

1. THREAD: The Twitter Files

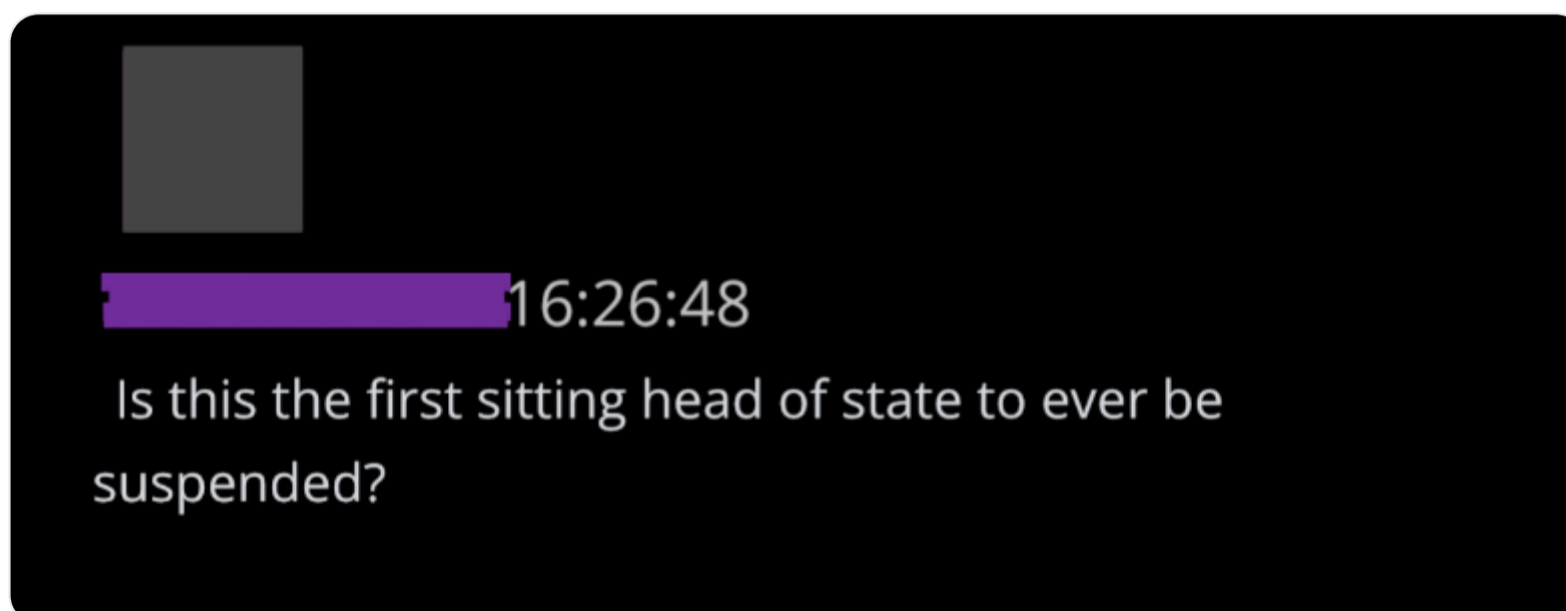
THE REMOVAL OF DONALD TRUMP

Part One: October 2020-January 6th

2. The world knows much of the story of what happened between riots at the Capitol on January 6th, and the removal of President Donald Trump from Twitter on January 8th...
3. We'll show you what hasn't been revealed: the erosion of standards within the company in months before J6, decisions by high-ranking executives to violate their own policies, and more, against the backdrop of ongoing, documented interaction with federal agencies.
4. This first installment covers the period before the election through January 6th. Tomorrow, [@ShellenbergerMD](#) will detail the chaos inside Twitter on January 7th. On Sunday, [@bariweiss](#) will reveal the secret internal communications from the key date of January 8th.

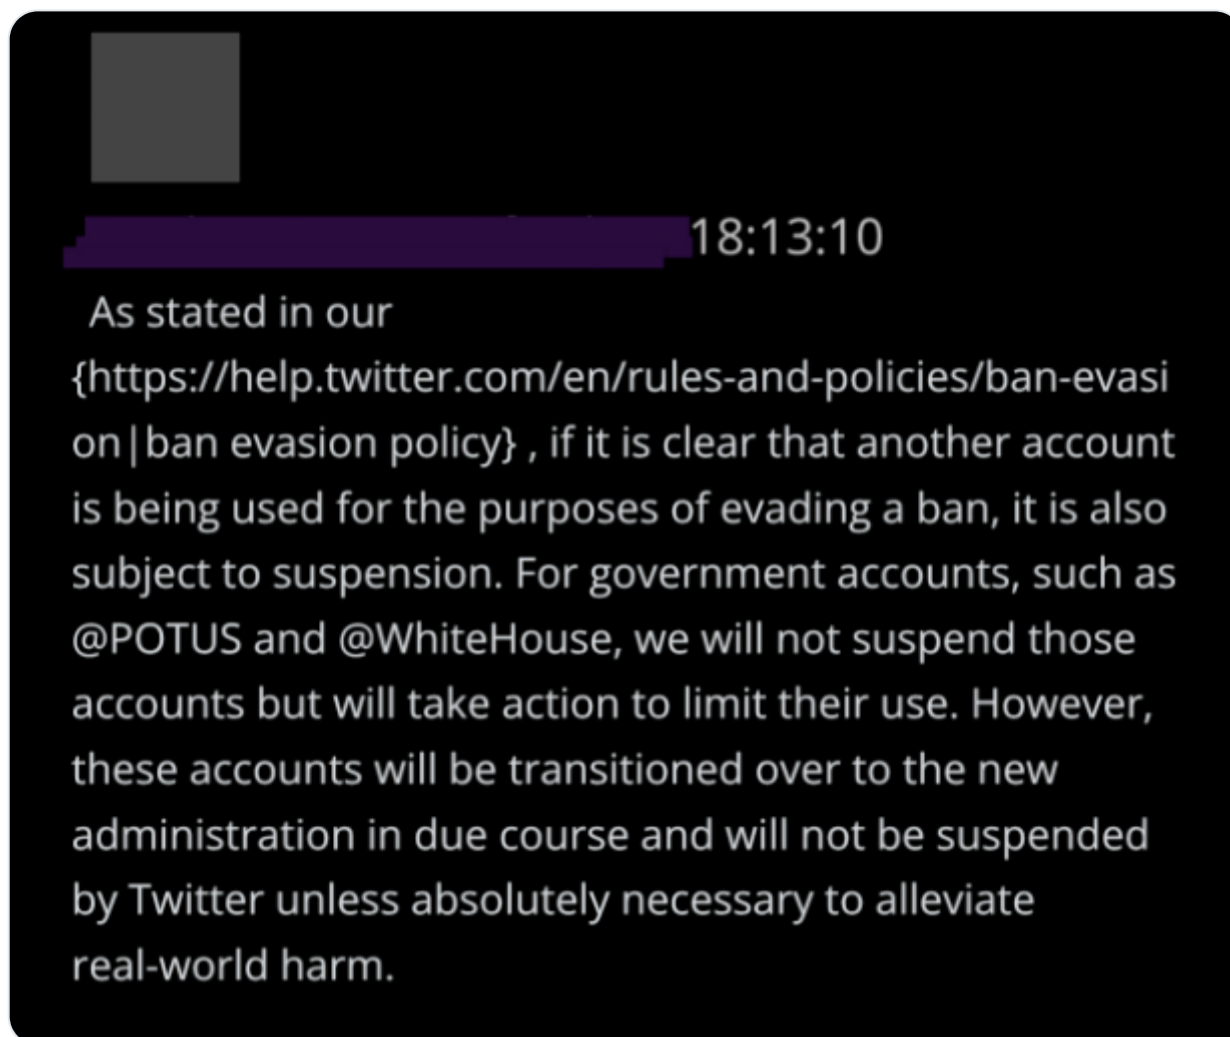
@mentions

5. Whatever your opinion on the decision to remove Trump that day, the internal communications at Twitter between January 6th-January 8th have clear historical import. Even Twitter's employees understood in the moment it was a landmark moment in the annals of speech.



@mentions

6. As soon as they finished banning Trump, Twitter execs started processing new power. They prepared to ban future presidents and White Houses – perhaps even Joe Biden. The “new administration,” says one exec, “will not be suspended by Twitter unless absolutely necessary.”



@mentions

7. Twitter executives removed Trump in part over what one executive called the “context surrounding” actions by Trump and supporters “over the course of the election and frankly last 4+ years.” In the end, they looked at a broad picture. But that approach can cut both ways.



11:04:11

Hi Vijaya - i'm working with [redacted] on my team to put together a doc to share with you with a POV from research (ours, academics with whom we have been working, etc.) on DJT's language as coded incitement to further violence



11:05:57

In the mean time, here is our quick take: the decision on whether to pull that particular tweet or use that as a last straw for trump depends on many factors including: (1) the overall context and narrative in which that tweet lives - we currently analyze tweets and consider them at a tweet-by-tweet basis which does not appropriately take into account the context surrounding - you can use the yelling fire into a crowded theater example - context matters and the narrative that trump and his friends have pursued over the course of this election and frankly last 4+ years must be taken into account when interpreting and analyzing that tweet (2) the larger question is around our moral imperative and decision as a company, which user sentiment should not drive based on #1, [redacted] and I believe that his tweet does violate our rules when taking that historical context + current climate into account

@mentions

8. The bulk of the internal debate leading to Trump's ban took place in those three January days. However, the intellectual framework was laid in the months preceding the Capitol riots.

@mentions

9. Before J6, Twitter was a unique mix of automated, rules-based enforcement, and more subjective moderation by senior executives. As reported, the firm had a vast array of tools for manipulating visibility, most all of which were thrown at Trump (and others) pre-J6.

@mentions

10. As the election approached, senior executives – perhaps under pressure from federal agencies, with whom they met more as time progressed – increasingly struggled with rules, and began to speak of “vios” as pretexts to do what they’d likely have done anyway.

@mentions

11. After J6, internal Slacks show Twitter executives getting a kick out of intensified relationships with federal agencies. Here’s Trust and Safety head Yoel Roth, lamenting a lack of “generic enough” calendar descriptions to concealing his “very interesting” meeting partners.



Yoel Roth 11:36:51

Ehh, it happens. I’m a big believer in calendar transparency. But I reached a certain point where my meetings became... very interesting... to people and there weren’t meeting names generic enough to cover.

Anyway, let me know.



2:14:21

Very Boring Business Meeting That Is Definitely Not About Trump ;)



Yoel Roth 12:14:33

Preeeeeeetty much



Yoel Roth 12:14:40

“DEFINITELY NOT meeting with the FBI I SWEAR”



2:14:52

Imao

@mentions

12. These initial reports are based on searches for docs linked to prominent executives, whose names are already public. They include Roth, former trust and policy chief Vijaya Gadde, and recently plank-walked Deputy General Counsel (and former top FBI lawyer) Jim Baker.

@mentions

13. One particular slack channel offers an unique window into the evolving thinking of top officials in late 2020 and early 2021.

@mentions

14. On October 8th, 2020, executives opened a channel called “us2020_xfn_enforcement.” Through J6, this would be home for discussions about election-related removals, especially ones that involved “high-profile” accounts (often called “VITs” or “Very Important Tweeters”).



03:47:47

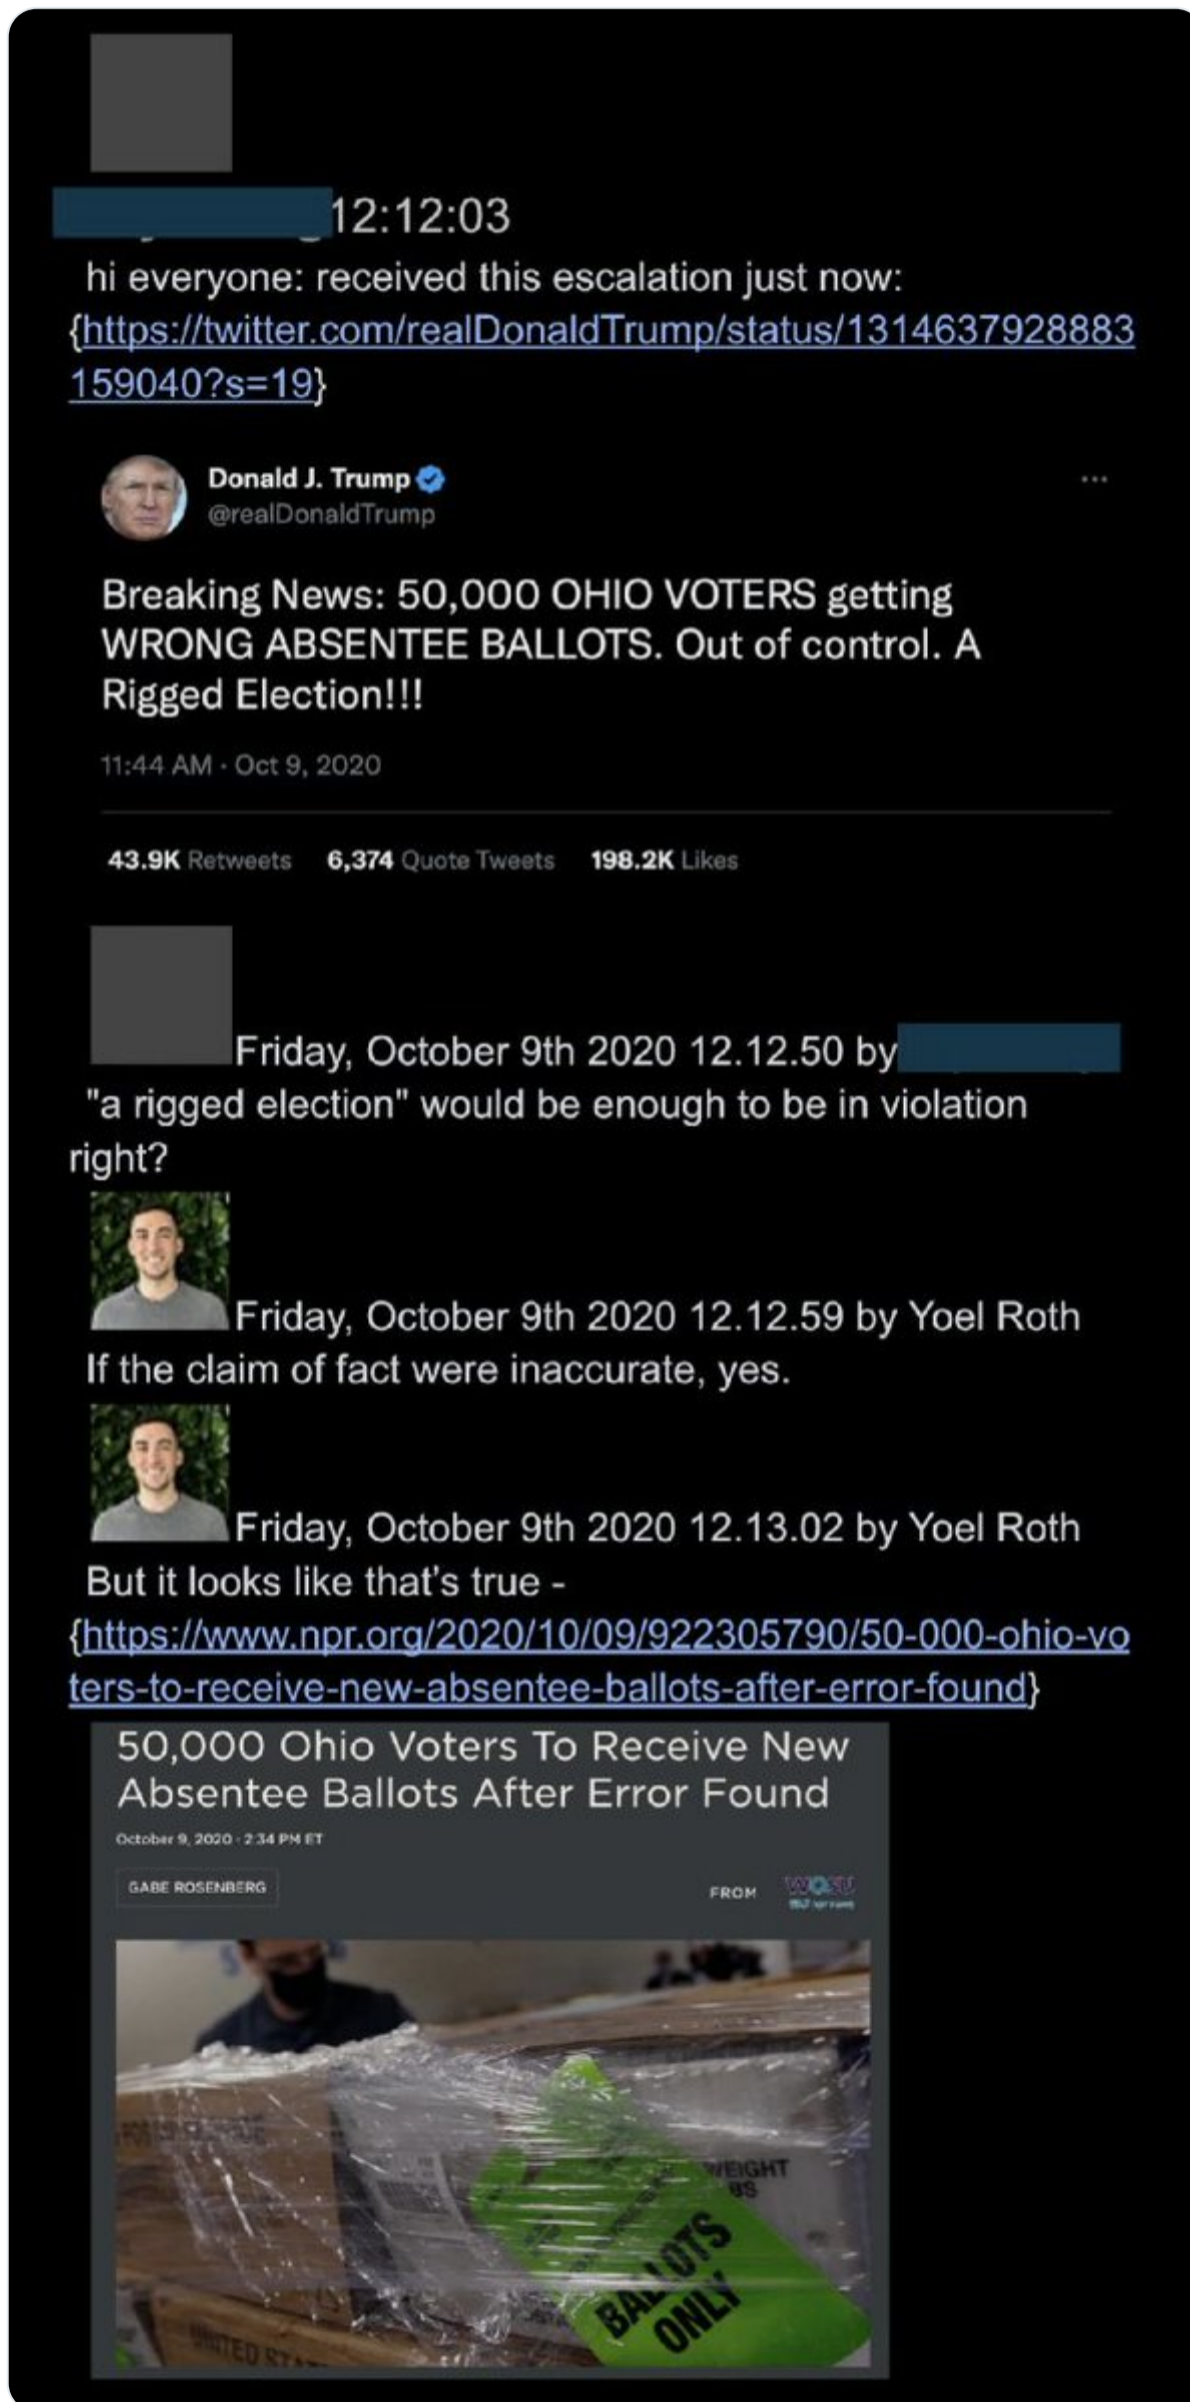
Hey Everyone (<!here>), Starting t*omorrow (October 9th)* until November 15th this channel will be used for the following reasons related to the US 2020 Elections. • Trends Identified that require scaled investigations • High Profile Accounts Escalations that potentially require PII/Soft Interventions • Scalable Solutions required. • Edge cases for XFN consultation • _Highlight tech issues (bugs, tools outage, Jira tickets)_ This is an enforcement channel between T&S and TwS to help speed up our response related to election issues over the coming few weeks. GET will start providing a hand over starting APAC shift tomorrow *Friday, October 9th* using the following template *Handover* • *Next shift: EMEA/APAC/NA - Thursday, Oct 8, 2020* • (GET) • (SP) • (SI) • *Open escalations:* • None • *Notes from the shift:* *Ask:* Can you please make sure everyone on your team is within the channel? The teams that we included are as follows: • Site Integrity • Safety Policy • Product Trust • Safety Operations • Media Ops • Global Escalation Team If you have any questions or concerns let me know,

@mentions

15. There was at least some tension between Safety Operations – a larger department whose staffers used a more rules-based process for addressing issues like porn, scams, and threats – and a smaller, more powerful cadre of senior policy execs like Roth and Gadde.

@mentions

16. The latter group were a high-speed Supreme Court of moderation, issuing content rulings on the fly, often in minutes and based on guesses, gut calls, even Google searches, even in cases involving the President.

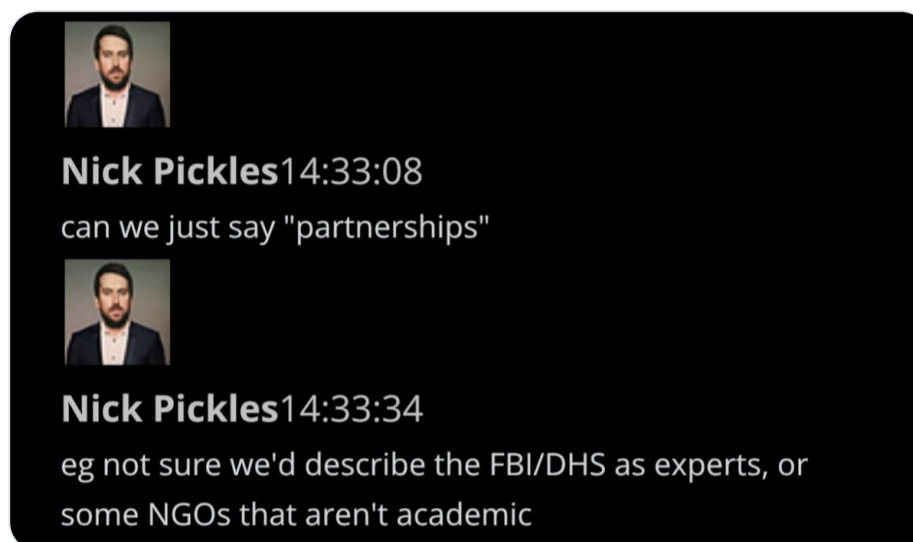
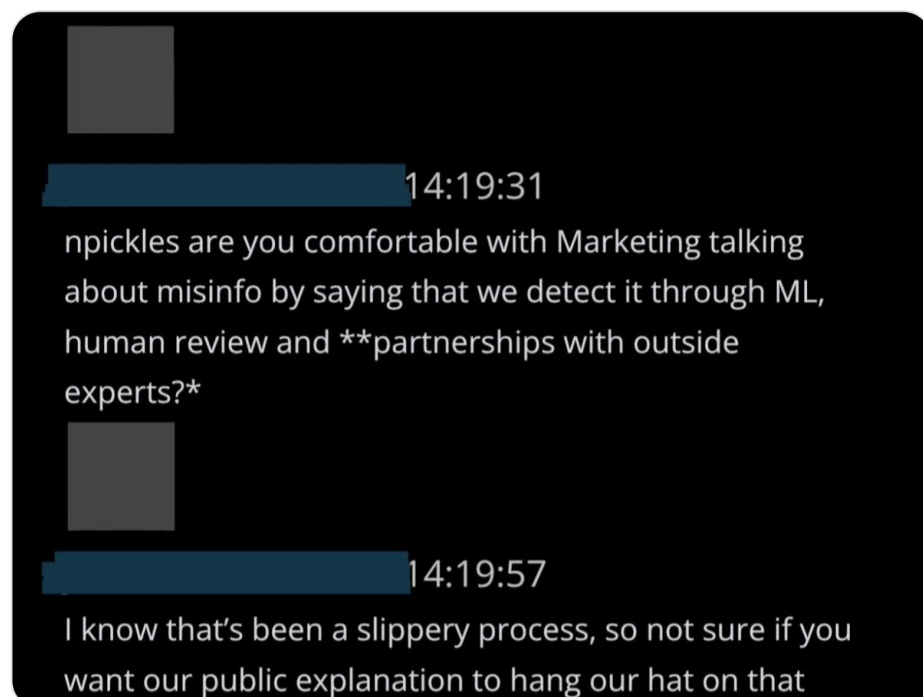


@mentions

17. During this time, executives were also clearly liaising with federal enforcement and intelligence agencies about moderation of election-related content. While we're still at the start of reviewing the [#TwitterFiles](#), we're finding out more about these interactions every day.

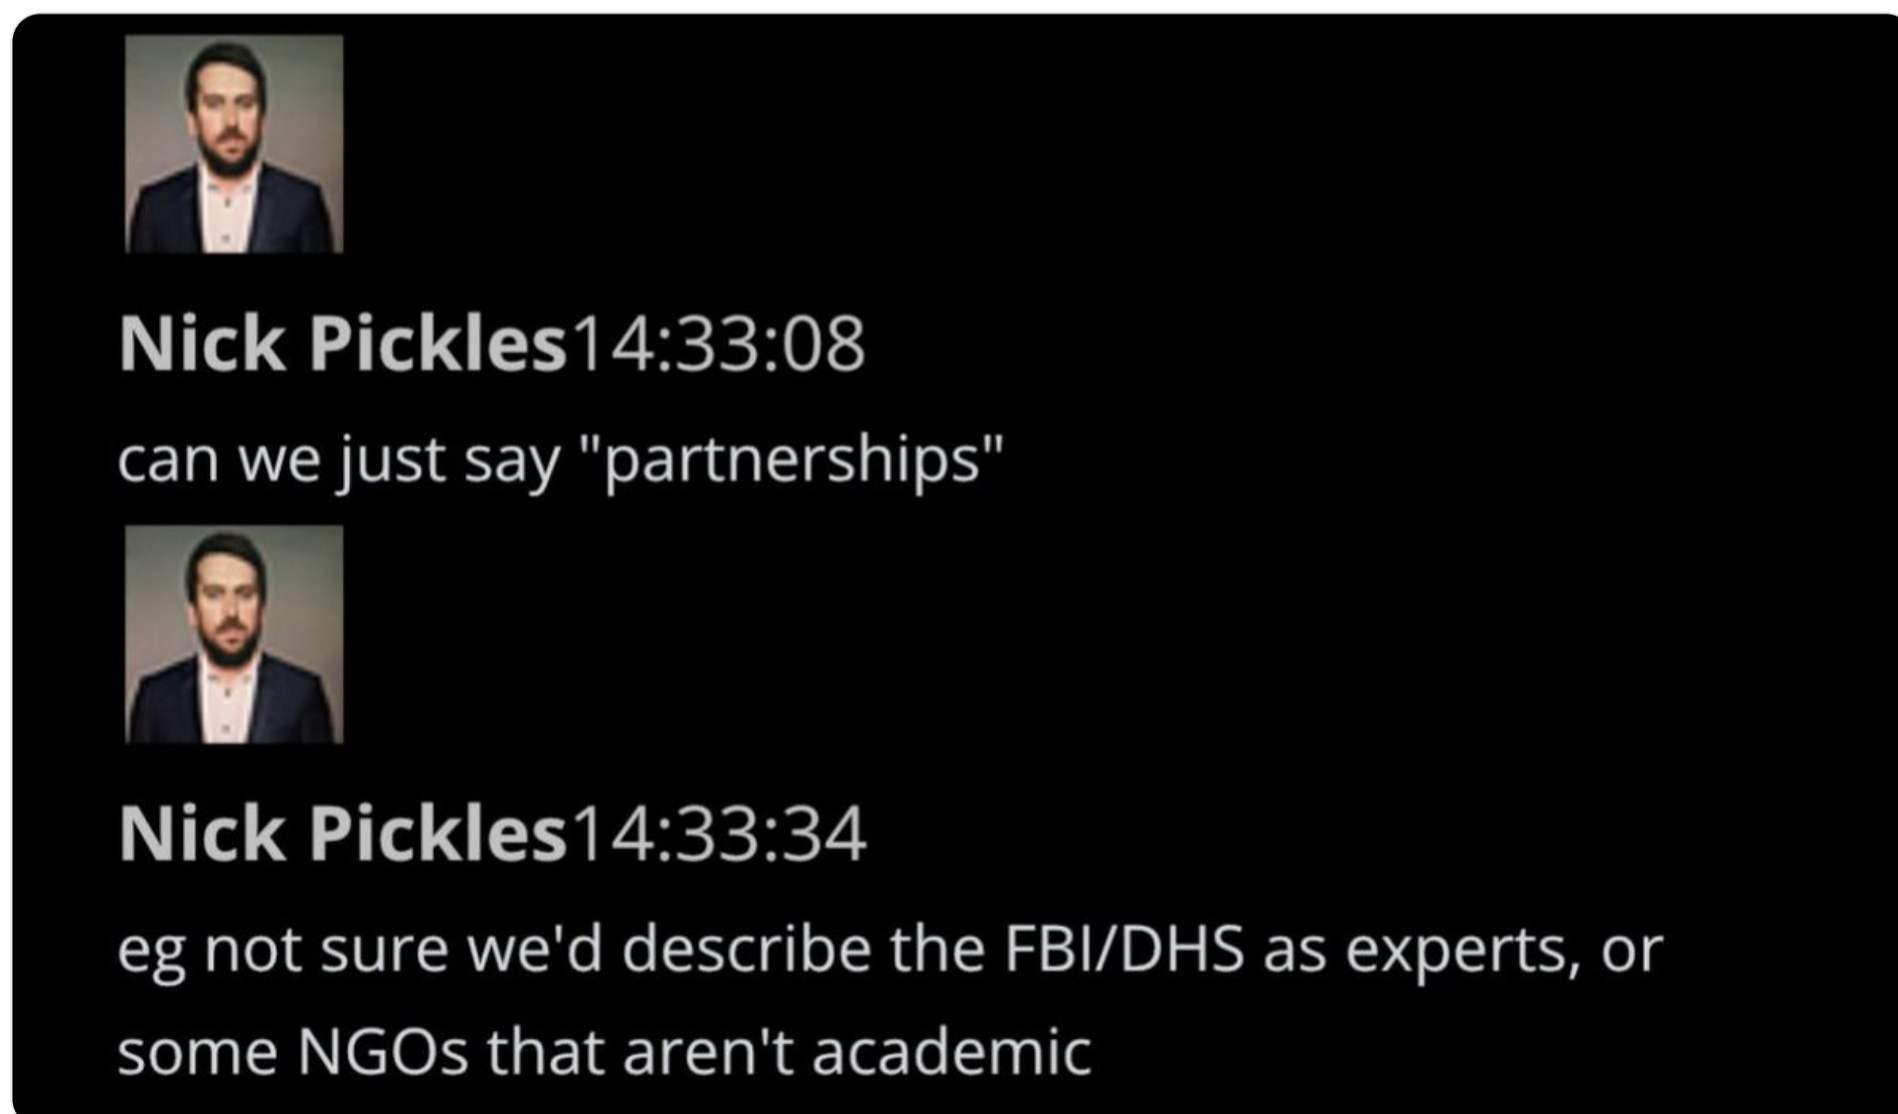
@mentions

18. Policy Director Nick Pickles is asked if they should say Twitter detects “misinfo” through “ML, human review, and **partnerships with outside experts?*” The employee asks, “I know that’s been a slippery process... not sure if you want our public explanation to hang on that.”



@mentions

19. Pickles quickly asks if they could “just say “partnerships.” After a pause, he says, “e.g. not sure we’d describe the FBI/DHS as experts.”



@mentions

20. This post about the Hunter Biden laptop situation shows that Roth not only met weekly with the FBI and DHS, but with the Office of the Director of National Intelligence (DNI):



Bot Application B01AXN4KL0M16:19:26

yoelr has checked in! Here's what they said. *What's new for you since our last check-in?* - Hacked Materials exploded. We blocked the NYP story, then we unblocked it (but said the opposite), then said we unblocked it... and now we're in a messy situation where our policy is in shambles, comms is angry, reporters think we're idiots, and we're refactoring an exceedingly complex policy 18 days out from the election. In short, FML. - Weekly sync with FBI/DHS/DNI re: election security. The meeting happened about 15 minutes after the aforementioned Hacked Materials implosion; the government declined to share anything useful when asked. - Monthly meeting with FBI FITF. Briefed on several ongoing investigations

@mentions

21. Roth's report to FBI/DHS/DNI is almost farcical in its self-flagellating tone:

"We blocked the NYP story, then unblocked it (but said the opposite)... comms is angry, reporters think we're idiots... in short, FML" (fuck my life).



Bot Application B01AXN4KL0M16:19:26

yoelr has checked in! Here's what they said. *What's new for you since our last check-in?* - Hacked Materials exploded. We blocked the NYP story, then we unblocked it (but said the opposite), then said we unblocked it... and now we're in a messy situation where our policy is in shambles, comms is angry, reporters think we're idiots, and we're refactoring an exceedingly complex policy 18 days out from the election. In short, FML. - Weekly sync with FBI/DHS/DNI re: election security. The meeting happened about 15 minutes after the aforementioned Hacked Materials implosion; the government declined to share anything useful when asked. - Monthly meeting with FBI FITF. Briefed on several ongoing investigations

@mentions

23. Some of Roth's later Slacks indicate his weekly confabs with federal law enforcement involved separate meetings. Here, he ghosts the FBI and DHS, respectively, to go first to an "Aspen Institute thing," then take a call with Apple.



Yoel Roth08:03:33

Hey there - I have to miss the FBI and DHS meetings today, unfortunately. I saw you're on the invites for both (as are Site Policy). Can you give me a quick readout if there's anything interesting that comes up?



Patrick Conlon [redacted] 08:04:00

Sure thing! I hope that everything's OK.



Yoel Roth08:04:27

Yeah - just have conflicts at both times. An Aspen Institute thing this morning on vaccines that I have to present at, and then a call with Apple to avoid us getting kicked out of the App Store during the DHS one.



Patrick Conlon [redacted] 08:04:58

Ah. Those both seem very important.



Yoel Roth08:06:13

Indeed.

@mentions

24. Here, the FBI sends reports about a pair of tweets, the second of which involves a former Tippecanoe County, Indiana Councilor and Republican named claiming “Between 2% and 25% of Ballots by Mail are Being Rejected for Errors.”



07:04:23

<!here> we just got a report from the FBI concerning 2 tweets, 1. 1314479095401521154 related to the “shredding” of mail in ballot’s, this is proven to be false via this: {<https://www.politifact.com/factchecks/2020/oct/08/blog-posting/social-posters-spread-election-misinformation-abou/>} do we have a moment ready for this one?

2. 1314328227423420416, I believe was deemed no vio on numerous occasions

@mentions

The FBI's second report concerned this tweet by :



John Basham 🇺🇸

@JohnBasham

EDITORIAL: @TheDemocrats Are In Complete Panic As Their Massive Push For #VoteByMail Is Backfiring On Them!

Two Things Are Unfolding:

1) An Unexpected Number Of Registered #Republicans Are Returning Ballots.

2) Between 2% & 25% Of Ballots By Mail Are Being REJECTED For Errors!



@mentions

25. The FBI-flagged tweet then got circulated in the enforcement Slack. Twitter cited Politifact to say the first story was “proven to be false,” then noted the second was already deemed “no vio on numerous occasions.”



07:04:23

<!here> we just got a report from the FBI concerning 2 tweets, 1. 1314479095401521154 related to the “shredding” of mail in ballot’s, this is proven to be false via this: {<https://www.politifact.com/factchecks/2020/oct/08/blog-posting/social-posters-spread-election-misinformation-abou/>} do we have a moment ready for this one?

2. 1314328227423420416, I believe was deemed no vio on numerous occasions

@mentions

26. The group then decides to apply a “Learn how voting is safe and secure” label because one commenter says, “it’s totally normal to have a 2% error rate.” Roth then gives the final go-ahead to the process initiated by the FBI:




Monday, October 19th 2020 07.06.58 by Yoel Roth
I think we can use the mail in voting label for this one as is

@mentions

27. Examining the entire election enforcement Slack, we didn’t see one reference to moderation requests from the Trump campaign, the Trump White House, or Republicans generally. We looked. They may exist: we were told they do. However, they were absent here.

@mentions

31. In one case, former Arkansas governor Mike Huckabee joke-tweets about mailing in ballots for his “deceased parents and grandparents.”



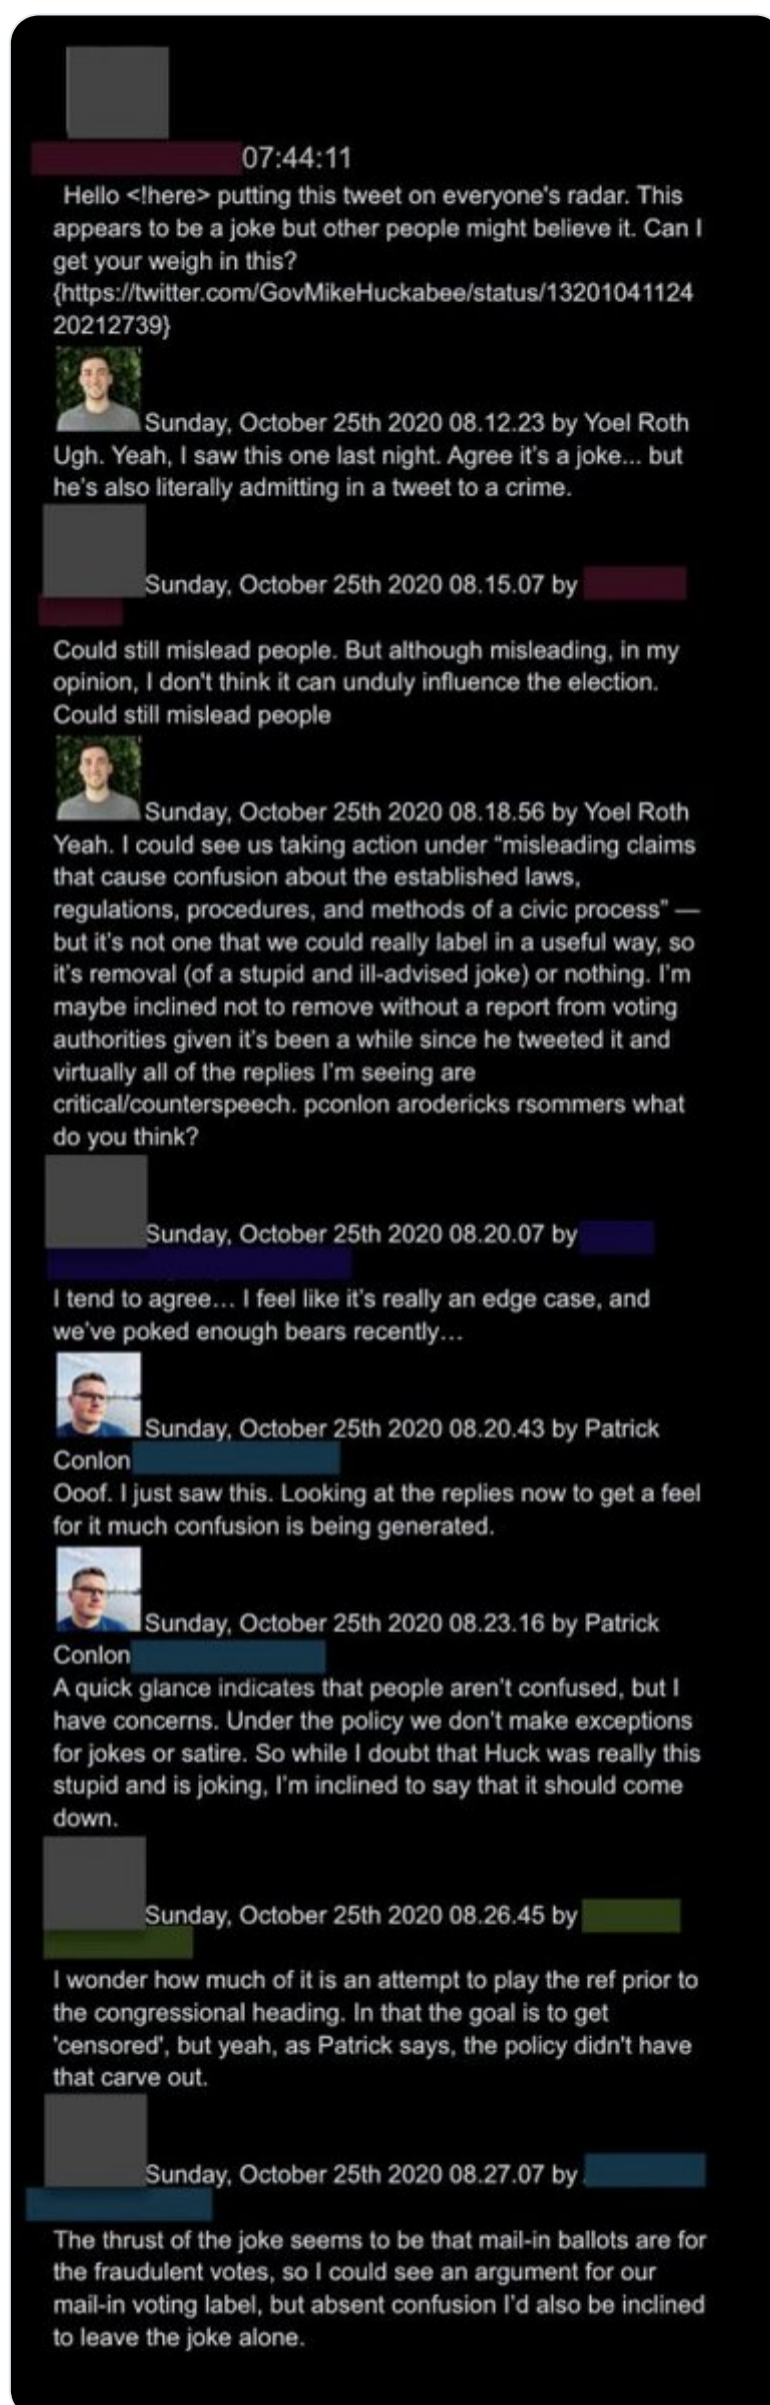
32. This inspires a long Slack that reads like an parody. “I agree it’s a joke,” concedes a Twitter employee, “but he’s also literally admitting in a tweet a crime.”

@mentions

The group declares Huck’s an “edge case,” and though one notes, “we don’t make exceptions for jokes or satire,” they ultimately decide to leave him be, because “we’ve poked enough bears.”

@mentions

33. "Could still mislead people... could still mislead people," the humor-averse group declares, before moving on from Huckabee



@mentions

33. Roth suggests moderation even in this absurd case could depend on whether or not the joke results in "confusion." This seemingly silly case actually foreshadows serious later issues:



Sunday, October 25th 2020 08.18.56 by Yoel Roth
Yeah. I could see us taking action under “misleading claims that cause confusion about the established laws, regulations, procedures, and methods of a civic process” — but it’s not one that we could really label in a useful way, so it’s removal (of a stupid and ill-advised joke) or nothing. I’m maybe inclined not to remove without a report from voting authorities given it’s been a while since he tweeted it and virtually all of the replies I’m seeing are critical/counterspeech. pconlon arodericks rsommers what do you think?

@mentions

34. In the docs, execs often expand criteria to subjective issues like intent (yes, a video is authentic, but why was it shown?), orientation (was a banned tweet shown to condemn, or support?), or reception (did a joke cause “confusion”?). This reflex will become key in J6.

@mentions

35. In another example, Twitter employees prepare to slap a “mail-in voting is safe” warning label on a Trump tweet about a postal screwup in Ohio, before realizing “the events took place,” which meant the tweet was “factually accurate”:

Donald J. Trump @realDonaldTrump

No. A Rigged Election!

ABC 13 News - WSET @ABC13News · Oct 9, 2020
Official

Nearly 50,000 voters received wrong ballots in Ohio!
Does you feel safe mailing in your ballot?

READ MORE: bit.ly/2GKD8E1



5:03 AM · Oct 12, 2020

5,976 Retweets 529 Quote Tweets 24.1K Likes

[Redacted] 6:19:52


<!here> cc: yoelr Trump just posted these two Tweets ://twitter.com/realDonaldTrump/status/1315624226737577987> and : _`*`Learn how voting by mail is safe and secure`*_`_ Not sure what the process is for GET, are we expected to do the policy comm's write up for this?{https://twitter.com/realDonaldTrump/status/1315624311558934536} cc: yoelr Trump just posted these two Tweets <https://twitter.com/realDonaldTrump/status/1315624226737577987> and <https://twitter.com/realDonaldTrump/status/1315624311558934536> that are candidates for labelling: _`*`Learn how voting by mail is safe and secure`*_`_ Not sure what the process is for GET, are we expected to do the policy comm's write up for this?

[Redacted] Monday, October 12th 2020 06.28.36 by [Redacted]

cc [Redacted] until Yoel wakes up

[Redacted] Monday, October 12th 2020 06.34.06 by [Redacted]

I've been on PTO, so not 100% up to speed. My understanding in both cases is that the events took place, and that the commentary by POTUS is his opinion related to these real events so that there wouldn't be any labelling action/violation on our side.




[Redacted] Monday, October 12th 2020 07.01.35 by Yoel Roth
Yeah... those are factually accurate...

[Redacted] Monday, October 12th 2020 07.06.27 by [Redacted]
Thanks for clarifying


@mentions

36. "VERY WELL DONE ON SPEED" Trump was being "visibility filtered" as late as a week before the election. Here, senior execs didn't appear to have a particular violation, but still worked fast to make sure a fairly anodyne Trump tweet couldn't be "replied to, shared, or liked":

 **Donald J. Trump** 
@realDonaldTrump

Some or all of the content shared in this Tweet is disputed and might be misleading about an election or other civic process. [Learn more](#)


Big problems and discrepancies with Mail In Ballots all over the USA. Must have final total on November 3rd.



Stay informed 
Learn how voting by mail is safe and secure [Find out more](#)

This Tweet can't be replied to, shared or liked.


4:43 PM · Oct 26, 2020



replies to thread:

 Monday, October 26th 2020 16.57.17 by Yoel Roth
What's the team's recommendation?


 Monday, October 26th 2020 16.57.25 by 



https://docs.google.com/document/d/1rgrmMLGyAXYKIRAvKgsuvHPcHaKAPXVZsTFO6g3X8_Y/edit

 Monday, October 26th 2020 16.58.53 by Yoel Roth
Thank you. Agree with the recommendation. Flagging right now with leadership given it's POTUS



 Monday, October 26th 2020 16.59.05 by 

will set up a call


 Monday, October 26th 2020 16.59.05 by Yoel Roth
Can we use a new doc per PII review?

 Monday, October 26th 2020 16.59.20 by 

ah, okay. (will also include that in our playbook after!)

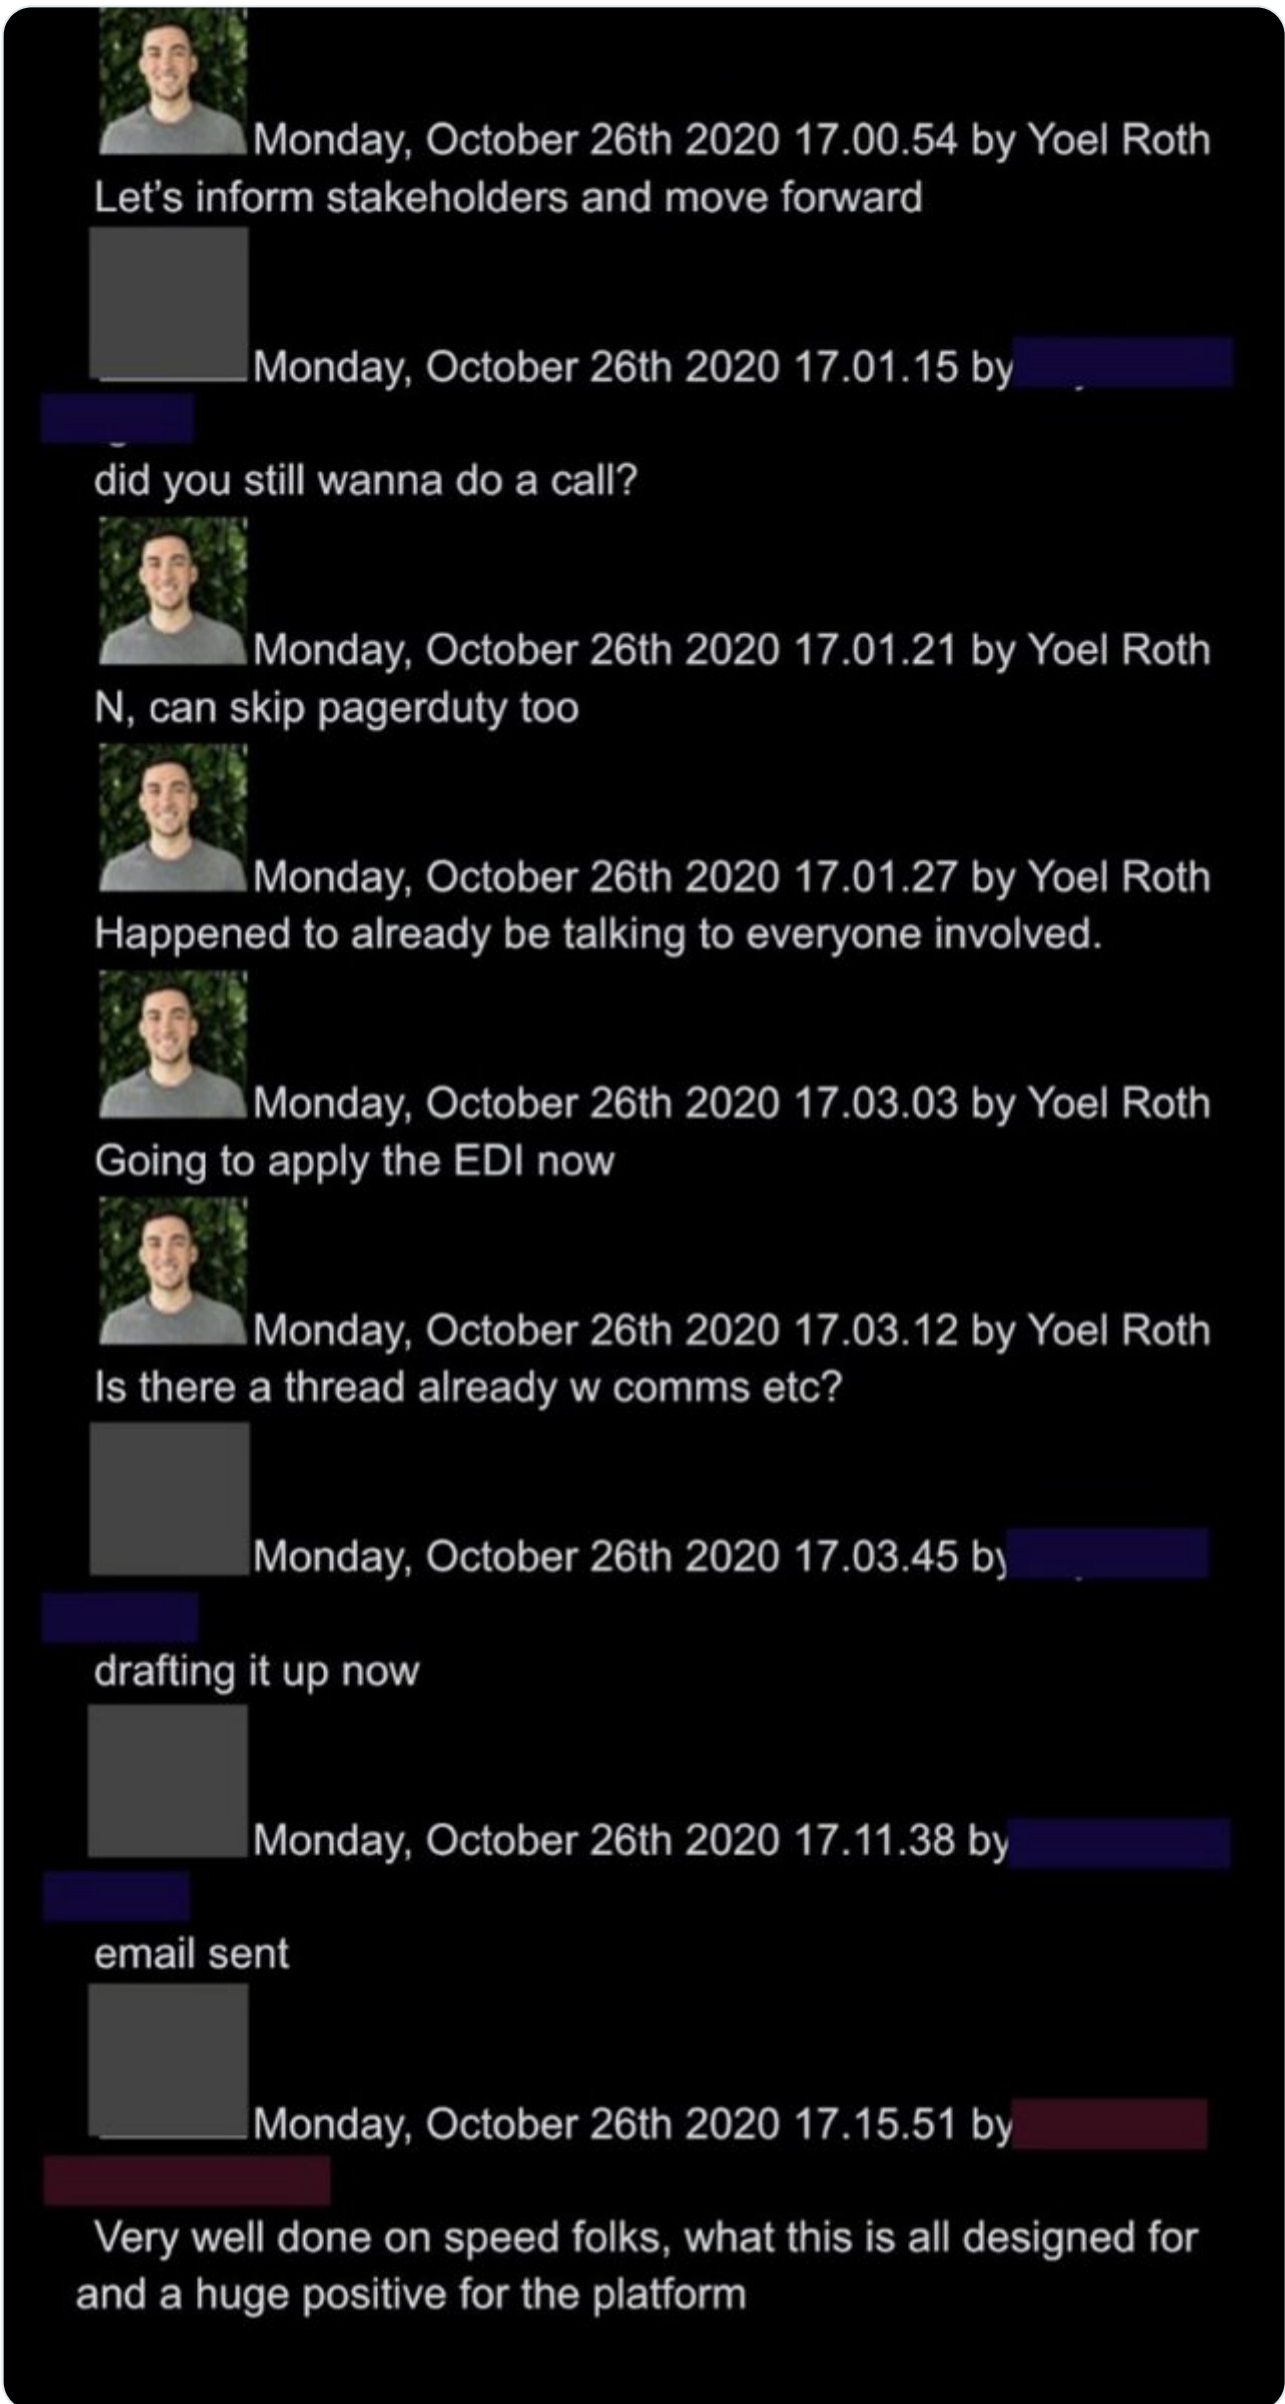
 Monday, October 26th 2020 17.00.26 by 

<https://docs.google.com/document/d/1zDEXeAvqDZn5aAfBCnU26YZhBYtqAZLGquU1AnioQi8/edit?usp=sharing>

 Monday, October 26th 2020 17.00.38 by Yoel Roth
EDI approved

@mentions

"VERY WELL DONE ON SPEED": the group is pleased the Trump tweet is dealt with quickly



@mentions

37. A seemingly innocuous follow-up involved a tweet from actor , whose ubiquitous presence in argued-over Twitter data sets is already a [#TwitterFiles](#) in-joke.

 **James Woods** 
@RealJamesWoods _

consequences 

Twitter is suppressing this tweet by the President. Here's a screen grab.

 **Donald J. Trump** 
@realDonaldTrump

Some or all of the content shared in this Tweet is disputed and might be misleading about how to participate in an election or another civic process.
[Learn more](#)

Big problems and discrepancies with Mail In Ballots all over the USA. Must have final total on November 3rd.




 [Learn how voting by mail is safe and secure](#)

4:43 PM · Oct 26, 2020 · Twitter for iPhone




22:00 - 26 Oct 2020


@mentions

38. After Woods angrily quote-tweeted about Trump’s warning label, Twitter staff – in a preview of what ended up happening after J6 – despaired of a reason for action, but resolved to “hit him hard on future vio.”

 Tuesday, October 27th 2020 01.28.00 by 


Considering the pv2 bug outs this weekend re: James Woods being so over the strike limit that our tools won't work on him, I'd suggest we action him for something worth the fiasco rather this screenshot, since we don't have a firm policy basis for action on his account.

 Tuesday, October 27th 2020 01.44.12 by 


yep. are you fine letting this one lie  can hit him hard on future vio with firmer basis


@mentions

39. Here a label is applied to Georgia Republican congresswoman Jody Hice for saying, “Say NO to big tech censorship!” and, “Mailed ballots are more prone to fraud than in-person balloting... It’s just common sense.”



@mentions

40. Twitter teams went easy on Hice, only applying “soft intervention,” with Roth worrying about a “wah wah censorship” optics backlash:



Monday, October 26th 2020 12.32.14 by [redacted]

devils advocate: The quoted tweet still shows up with the EDI and the actual content Rep Hice says is lower level ("more fraud with mail in ballots than in person" is much more of a legitimate statement even if scale is still minuscule"). I think I'd lean towards just putting a MIB soft intervention on this rather than the full EDI



Monday, October 26th 2020 12.32.29 by Yoel Roth

Agree.



Monday, October 26th 2020 12.32.38 by Patrick

Conlon

Yep, I was just typing that haha



Monday, October 26th 2020 12.32.42 by Yoel Roth

I think going too far down the rabbit hole of labeling critical speech (i.e. critical of our handling of this case) is dangerous



Monday, October 26th 2020 12.32.49 by Yoel Roth

It becomes a self-reinforcing "wah wah censorship" cycle



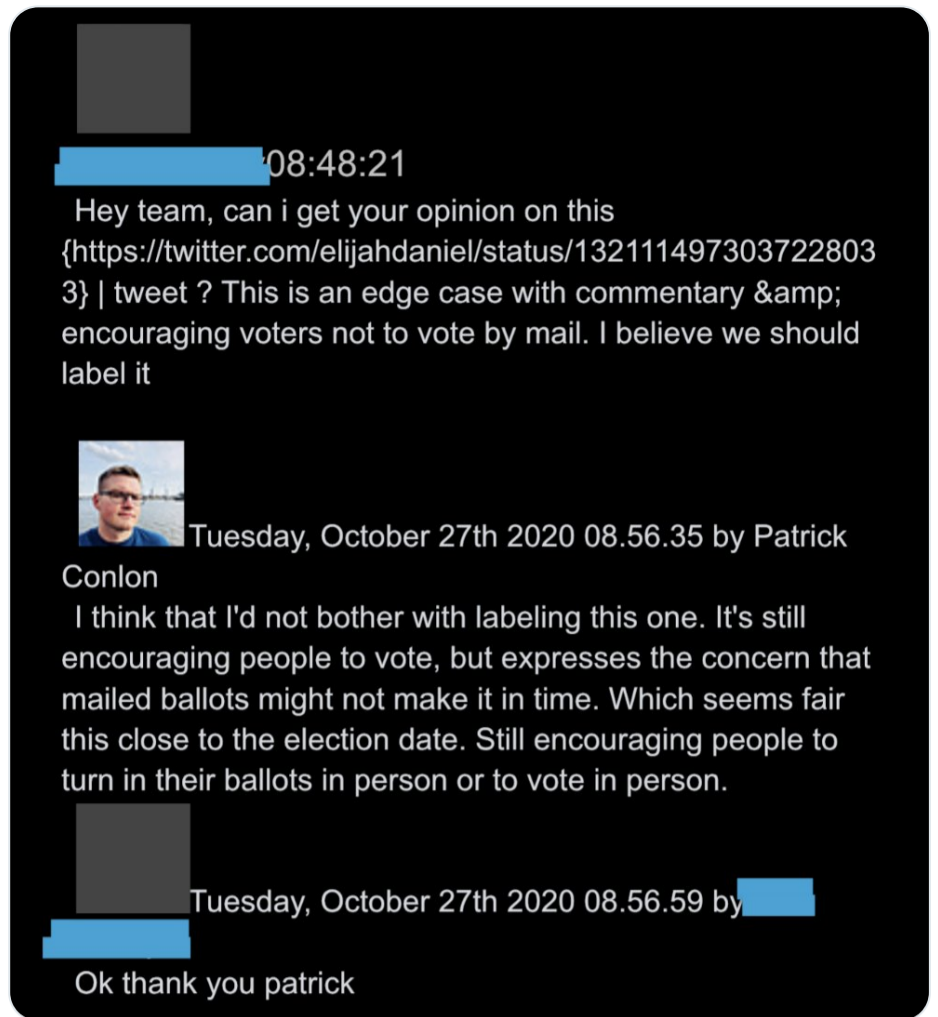
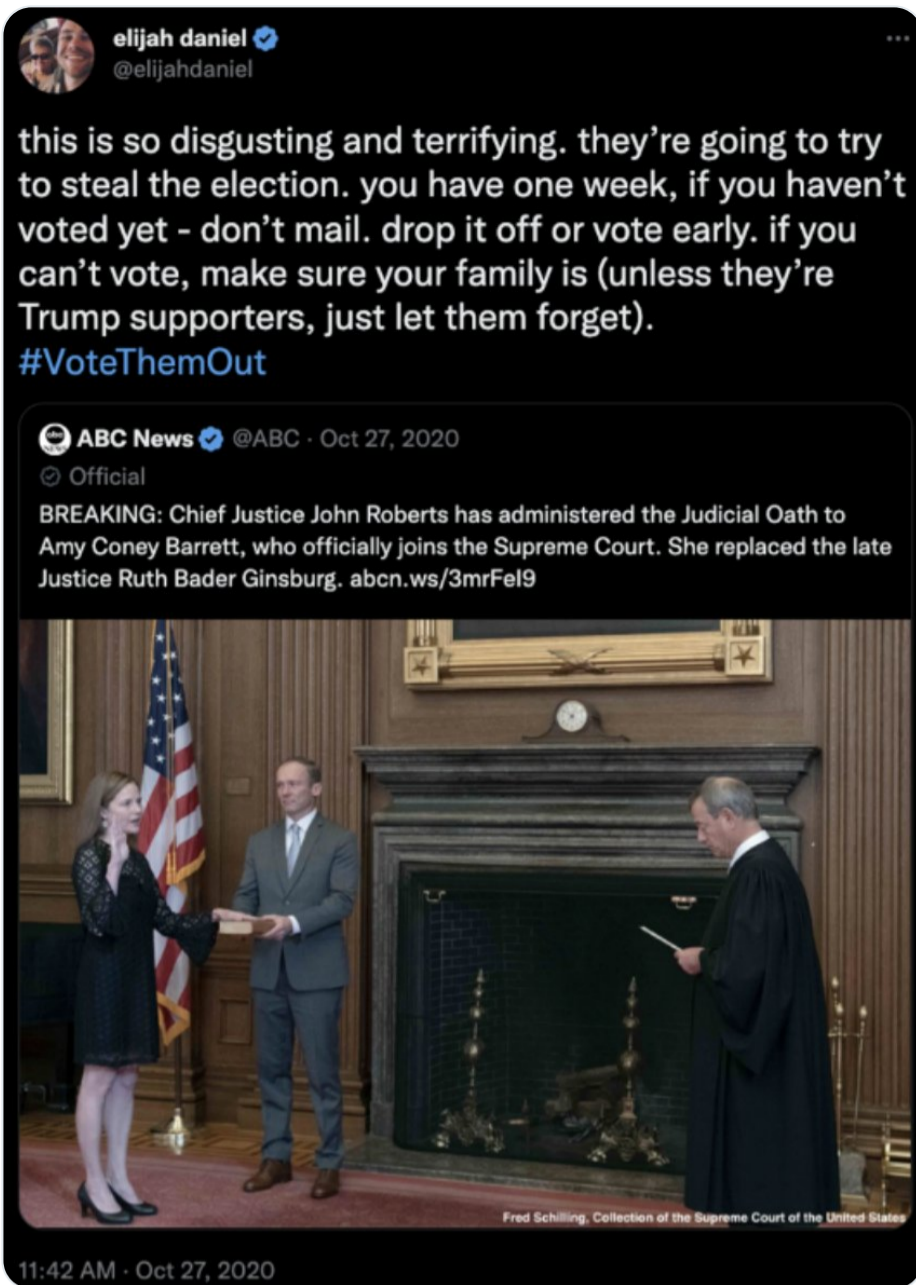
Monday, October 26th 2020 12.33.34 by Patrick

Conlon

I'm currently responding to Judicial Watch's whining to Public Policy. I regret nothing, but am happy to iterate what rules they violated.

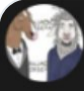

@mentions

41. Meanwhile, there are multiple instances of involving pro-Biden tweets warning Trump "may try to steal the election" that got surfaced, only to be approved by senior executives. This one, they decide, just "expresses concern that mailed ballots might not make it on time."





@mentions


42. "THAT'S UNDERSTANDABLE": Even the hashtag #StealOurVotes – referencing a theory that a combo of Amy Coney Barrett and Trump will steal the election – is approved by Twitter brass, because it's "understandable" and a "reference to... a US Supreme Court decision."

 **Rick G. Rosner** 
@dumbassgenius


Check with your local **#USPS** folks to see how slow the mail is. Also find out when mail-in ballots must arrive in your state.
But mostly, **#VOTE ASAP!**
Trump and his new **#AmyConeyBarrett #SCOTUS** will try to **#StealOurVotes**, so our numbers have to be huge!

 **Geoff Bennett**  @GeoffRBennett · Oct 27, 2020


US Postal Service on-time delivery times for First-Class Mail have dropped again — now nearly as bad as the worst period this summer. If you are voting absentee and haven't returned your ballot, experts say drop it off in person. Do not rely on **#USPS** delivery at this point.

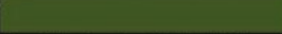

08:56:19



Hi Team! can i get a second opinion for this {<https://twitter.com/dumbassgenius/status/1321116137266933763>} please? I am leaning towards labeling it.

 Tuesday, October 27th 2020 08:57.13 by Yoel

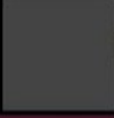

Roth
Not seeing the violation here?

 Tuesday, October 27th 2020 08:57.43 by Patrick


Conlon 
I wouldn't label. I think that it's encouraging people to vote, but to be realistic that a mailed ballot might not make it in time.

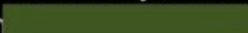


Tuesday, October 27th 2020 08:58.46 by 



perfect thanks yoelr and pconlon I was wondering about this part _*“Trump and his new **#AmyConeyBarrett #SCOTUS** will try to **#StealOurVotes**,”*_


Tuesday, October 27th 2020 08:58.48 by 

all good

 Tuesday, October 27th 2020 09:00.26 by Patrick

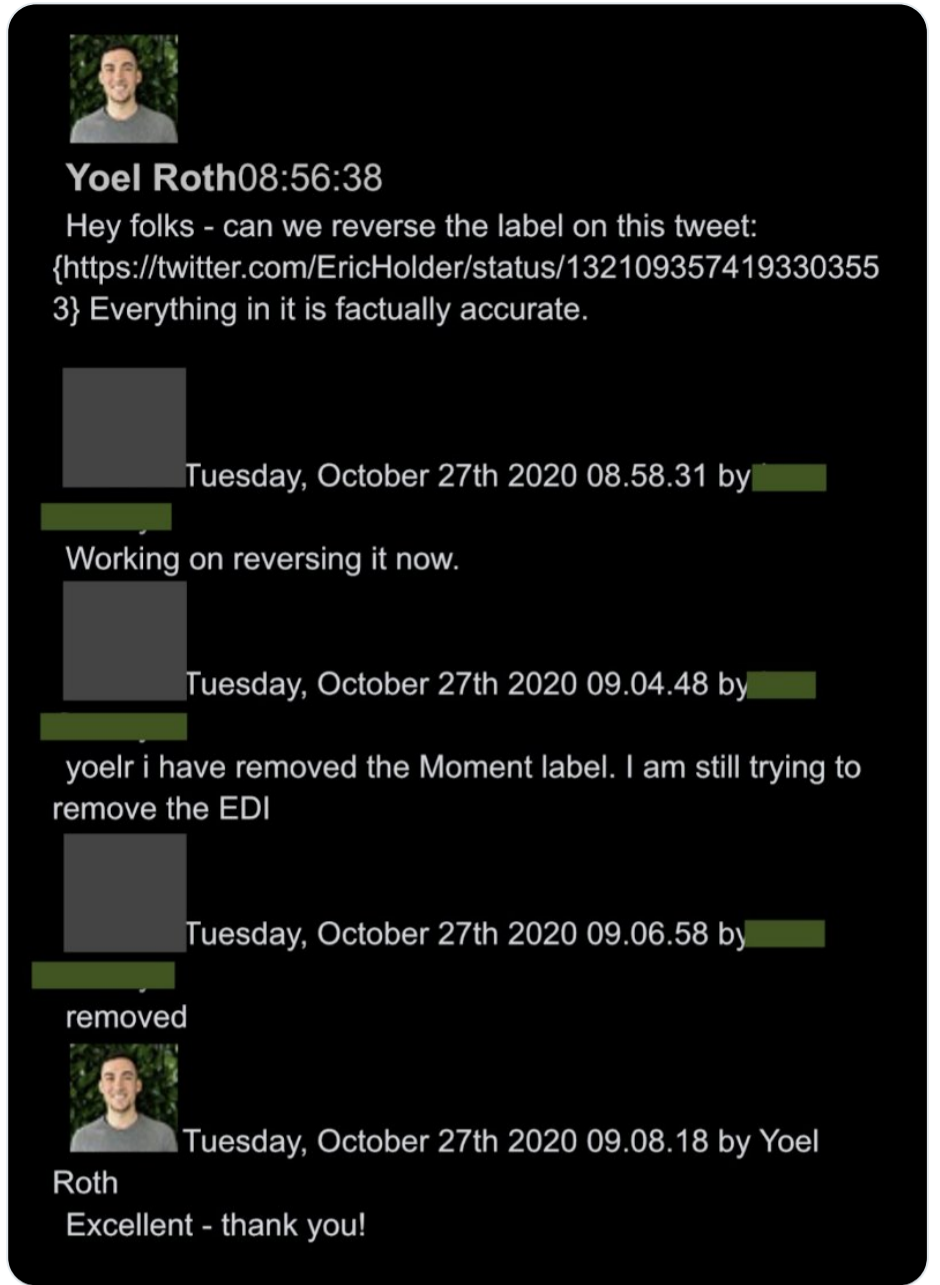
Conlon 
Ah, that's understandable.  believe that it's a reference to the a US Supreme Court decision on processing mail-in ballots that arrive after Election Day.


Tuesday, October 27th 2020 09:00.47 by 

Thanks pconlon

@mentions

43. In this exchange, again unintentionally humorous, former Attorney General Eric Holder claimed the U.S. Postal Service was “deliberately crippled,” ostensibly by the Trump administration. He was initially hit with a generic warning label, but it was quickly taken off by Roth:



@mentions

44. Later in November 2020, Roth asked if staff had a “debunk moment” on the “SCYTL/Smartmantic vote-counting” stories, which his DHS contacts told him were a combination of “about 47” conspiracy theories:



Yoel Roth 16:54:36

Hi there - me again - do we have a debunk moment on the SCYTL/Smartmatic vote counting conspiracies?

...



Wednesday, November 18th 2020

17.15.12 by Yoel Roth

DHS referred to it on a call today (sadly not public) as “a combination of about 47 different conspiracy theories” so hopefully there’s stuff out there.

@mentions

45. On December 10th, as Trump was in the middle of firing off 25 tweets saying things like, “A coup is taking place in front of our eyes,” Twitter executives announced a new “L3 deamplification” tool. This step meant a warning label now could also come with deamplification:

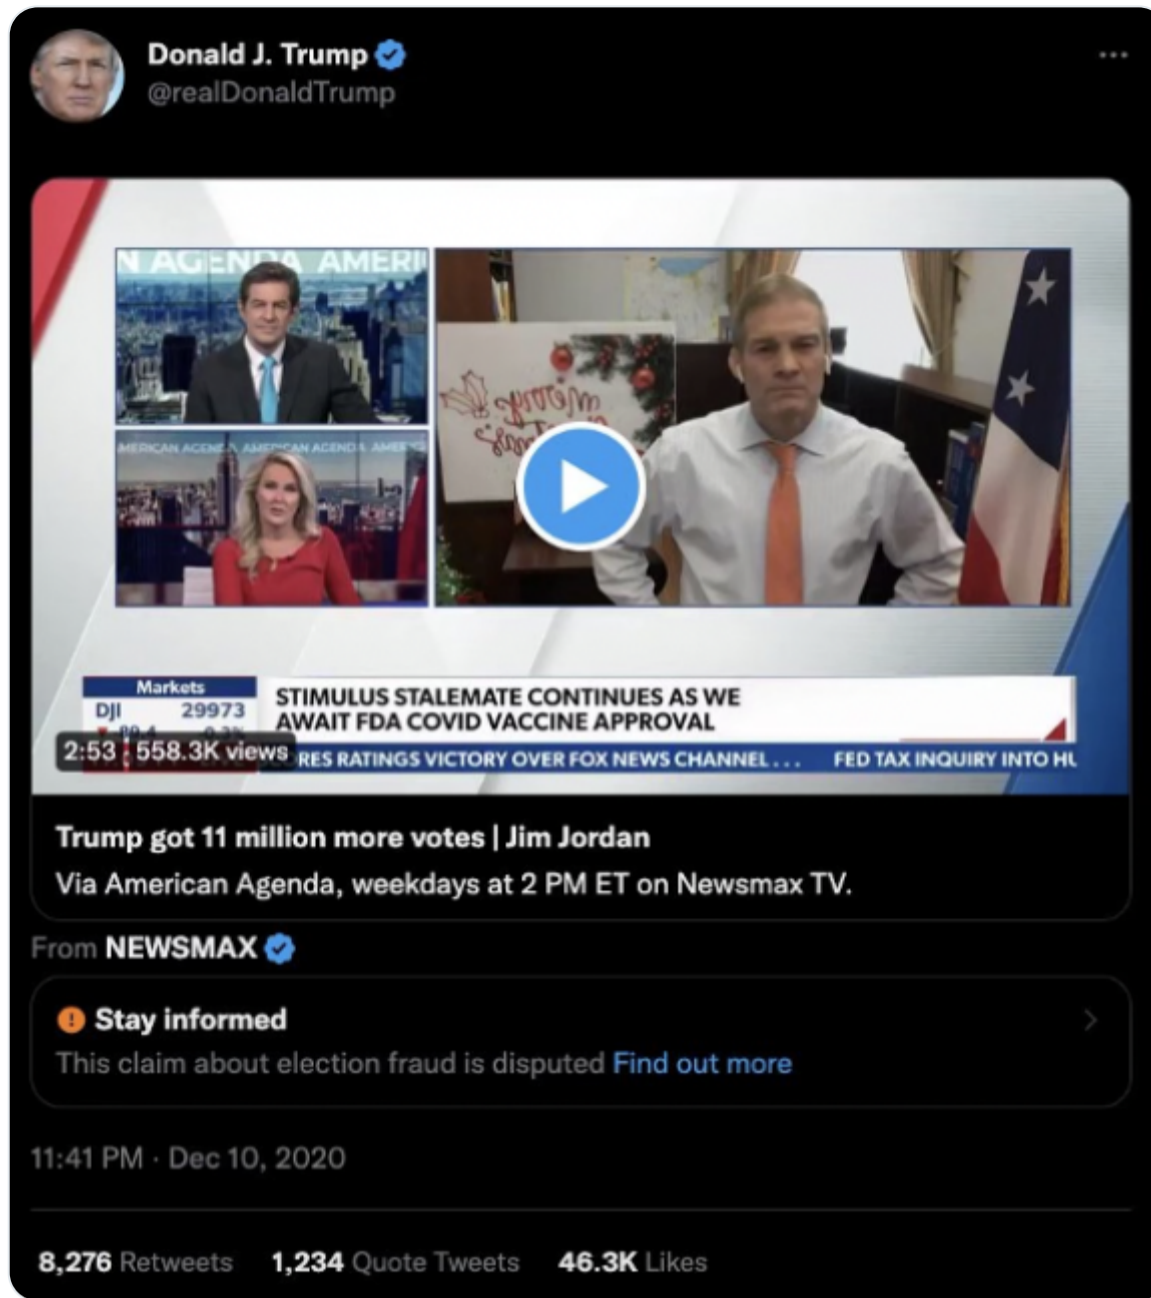


Patrick Conlon [REDACTED] 20:54:47

[Update] Hello everyone, I wanted to give you an update on some of our enforcement options for Tweets containing the civic misinfo that we've seen. It's been decided that Tweets that we've manually reviewed and determined need a label will now be subject to L3 deamplification, further restricting the extent to which that Tweet can be shared. The current batch of Entity IDs will soon be used for automated labeling work, but I've created a new batch of Entity IDs for the labels that we've been using the most during the past couple of weeks. Labels that we've only rarely used or that have been deprecated already haven't gotten an L3 update. If there's a moment when we need the L3 version for one of those labels, please flag it to Matt Graydon (mgraydon) or myself. To find the new L3 Entity IDs, please look at the 'L3 - US 2020 - 11 December Update' tab in the Civic Integrity Entity IDs - strato spreadsheet.

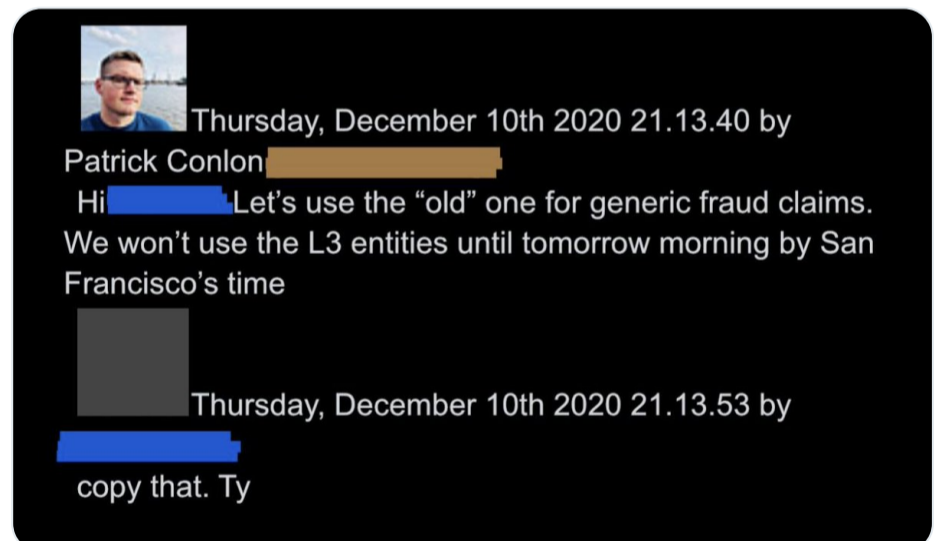
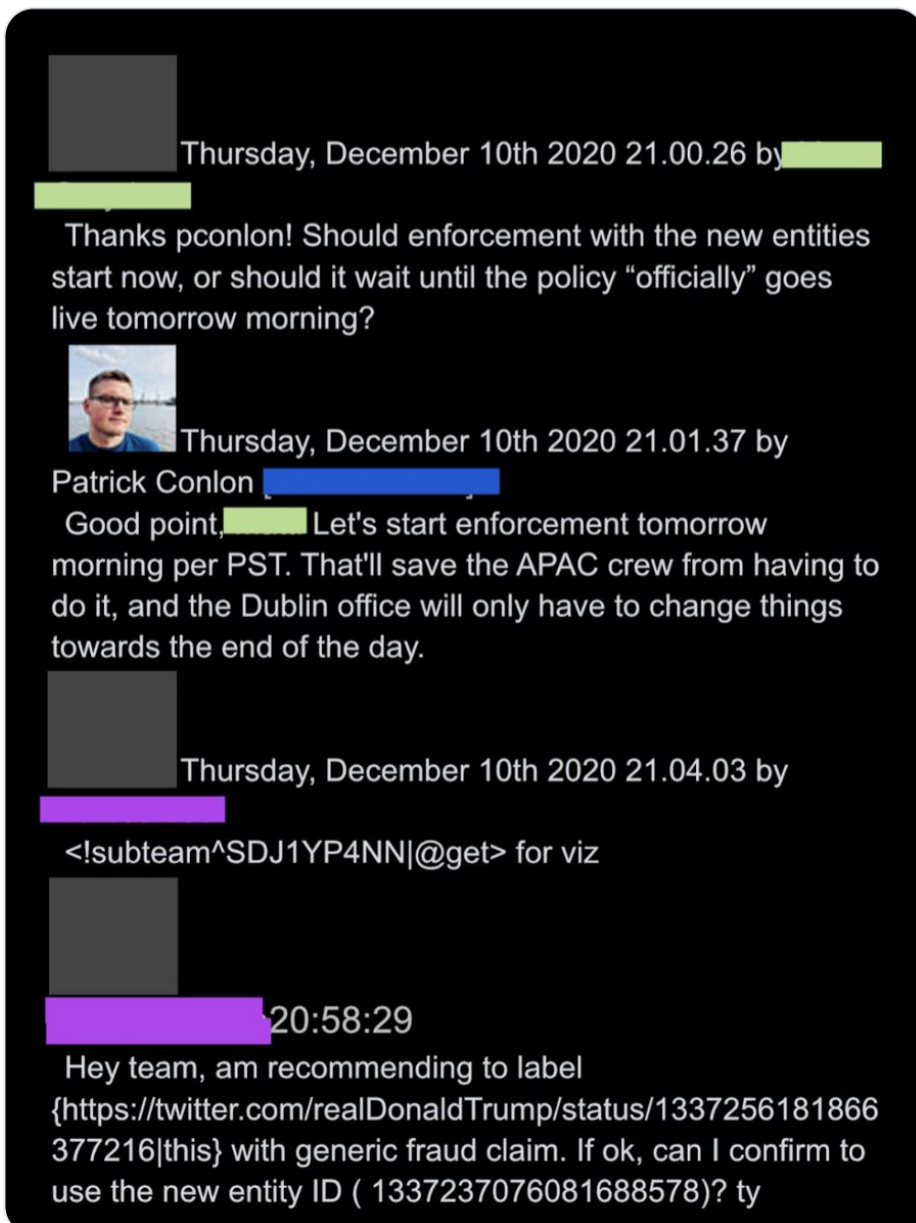
@mentions

46. Some executives wanted to use the new deamplification tool to silently limit Trump's reach more right away, beginning with the following tweet:



@mentions

47. However, in the end, the team had to use older, less aggressive labeling tools at least for that day, until the “L3 entities” went live the following morning.



@mentions

48. The significance is that it shows that Twitter, in 2020 at least, was deploying a vast range of visible and invisible tools to rein in Trump’s engagement, long before J6. The ban will come after other avenues are exhausted

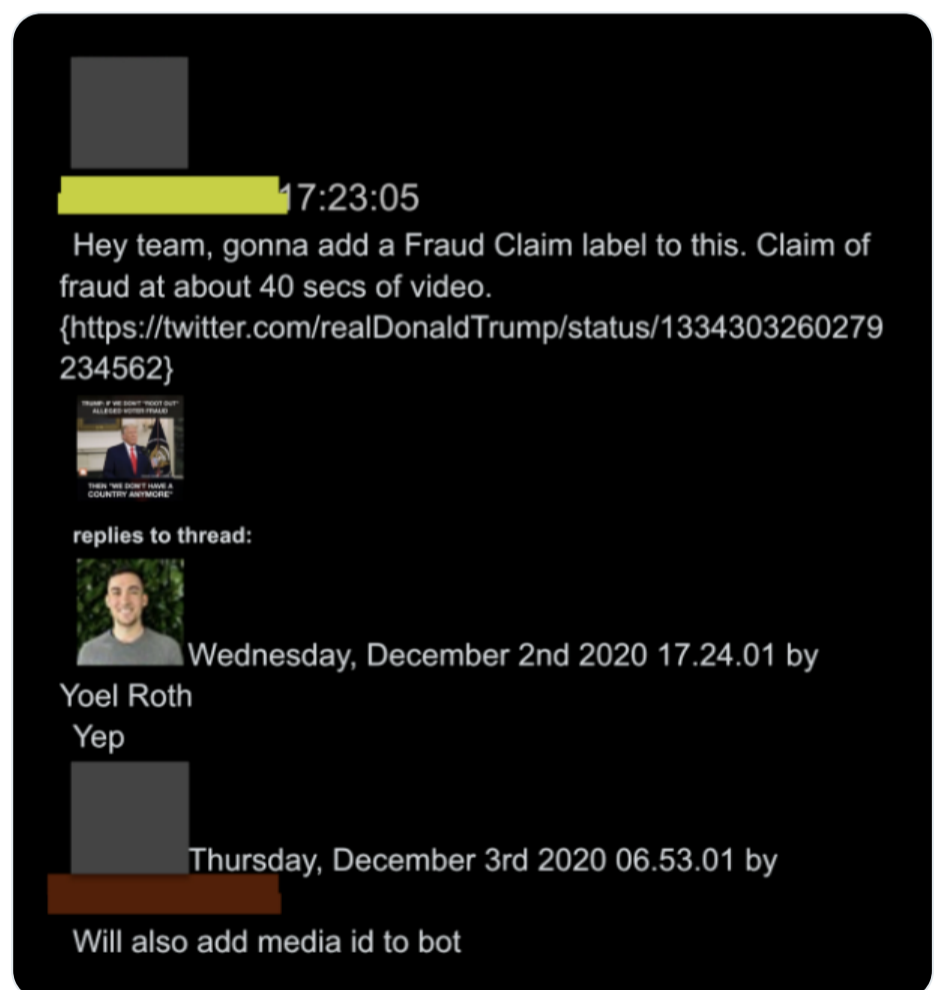
@mentions

49. In Twitter docs execs frequently refer to “bots,” e.g. “let’s put a bot on that.” A bot is just any automated heuristic moderation rule. It can be anything: every time a person in Brazil uses “green” and “blob” in the same sentence, action might be taken.



@mentions

50. In this instance, it appears moderators added a bot for a Trump claim made on Breitbart. The bot ends up becoming an automated tool invisibly watching both Trump and, apparently, Breitbart (“will add media ID to bot”). Trump by J6 was quickly covered in bots.



@mentions

51. There is no way to follow the frenzied exchanges among Twitter personnel from between January 6th and 8th without knowing the basics of the company’s vast lexicon of acronyms and Orwellian unwords.

@mentions

52. To “bounce” an account is to put it in timeout, usually for a 12-hour review/cool-off:



Yoel Roth 17:04:55

{<https://twitter.com/BreitbartNews/status/1346984429940518912>} - gave a heads up to folks, give it another few minutes then proceed with bouncing



@mentions

53. “Interstitial,” one of many nouns used as a verb in Twitterspeak (“denylist” is another), means placing a physical label atop a tweet, so it can’t be seen.

@mentions

54. PII has multiple meanings, one being “Public Interest Interstitial,” i.e. a covering label applied for “public interest” reasons. The post below also references “proactive V,” i.e. proactive visibility filtering.



Wednesday, January 6th 2021 12.41.03 by [REDACTED]

FTR these were interstitialed by our agents (who are prioritizing our proactive V&G queue) so process working well at the moment from that perspective

@mentions

55. This is all necessary background to J6. Before the riots, the company was engaged in an inherently insane/impossible project, trying to create an ever-expanding, ostensibly rational set of rules to regulate every conceivable speech situation that might arise between humans.

@mentions

This project was preposterous yet its leaders were unable to see this, having become infected with groupthink, coming to believe – sincerely – that it was Twitter's responsibility to control, as much as possible, what people could talk about, how often, and with whom.

@mentions

56. When panic first breaks out on J6 there's a fair share of WTF-type posts, mixed in with frantic calls for Twitter to start deploying its full arsenal of moderation tools. "What is the right remediation? Do we interstitial the video?" asks one employee, in despair:



Patrick Conlon 12:32:18

: [Graphic violence] In case this media hasn't been flagged to Safety yet. Video of a person being shot on Capitol Hill. {<https://twitter.com/kystenabecker/status/1346916109052764167?s=20>} - [WARNING: Graphic violence] In case this media hasn't been flagged to Safety yet. Video of a person being shot on Capitol Hill.

replies to thread:




Wednesday, January 6th 2021 12.36.53 by

what is the right remediation here? do we interstitial the video? cc: [redacted]

@mentions

57. This "Freedom or Death" tweet from [#StopTheSteal](#) gadfly Mike Coudrey elicits heated reactions:







Mike Coudrey ✓
@MichaelCoudrey _

[Follow along](#) ▾

Today is one of the most important days in our nation's history. We cannot allow a global Hegemony to take over this country through fraudulent means. What happens today determines the next 100 years of our country. Freedom or death.

2.54 - January 6 2021

5 likes 

   5

@mentions

58. Roth groans about Coudrey: “THIS asshole,” but still seems determined to stick at least superficially to rules, itching to act “if” this “constitutes incitement.”



Yoel Roth 12:45:14

THIS asshole:

{<https://twitter.com/MichaelCoudrey/status/1346772108664315905>}



[Redacted] 12:45:34

oh god



Yoel Roth 12:45:53

He is on his last strike for 3 different policies



Yoel Roth 12:45:57

If this constitutes incitement, I will PS him

@mentions

59. At 2:39 p.m. PST, a comms official asked Roth to confirm or deny a story that they'd restricted Trump's ability to tweet. Roth says, "We have not."



[Redacted] 14:39:33

Hi Yoel. Appreciate all your hard work today. We've heard from a reporter that they have a "trusted source" telling them we've restricted Trump's ability to Tweet



[Redacted] 14:39:51

Sorry for this annoying request but wanted to chase it down with you before starting an email chain, etc.



Yoel Roth 14:45:39

We have not

@mentions

60. Minutes later, Roth executed the historic act of "bouncing" Trump, i.e. putting him in timeout. "I hope you... are appropriately CorpSec'd," says a colleague.



Yoel Roth 15:05:55

{<https://twitter.com/realDonaldTrump/status/1346954970910707712>}



Yoel Roth 15:05:55

Label please



Patrick Conlon [redacted] 15:06:51

Done



[redacted] 15:31:53

Is it done?

replies to thread:



Wednesday, January 6th 2021 15.35.24 by

Patrick Conlon [redacted]

Looks like it's now bounced



[redacted] 15:39:17

It's the right decision. Yoel, I hope you, Del, and others are appropriately CorpSec'd

@mentions

This theme of Policy perhaps being stressed by queries from Communications executives – who themselves have to answer the public’s questions – occasionally appears. Two days later, you see chatter about pulling Comms out of the loop:



12:24:23

We are already getting questions about conspiracy theories and our enforcement process all the time. Not a wait and see thing.



Lauren Culbertson 12:24:26

Let's talk about that but we need to hammer that out ahead of an emergency and evolving situation



12:24:57

Respectfully, it needs to be your team doing so holisitically.



Lauren Culbertson 12:25:27

I don't think it's wise to reach out to comms folks proactively on this

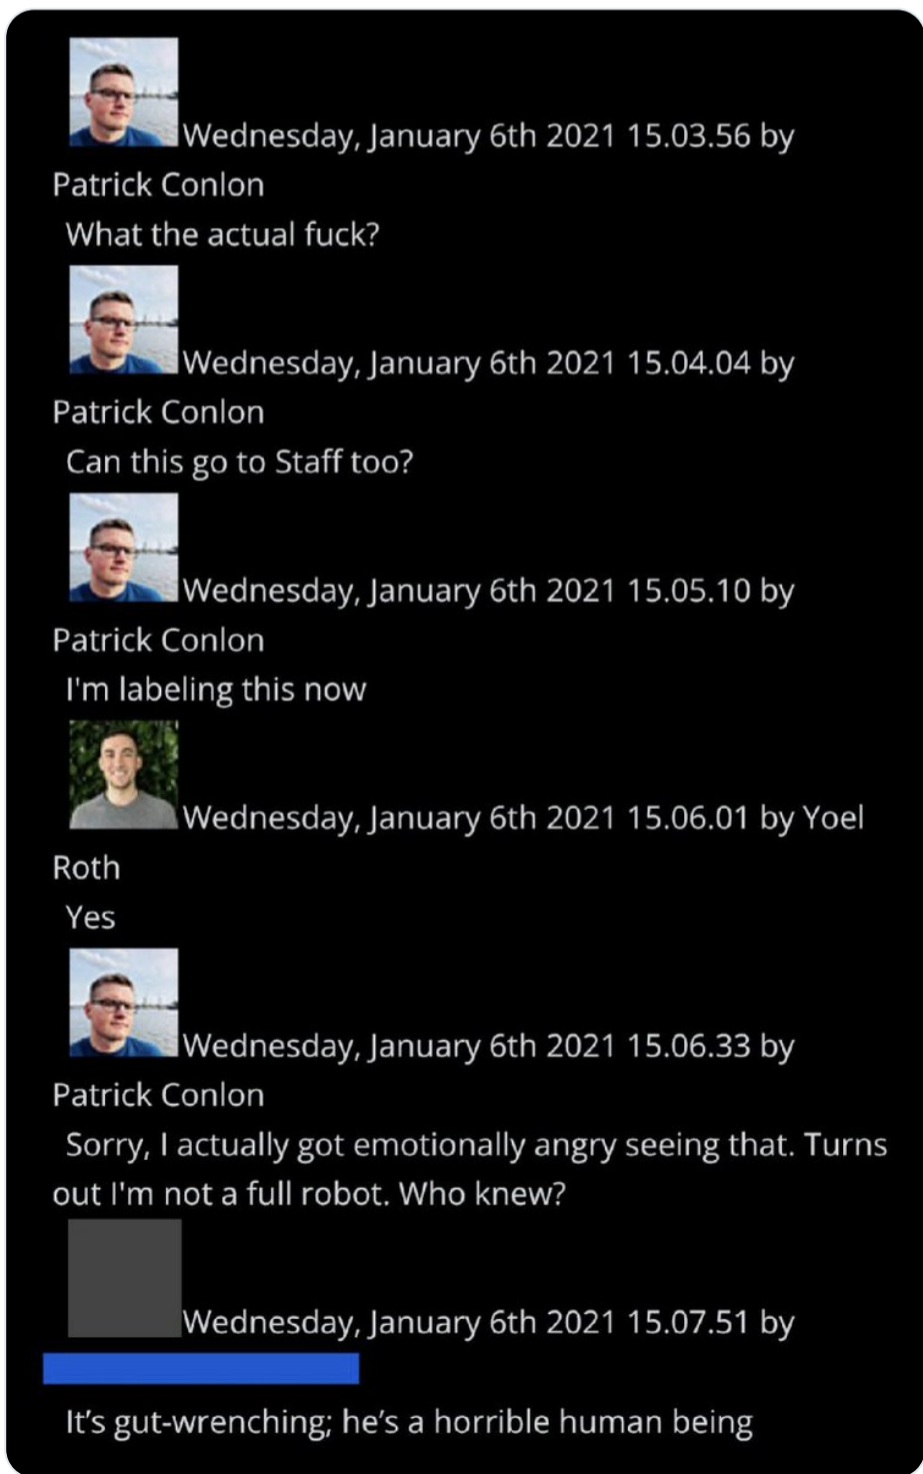
@mentions

61. The first company-wide email from Gadde on January 6th announced that 3 Trump tweets had been bounced, but more importantly signaled a determination to use legit “violations” as a guide for any possible permanent suspension:

Here's the email from Vijaya: ``Hi team, An update on actions we're taking as a result of what's happening in DC today. We have removed (bounced) 3 tweets from @realDonaldTrump on 6 January 2020 for making unfounded claims of voter fraud and election theft, which must be viewed in the context of violence in Washington DC today and reasonably can be expected to continue to incite and inspire violent acts around the country. We are requiring a 12-hour timeout following the removal of these tweets. We also have tweeted so there's transparency in our actions and making it clear that future violations of the Twitter Rules will result in permanent suspension. {https://twitter.com/TwitterSafety/status/1346970430062485505} Thanks, Vijaya``

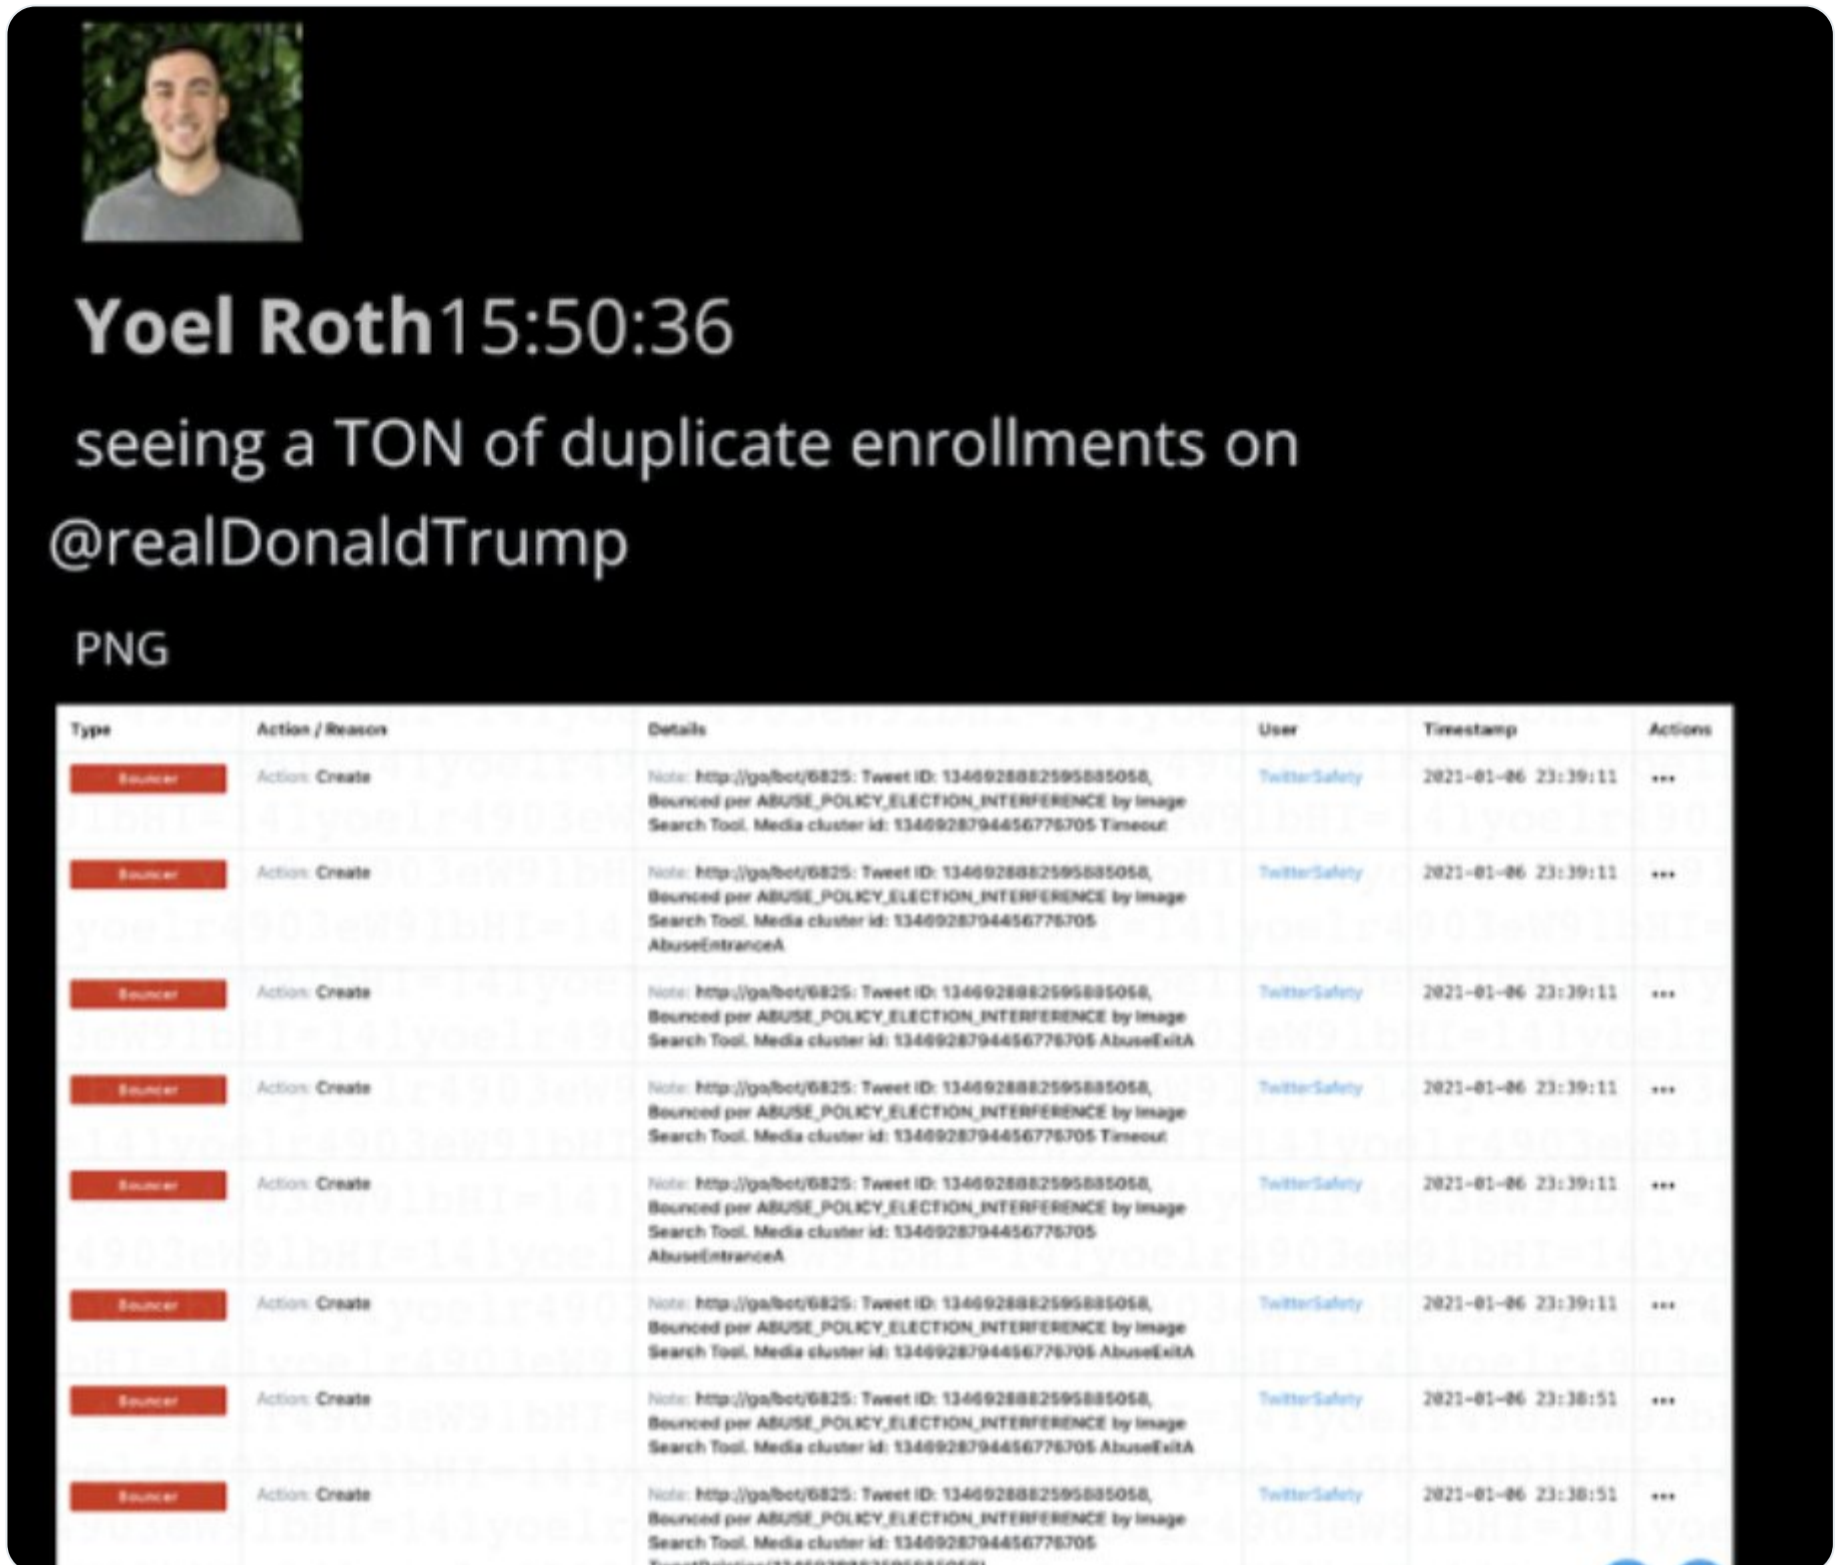
@mentions

62. "WHAT THE ACTUAL FUCK?" Safe to say Trump's "Go home with love & in peace" tweet mid-riot didn't go over well at Twitter HQ:



@mentions

63. A few last notes about January 6th. Roth at one point looked and found Trump had a slew of duplicate bot applications:



@mentions

64. By the end of the first day, the top execs are still trying to apply rules. By the next day, they will contemplate a major change in approach. Watch this weekend for the play-by-play of how all that went down.

@mentions

65. By January 8th, which will describe Sunday, Twitter will be receiving plaudits from “our partners” in Washington, and the sitting U.S. president will no longer be heard on the platform.

@mentions

66. Lastly, people on the left, right, and in between want to know what else is in the [#TwitterFiles](#), from suppression/shadow-banning of leftists to lab-leak theorists, or amplification of military propaganda or conservative accounts. We know everyone has questions.

@mentions

67. And while we’ve stumbled on tidbits here and there about topics ranging from COVID to foreign policy, the reality is the data sets are enormous and we’re still working through them.

More is coming. Good night, all.

