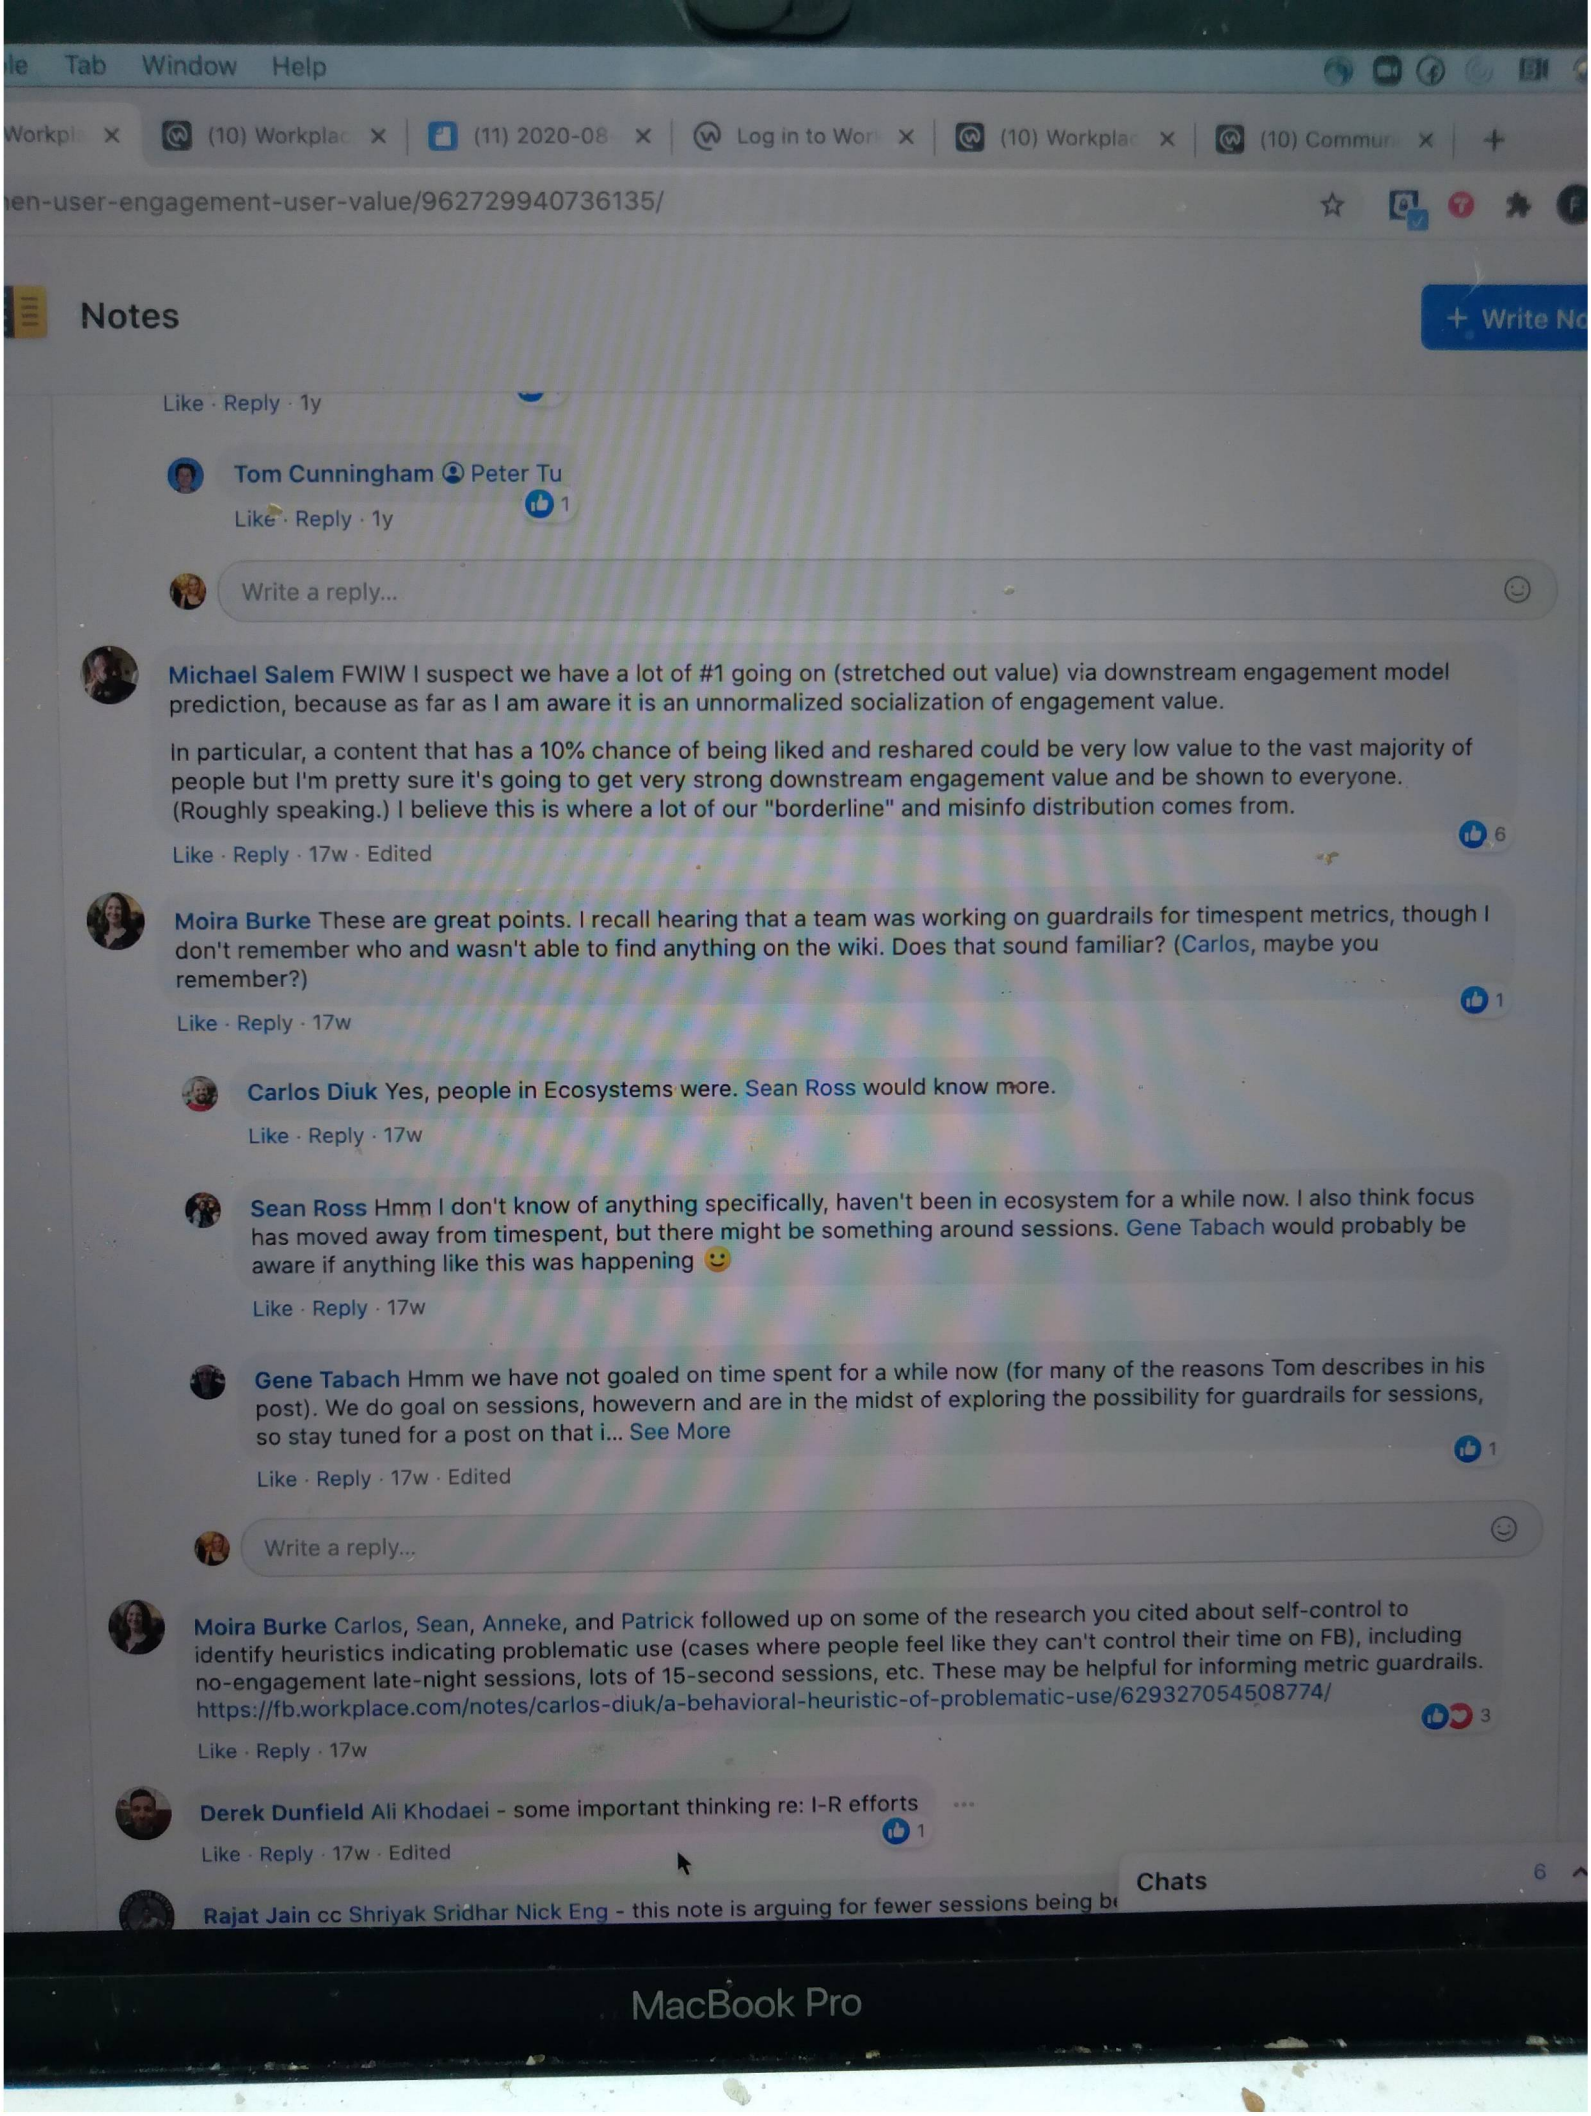
Tab Window Help
(10) Workplac
(11) 2020-08
Log in to Wor
(10) Workplac
(10) Commur
en-user-engagement-user-value/962729940736135/
Notes
+ Write No
Like · Reply · 1y
Tom Cunningham Peter Tu
1
Like · Reply · 1y
Write a reply...
Michael Salem FWIW I suspect we have a lot of #1 going on (stretched out value) via downstream engagement model prediction, because as far as I am aware it is an unnormalized socialization of engagement value.
In particular, a content that has a 10% chance of being liked and reshared could be very low value to the vast majority of people but I'm pretty sure it's going to get very strong downstream engagement value and be shown to everyone. (Roughly speaking.) I believe this is where a lot of our "borderline" and misinfo distribution comes from.
Like · Reply · 17w · Edited
6
Moira Burke These are great points. I recall hearing that a team was working on guardrails for timespent metrics, though I don't remember who and wasn't able to find anything on the wiki. Does that sound familiar? (Carlos, maybe you remember?)
Like · Reply · 17w
1
Carlos Diuk Yes, people in Ecosystems were. Sean Ross would know more.
Like · Reply · 17w
Sean Ross Hmm I don't know of anything specifically, haven't been in ecosystem for a while now. I also think focus has moved away from timespent, but there might be something around sessions. Gene Tabach would probably be aware if anything like this was happening
Like · Reply · 17w
Gene Tabach Hmm we have not goaled on time spent for a while now (for many of the reasons Tom describes in his post). We do goal on sessions, howevern and are in the midst of exploring the possibility for guardrails for sessions, so stay tuned for a post on that i... See More
1
Like · Reply · 17w · Edited
Write a reply...
Moira Burke Carlos, Sean, Anneke, and Patrick followed up on some of the research you cited about self-control to identify heuristics indicating problematic use (cases where people feel like they can't control their time on FB), including no-engagement late-night sessions, lots of 15-second sessions, etc. These may be helpful for informing metric guardrails.
https://fb.workplace.com/notes/carlos-diuk/a-behavioral-heuristic-of-problematic-use/629327054508774/
3
Like · Reply · 17w
Derek Dunfield Ali Khodaei - some important thinking re: I-R efforts
1
Like · Reply · 17w · Edited
Rajat Jain cc Shriyak Sridhar Nick Eng - this note is arguing for fewer sessions being b
Chats
6
MacBook Pro

Tab Window Help

(10) Workpla | (11) 2020-08 | Log in to Wor | (10) Workpla | (10) Commun

n-user-engagement-user-value/962729940736135/

Notes

+ Write No

Sean Ross Hmm I don't know of anything specifically, haven't been in ecosystem for a while now. I also think focus has moved away from timespent, but there might be something around sessions. Gene Tabach would probably be aware if anything like this was happening

Like · Reply · 17w

Gene Tabach Hmm we have not goaled on time spent for a while now (for many of the reasons Tom describes in his post). We do goal on sessions, howevern and are in the midst of exploring the possibility for guardrails for sessions, so stay tuned for a post on that i... See More

Like · Reply · 17w · Edited

1

Write a reply...

Moira Burke Carlos, Sean, Anneke, and Patrick followed up on some of the research you cited about self-control to identify heuristics indicating problematic use (cases where people feel like they can't control their time on FB), including no-engagement late-night sessions, lots of 15-second sessions, etc. These may be helpful for informing metric guardrails. https://fb.workplace.com/notes/carlos-diuk/a-behavioral-heuristic-of-problematic-use/629327054508774/

Like · Reply · 17w

3

Derek Dunfield Ali Khodaei - some important thinking re: I-R efforts

1

Like · Reply · 17w · Edited

Rajat Jain cc Shriyak Sridhar Nick Eng - this note is arguing for fewer sessions being better for overall experience

3

Like · Reply · 17w

Praveen Rathinavelu replied · 2 Replies

Albert Hamood A lot of (3) sounds like addiction. Addiction (as I was taught!) is characterized by a dissociation between "liking" and "wanting." The addict self-reports that they do drugs because they like it, but generally it can be revealed that addicts don't like the drug as much as they used to, before they were addicted. However they want it more, and compulsively seek it, despite suffering harm. Mapping to Tom's thinking here, wanting can be proxied by engagement, but liking is more like the value we're after. These dissociate all the time, and it can happen to anyone.

Addiction is challenging to define, and usually refers to behavior that persists despite adverse consequences. It seems clear from what's presented here that some of our users are addicted to our products. And I worry that driving sessions incentivizes us to make our product more addictive, without providing much more value. How to keep someone returning over and over to the same behavior each day? Intermittent rewards are most effective (think slot machines), reinforcing behaviors that become especially hard to extinguish -- even when they provide little reward, or cease providing reward at all.

1

Like · Reply · 17w

Noah Gaspar How do you define user value? Is it enough to say that "People spend a huge share of their time on Facebook apps"; therefore, they must be receiving some value? What's a more precise definition?

Like · Reply · 17w

Write a comment...

Chats 6

MacBook Pro