Center for American Progress

Results Not Found

Addressing Social Media's Threat to Democratic Legitimacy and Public Safety After Election Day

By Adam Conner and Erin Simpson   October 23, 2020

Update, October 27, 2020: Shortly after the publication of this issue brief, YouTube announced initial postelection policies to "prevent the abuse of [its] systems" and "provide access to authoritative information this election season." Google also updated its advertising policies for U.S. election ads and plans to implement a "sensitive events policy" after the polls close on November 3.

It may take days or weeks after Election Day to count the votes, declare winners in hotly contested races, and certify the results for the U.S. 2020 general election. And during this waiting period, social media platforms may be used in attempts to delegitimize the election process and cause chaos. Once all votes have been cast, social media platforms' priorities should shift from voter engagement to affirming democratic legitimacy and protecting public safety. This will require platforms to quickly remove content that baselessly attempts to delegitimize the election, stop disinformation about election results from going viral, and prevent their platforms from being used to facilitate threats of violence.

This issue brief offers substantive recommendations for product and policy interventions that social media companies should begin implementing immediately to prepare for the postelection period, from election night through the certified winners of the election assuming office.

During this time period, new threats will emerge as the information environment evolves with each phase of the election process—from initial counts and media projections to canvassing and recounts to the Electoral College and congressional certification. Malicious actors may use the information void created by the slower timeline of this year's initial results to spread false information. Candidates in local, state, or federal races might use confusion about the postelection timeline or process to inaccurately claim victory. Aggrieved groups may seek to organize disruption of recounts or the Electoral College process. Throughout the process, social media platforms—such as YouTube, TikTok, Nextdoor, Facebook, Reddit, Pinterest, Snapchat, and Twitter—will likely be used to sow confusion, stoke conflict, and further attempts to delegitimize the election.

Many Americans have already experienced unprecedented online chaos in their social media feeds during the lead-up to Election Day. Without proper planning and preparation by social media companies, it is easy to imagine continued confusion and disinformation after the polls close, with potentially severe consequences for public safety and democratic outcomes.

Background

The coronavirus pandemic is expected to result in record levels of votes cast by mail.[1] The public may face an extended period of counting, recounting, and certification of votes, potentially creating weeks of time before election outcomes in state, federal, and local races are determined.[2] The threats the United States has seen preelection—including disinformation, attempts to delegitimize the election, and calls for violent behavior—are not likely to cease but may evolve to exploit the uncertainty that lingers in the days and potentially weeks following Election Day.

Since the U.S. 2016 general election, the conversation about preventing social media platforms from once again becoming a threat to the democratic process has focused primarily on concerns and threats leading up to the election. While the companies operating the social media platforms are doing more than ever to surface essential information about how to vote and prevent foreign interference, the product features that aid and abet democratic threats have yet to be reined in. Social media companies should make every possible effort to prevent their platforms from contributing to voter suppression and calls to violence. Proposals from the Brennan Center for Justice, Stop Online Violence Against Women, Stop Hate for Profit, Accountable Tech, New America's Open Technology Institute, Facebook's Civil Rights Audit, the Berkman Klein Center for Internet and Society, and others have put forth numerous preelection resources to that end.[3] These efforts range from tabletop exercises that stress-test existing social media content policies[4] to recommendations on preventing demographically targeted voter suppression campaigns.[5]

Lesser attention, though, has been paid to the time period after polls close on election night. Consideration of postelection activities is especially warranted by 1) this year's slower election results timeline; 2) concerted efforts by bad actors, including government officials, to preemptively delegitimize election results;[6] and 3) the United States' complex election processes.

Indeed, U.S. election processes vary substantially across state and local governments and up and down the ticket. The presidential election is particularly complex: The winner is determined not by popular vote, but rather a combination of outcomes across states, which are then inputted into the electoral

college—a body that convenes to select the next president and whose selection is subsequentially submitted for certification by a joint session of Congress. Social media platform policies should consider the postelection period for the U.S. presidential election to be from after the polls close on November 3 to the inauguration on January 20. This time will likely be comprised of four distinct periods, which could progress at different speeds for different states: 1) after the polls close while votes are initially being counted, 2) initial public results posted from election officials and initial declarations of winners from media organizations, 3) the period of election certification, including legal challenges and recounts, and 4) the electoral college process for the presidential election. This 2 1/2-month period may include specialized events such as recounts, legal challenges, and potential issues with the electoral college process.

In the postelection period, candidates, government officials, and other actors—domestic or foreign—may seize on the uncertainty to baselessly delegitimize the results, prematurely declare victory, or mobilize supporters to interrupt legitimate election processes or commit violence. Social media platforms need heightened rules and scrutiny, not just ahead of elections but throughout the postelection period, and they should actively coordinate with each other to address these issues. The prevention of violence and protection of democratic legitimacy must be the guiding values for platforms once polls close on November 3. Only weeks before the U.S. 2020 general election, no social media platforms had sufficient standards for grappling with election delegitimization attempts and postelection conflict.[7]

As companies rapidly develop dedicated policies for this time period, the Center for American Progress urges them to consider the suggestions below.

Evolving threats once the polls close

Once polls close, voters are no longer seeking information to inform their vote, and as such, there is a compelling public interest in aggressive removal of inaccurate or inflammatory content that seeks to delegitimize the election. Such action cannot affect votes that have already been cast. After voting ends, platforms' content moderation choices can instead influence perception around the process of counting votes and the legitimacy of the election—and, in turn, what actions individuals may choose to exercise, such as taking to the streets to protest—but can no longer influence a voter's choice of candidate because there is no longer the ability to cast a vote.

At present, social media sites are prioritizing nonintervention toward politicians, even going so far as to exempt content from politicians and elected officials from their community guidelines and fact-checking.[8] For example, Facebook has exempted politicians' ads and most posts from standard fact-checking processes.[9] Platforms have justified nonintervention on the basis of a narrow interpretation

of freedom of expression while ignoring the numerous associated harms, such as harassment, hate speech, voter suppression, and violence. While these companies may have underused their content moderation capabilities due to the concerns in the lead-up to the election, they should reevaluate the calculus of risk management in the postelection period, which demands aggressive action.

Furthermore, platforms must stop their services from being a vector for violence before, during, and after the election. Just this month, an effort to kidnap Michigan Gov. Gretchen Whitmer (D) and take "violent action against multiple state governments" had been at least partially organized on Facebook before being disrupted by the FBI.[10] In a heightened and increasingly polarized political environment, social media platforms will have to not only remove content and accounts that incite or inflict violence but also take proactive action to detect and disrupt activity that could lead to violence.

Existing rules on election delegitimization

In August, the Election Integrity Partnership's initial analysis of social media policies around elections noted that none of the platforms studied had "clear, transparent" policies on how they would respond to attempts at election delegitimization.[11] As of mid-October, however, several platforms had introduced policies addressing these issues; and Facebook and Twitter, who were first out with specific policies, further developed their standards. While standards are only the first step of effective action—accompanying enforcement is required and yet, due to most platforms' opacity, impossible to evaluate—they are an essential foundation.

Facebook and Twitter have developed content moderation labels that have been increasingly deployed around COVID-19 misinformation and election disinformation. Content moderation labels can take a variety of forms—ranging from discrete text and icons appended to a social media post to labels that cover and obscure a post until a reader clicks through to view the original content. Label text varies from the generic and discrete to specific, direct refutations of the original content. Labels alone are often insufficient to prevent the spread of misinformation unless they are accompanied by algorithmic changes such as downranking or the disabling of sharing features.

For example, Facebook and Instagram—the latter of which is owned by Facebook, Inc.—have stated that they will attach an informational label providing "authoritative information about the integrity of the election and voting methods" on "content that seeks to delegitimize the outcome of the election or discuss the legitimacy of voting methods, for example, by claiming that lawful methods of voting will lead to fraud."[12] Facebook's labels have thus far tended toward discrete; but direct, rather than generic, text has been applied to election delegitimization claims. And little has been shared on potential algorithmic limitations that may

accompany labels and whether limitations would apply to content from a politician or elected official who is not subject to fact-checking.

In October, Facebook published additional details of their postelection information efforts, noting that "if a candidate or party declares premature victory before a race is called by major media outlets, we will add more specific information in the notifications that counting is still in progress and no winner has been determined."[13] The company also announced its intention to post an announcement at the top of Facebook and Instagram if a winner is declared by major media outlets. Finally, the company announced advertising restrictions, including prohibitions on political advertisements that make "premature claims of victory" or "delegitimize the outcome of an election" as well as a pause on running all social issue, electoral, or political ads after polls close to "reduce opportunities for confusion or abuse."[14] Their conditions for ending such a pause are unclear. This is an important step for paid content, even though it is expected to play a much smaller role than organic content in any potential postelection delegitimization efforts.
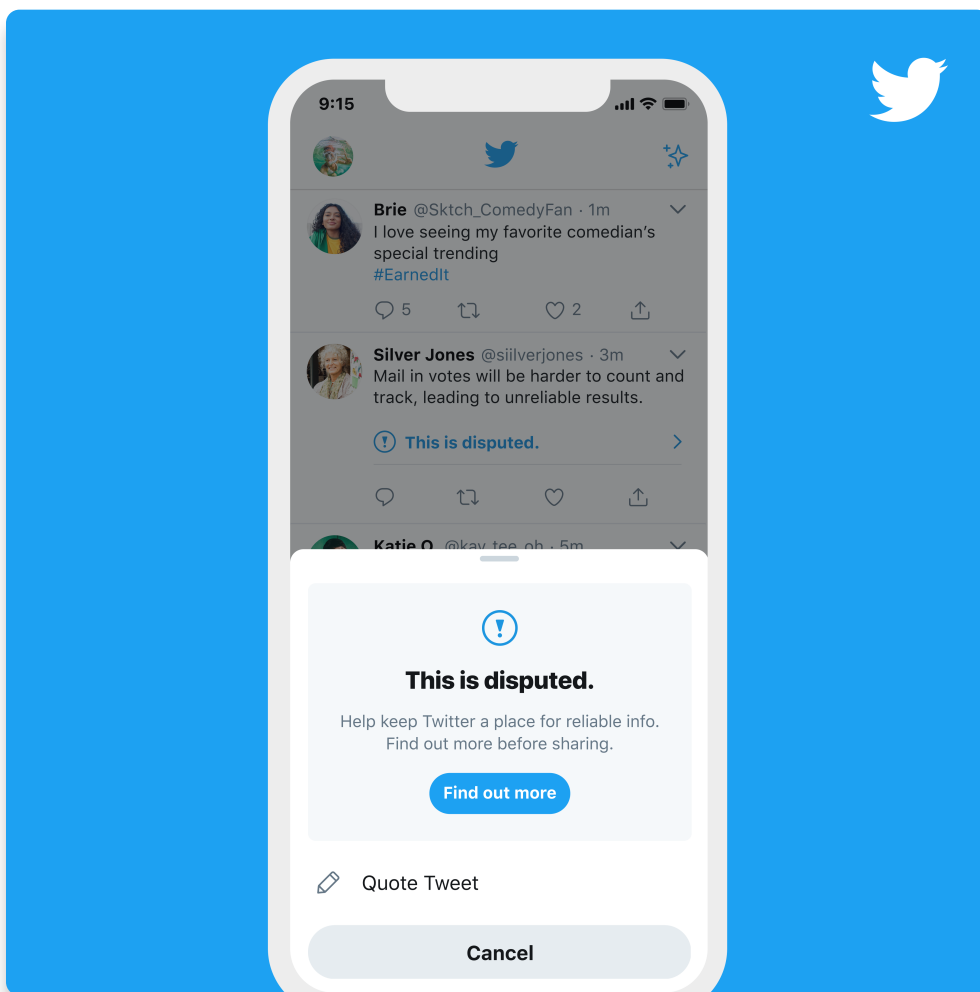

Image of Twitter's new credible information prompt for attempted retweets of misleading information.
Credit: Twitter

Similarly, Twitter has updated its civic integrity policy and gone one step further in that it "will label or remove false or misleading information intended to undermine public confidence in an election or other civic process."[15] An October update noted that Twitter "will label Tweets that falsely claim a win for any candidate and will remove Tweets that encourage violence or call for people to interfere with election results or the smooth operation of polling places" and that "people on Twitter, including candidates for office, may not claim an election win before it is authoritatively called."[16] Tweets labeled under this policy are de-amplified in algorithmic recommendation systems and attempts to retweet these posts will now be intercepted by an interstitial prompt pointing to credible information on the topic before a user can proceed. (see image above) Importantly, Twitter has also clarified that its civic integrity policies will apply to world leaders and has enforced them throughout 2020.[17] Indeed, Twitter's October update adds extra steps for misleading information from U.S. political figures, high-reach accounts, and high-engagement accounts: If given a misleading information rating, tweets from these groups will be labeled and covered such that a user has to click through a warning to see the original content, and engagement options will be limited. Twitter has previously decided to prohibit all political advertisements on its platform.[18]

As rated by the Election Integrity Partnership, Twitter has the most comprehensive policies on election delegitimization thus far developed, followed by Facebook. Meanwhile, YouTube, Pinterest, Nextdoor, Snapchat, and TikTok have more "non-comprehensive policies," and Reddit has no dedicated policies.[19] Google did announce, however, that it would be prohibiting political advertising on YouTube for at least a week after polls close on election night.[20] Given the size and reach of YouTube, it is important that this platform further develops and announces these policies ahead of the election. Hopefully, in the weeks ahead, all of the platforms will adopt clearer, more direct policies along these lines.

Recommendations for social media platforms

The companies that operate some of the United States' largest social media platforms must rapidly develop strong rules, proactive policy changes, and more effective enforcement mechanisms to prevent their products from being used to harm democratic legitimacy and incite violence following the election. Below are a set of recommended general product and rule changes for the postelection period, which can be adapted to the unique features and operation of various platforms.

Remove posts that baselessly delegitimize
Social media platforms need to remove information that baselessly delegitimizes the election. Labeling is insufficient in preventing platform affordances from being used to destabilize the election: Malicious actors seek to sow doubt and confusion,

not necessarily persuade with facts. Fact checks and explicit, bold labels do not prevent the distribution of information that seeks to delegitimize the election. Moreover, as candidates aim to paint themselves as unfairly treated, labels themselves can be inappropriately weaponized as false evidence that the system is against them.

In order to effectively mitigate election delegitimization attempts, social media platforms must develop careful standards for its definition, preferably in coordination with one another and with advance input of democracy experts and representatives from civil society.

If platforms will not remove baseless delegitimization, they should at least obscure the content of such posts behind a strong warning label and reduce algorithmic amplification. As noted, Twitter has led the way in developing labels that also reduce algorithmic distribution. And if platforms will not impose click-through covers and algorithmic de-amplification on delegitimizing posts, they should at least develop visually bold labels that clearly and plainly contest the content. Such labels are preferred to discrete, minimalist labels that subtly contest or generically label the claims. Even if such policies are not universally applied to social media users, they must at least be applied to the social media accounts of candidates themselves. Legitimate journalism covering postelection events remains essential and should be unaffected by new policies.

Develop consistent, collaborative standards for determining election results

In order to effectively moderate disinformation around election results, platforms should develop a standard, public methodology, potentially in collaboration with one another and with relevant experts. This standard should appropriately weigh primary sources to make this determination, including initial public vote counts from election officials as well as media outlets with specialized election expertise. Ideally, platforms would develop and publish these standards to minimize public confusion during the postelection period. However, a consistent internal standard is preferable to no standard.

Starting election night, election administrators post initial vote counts to their official websites. In some states, these initial counts will at first include only in-person votes, with mail-in ballots counted more slowly. Counts are later updated during the initial canvass and then certified as the final election results.

U.S. media outlets that feature around-the-clock coverage on election night rely on these initial public counts, in combination with several other sources, to publicly declare the outcome of races. The National Election Pool is a consortium consisting of ABC News, CBS News, CNN, and NBC News that includes national and state in-person exit polls, vote tabulations, and election projections distributed by Reuters.[21] Each individual news entity in the consortium receives

the same data but makes its own independent analysis and decision to make election declarations.[22] Facebook has partnered with the National Election Pool and Reuters for their election efforts.[23] The Associated Press maintains the only election service that collects and verifies election returns in more than 7,046 down-ticket races in addition to the presidential election.[24] Numerous tech and social media companies, including Google, have subscribed to the Associated Press' service in order to incorporate election results into their consumer-facing technology products.[25] Finally, Fox News maintains a decision desk that is not part of the National Election Pool.[26]

Platforms should craft a standard for election results in individual races that considers the following three factors, giving initial counts from election administrators the highest weight:

1. Initial public vote counts from election administrators, with a focus on total outstanding uncounted ballots—if the total number of outstanding votes is not enough to change the election results, the outcome may be clear even if all votes have not yet been initially counted

2. Declarations from organizations with professional election decision desks with which platforms have officially partnered

3. A majority of major news organizations with professional election decision desks having publicly declared a winner

Given the volatile election environment, it is critical to err on the side of caution. For example, in a state where a candidate is projected to win handily based on polling and initial results, the declaration of a winner by all of the major media organizations can be used as a standard. In states expected to be highly competitive, the combination of all three factors should be used before determining a winner.

This standard should be used to help determine the legitimacy of victory claims for elections whose winner can be determined immediately after Election Day, while also providing appropriate discretion for elections whose outcomes may take longer to determine. Twitter has recently announced a similar standard, stating that an election determination requires "either an announcement from state election officials, or a public projection from at least two authoritative, national news outlets that make independent election calls."[27]

Given the highly politicized decision-making environment for platforms, companies should consider collaboration, improving on the model developed by the Global Internet Forum to Counter Terrorism, which coordinates anti-terrorism efforts between major internet platforms, including Facebook, YouTube, and Pinterest.

A similar approach was recently suggested by David Kaye, former U.N. special rapporteur on freedom of opinion and expression, to *The New York Times* editorial board.[28] He reasons: "If you had the platforms together making a statement of their values, then when they take action, it creates a permission structure for reticent platform executives to make difficult decisions quickly." Transparent, consistent standards adopted by a group of major platforms would mirror the strategy adopted by major media organizations in their long-standing agreement not to declare a winner in the presidential race until polls in the continental 48 states have closed. While there are some concerns regarding this approach to platform collaboration— for example, those raised by Evelyn Douek's scholarship on "content cartels," which she describes as "arrangements between platforms to work together to remove content or actors from their services without adequate oversight"[29]—transparent standards that enable action are preferable to the status quo.

Fact-check election result claims

As noted above, incorrect claims that seek to delegitimize the election should ideally be removed. Beyond these cases, platforms with fact-checking programs should fact-check all claims about election results, regardless of source, according to the standards outlined above. As proposed by experts at Avaaz, when users interact with or view a post that is later fact-checked, platforms should "correct the record" by notifying users of the associated fact check.[30] Moreover, because of the nature of election disputes, where candidates are increasingly willing to decry the electoral process outcome even if the process is fair, it is critical that platforms provide additional context where possible in clear and plain terms, in multiple languages, and accessible to screen readers or other accessibility aids.

The fact-checking of results claims published on social media sites should include opinion pieces and, especially, content from politicians. Platforms should consider funding fact checkers with specific expertise on elections and legal process to ensure the most accurate interpretation and context. As processes progress, fact checks will likewise need to target not only claims about results but also claims about recounts, the Electoral College, and other election processes. Platforms may need to engage the services of specialized partners to successfully fact-check the range of claims that will arise across various states and up and down the ballot in the postelection period.

Build viral circuit breakers

Many strategies for tackling harmful published information are reactive, yet once a harmful post is seen, the damage has already been done. Even with fact-checking and corrections, the lie travels farther than the truth.[31] In the short term, platforms need to take additional steps to help prevent false, harmful information from going viral in the first place. Experts have long recommended adding more context and strategically increasing friction—anything that inhibits user action within a digital interface—within social media products. Building on an idea first

raised by professor Ellen Goodman,[32] the Center for American Progress previously proposed a viral circuit breaker for disinformation around the coronavirus crisis: Social media platforms would program a pause in the algorithmic amplification of fast-growing content about the coronavirus and prioritize that content for rapid human review.[33] Much of this content would be flagged and reviewed by moderators anyway, but evidenced by the false coronavirus conspiracies, even swift review can be too late if algorithms are already working to amplify and recommend the harmful content.

Slowing the spread of fast-growing content about the election in order to perform effective review before it generates millions of views could help stop the spread of unfounded claims and impede efforts to sow doubt or cause chaos. Imagine a circuit breaker that is triggered whenever a potentially false claim about election results begins to go viral. Platforms would prioritize the content in internal human review and fact-checking, post a warning sign that it has yet to be verified, and suspend amplification in recommendation algorithms—while allowing individual posting and message sharing to continue—until it is reviewed.

Before the election, all platforms should prioritize the development and transparent deployment of content moderation systems in the spirit of viral circuit breakers and keep them active through the postelection period. Such a product alone would not be sufficient unless paired with updated content moderation procedures and increased, specialized staffing around the clock. Sufficient staffing increases to make rapid, context-specific calls are needed not only in the lead-up to the election and on election night but also throughout the postelection period in the following weeks and months.

Take steps to prevent violence on and off platform

Most online platforms generally only take action against violence-provoking content when it is posted on their platform or when it is linked to from their platform.[34] The promotion or glorification of violence rightly triggers the harshest penalties from these platforms, such as account suspension or removal. The Change the Terms coalition, of which the Center for American Progress is a co-chair, has long called for this: "Terms of service or acceptable use policies should, at a minimum, make it clear that using the service to engage in hateful activities on the service or to facilitate hateful activities off the service shall be grounds for terminating the service for a user."[35] However, when a group or actor calls for, organizes, or commits an act of violence, social media platforms should not wait for actors to post about that violence on their site specifically in order to violate their terms. This obviously applies when an act of violence is committed, but it also applies when violence is being called for or organized elsewhere. Platforms should take swift and proactive action to remove accounts, groups, networks, or events associated with acts of violence generally, not just violent content posted to their platforms specifically.

For example, if an individual or group makes a call to arms in a video or podcast, accounts of those involved should be removed on social media platforms, even before such content is reposted on platforms.[36] Malicious actors are experts in toeing the line on major platforms' terms, as they do not want to lose their megaphone to recruit and radicalize wider audiences. By proactively removing accounts associated with calls to violence elsewhere online, platforms can help prevent those efforts. Such a policy should apply as soon as possible and continue during the postelection period.

Interventions along the lines of a viral circuit breaker are ill-suited for social media accounts with millions of followers, whose posts are highly visible from the start. For these, platforms should take a cue from live network television, in which live feeds are put on short delays to prevent unacceptable content from being broadcast. Accountable Tech developed a proposal for a violence prevention preclearance system, in which social media platforms institute a short delay on high-reach accounts that have previously been sanctioned for election misinformation to manually review posts for content that instigates violence.[37] Platforms should implement such a system around calls to violence from when voting starts until certified winners of the election assume office.

Platforms should also go beyond removal and share information about potential threats and build relationships with relevant state and local officials, including governors, mayors, election officials, and state attorneys general. Platforms must have a heightened awareness that their organizing tools could be used to organize physical interference with election processes ranging from voting to vote counting and certification.

Build shutoff switches for product features that may contribute to violence

Finally, social media platforms should begin building "shutoff switches" for product features, such as Facebook's group recommendations, that could be used to organize violent action and/or attempts to baselessly contest the election. Many social media companies have never anticipated this need and have not built corresponding shutoff switches into their products for such an emergency; platforms must plan now to be able to instantly turn off any such features. Social media sites should build the ability to isolate and pause feature sets such as group recommendations,[38] local events recommendations, video recommendations, trending topics, or others that are identified as catalysts of ongoing problems.

Twitter has preemptively addressed this concern for some features by disabling the surfacing of "liked by" and "followed by" growth features and severely restricting content in the "For You" tab from before the election, beginning October 20, to a period of at least a week after Election Day.[39] This is in recognition that the use of such features has the potential to overwhelm timely moderation efforts at a critical moment and that the most responsible choice is to disable them, at least temporarily. Other platforms should take note and follow Twitter's leadership.

Certainly, some of these tools allow users to carry out legitimate election or organizing processes and could have unintended consequences in temporarily restricting other areas of life online. However, in a scenario in which social media could be utilized for widespread chaos, violence, or democratic harms, platforms should have the ability to temporarily pause feature sets until they can address the problem. Hopefully, it will not become necessary to undertake this kind of worst-case scenario, but social media platforms are urged to make advance preparations. In any situation, there will be a continued need to keep verified, authoritative, localized information prominently available.

Conclusion

The coronavirus crisis has catalyzed a historic shift in how Americans cast their votes, creating a longer election results timeline that may come as a surprise to many. Traditional expectations of election processes and election night results may result in confusion that could easily be weaponized by actors attempting to disrupt democratic process, especially in the information void created by the additional time that may be required to count all votes.

Social media companies must make every effort to prevent their products from being used in attempts to delegitimize the election or threaten public safety—both in the lead-up to the election and after the polls close on November 3, up until any newly elected officials have taken office. These challenges demand even more decisive action from platforms once the polls close.

Adam Conner is the vice president for Technology Policy at the Center for American Progress. Erin Simpson is the associate director of Technology Policy at the Center.

The authors wish to thank the following people for their time and expert input in the development of this work: Sam Berger, Nicole Gill, Jesse Lehrich, Jesse Littlewood, Ben Olinsky, Will Ragland, Danielle Root, Gus Rossi, Mara Rudman, Melissa Ryan, Michael Sozan, Sabrina Joy Stevens, Alex Tausanovitch, Dave Toomey, Alex Witt, Kip Wainscott, and our colleagues with CAP's Editorial and Art teams.

Endnotes

1 Adam Levy, Ethan Cohen, and Liz Stark, "Surge of ballot requests already setting records in the US," CNN, September 25, 2020, available at https://www.cnn.com/2020/09/25/politics/ballot-requests-voting-election-2020/index.html.

2 Brennan Center for Justice, "How Election Night 2020 Will Be Different," September 18, 2020, available at https://www.brennancenter.org/our-work/research-reports/how-election-night-2020-will-be-different; Elaine Kamarck, "Election night? Think again—election month might be more like it," Brookings Institution, August 4, 2020, available at https://www.brookings.edu/blog/fixgov/2020/08/04/election-night-think-again-election-month-might-be-more-like-it/.

3 Ian Vandewalker, "Digital Disinformation and Vote Suppression," Brennan Center for Justice, September 2, 2020, available at https://www.brennancenter.org/our-work/research-reports/digital-disinformation-and-vote-suppression; Stop Online Violence Against Women, "A threat to an American Democracy: Digital Voter Suppression" (2020), available at http://stoponlinevaw.com/wp-content/uploads/2020/02/7.pdf; Stop Hate for Profit, "Recommended Next Steps," available at https://www.stophateforprofit.org/productrecommendations (last accessed October 2020); Accountable Tech, "Election Integrity Roadmap for Social Media Companies," available at https://accountabletech.org/election-integrity-roadmap-for-social-media-companies/ (last accessed October 2020); Spandana Singh and Margerite Blase, "Protecting the Vote: How Internet Platforms Are Addressing Election and Voter Suppression-Related Misinformation and Disinformation" New America Open Technology Institute, September 30, 2020, available at https://www.newamerica.org/oti/reports/protecting-vote/; Laura W. Murphy and others, "Facebook's Civil Rights Audit – Final Report" (Menlo Park, CA: Facebook, 2020), available at https://about.fb.com/wp-content/uploads/2020/07/Civil-Rights-Audit-Final-Report.pdf; Oumou Ly and Jorhena Thomas, "US Elections Disinformation Tabletop Exercise Package" (Cambridge, MA: Berkman Klein Center for Internet and Society, 2020), available at https://dash.harvard.edu/handle/1/37365565.

4 Ly and Thomas, "US Elections Disinformation Tabletop Exercise Package."

5 Stop Online Violence Against Women, "A threat to an American Democracy."

6 Yochai Benkler and others, "Mail-In Voter Fraud: Anatomy of a Disinformation Campaign" (Cambridge, MA: Berkman Klein Center for Internet and Society, 2020), available at https://cyber.harvard.edu/publication/2020/Mail-in-Voter-Fraud-Disinformation-2020.

7 Election Integrity Partnership, "Evaluating Platform Election-Related Speech Policies," October 19, 2020, available at https://www.eipartnership.net/policy-analysis/platform-policies; Singh and Blase, "Protecting the Vote"; Mozilla Foundation, "US Election 2020: Platform Misinformation Policies Tracker," available at https://docs.google.com/document/d/1-b5b_AgJhz7rAXaSpekwBwZ-WHx_A7iSleoOGxK0lZto/edit (last accessed October 2020).

8 Nick Clegg, "Facebook, Elections and Political Speech," Facebook, September 24, 2019, available at https://about.fb.com/news/2019/09/elections-and-political-speech/; Tony Romm, "Zuckerberg: Standing For Voice and Free Expression," The Washington Post, October 17, 2019, available at https://www.washingtonpost.com/technology/2019/10/17/zuckerberg-standing-voice-free-expression/.

9 Mark Zuckerberg, Facebook, June 26, 2020, available at https://www.facebook.com/zuck/posts/10112048980882521.

10 Craig Timberg and Isaac Stanley-Becker, "Michigan kidnapping plot, like so many other extremist crimes, foreshadowed on social media," The Washington Post, October 8, 2020, available at https://www.washingtonpost.com/technology/2020/10/08/michigan-plot-kidnapping-boogaloo-socialmedia/.

11 Election Integrity Partnership, "Evaluating Platform Election-Related Speech Policies."

12 Facebook, "New Steps to Protect the US Elections," September 3, 2020, available at https://about.fb.com/news/2020/09/additional-steps-to-protect-the-us-elections/.

13 Guy Rosen, "Preparing for Election Day," Facebook, October 7, 2020, available at https://about.fb.com/news/2020/10/preparing-for-election-day/.

14 Ibid.; Zuckerberg, Facebook, June 26, 2020.

15 Twitter, "Expanding our policies to further protect the civic conversation," September 10, 2020, available at https://blog.twitter.com/en_us/topics/company/2020/civic-integrity-policy-update.html.

16 Vijaya Gadde and Kayvon Beykpour, "Additional steps we're taking ahead of the 2020 US Election," Twitter, October 9, 2020, available at https://blog.twitter.com/en_us/topics/company/2020/2020-election-changes.html.

17 Twitter, "World Leaders on Twitter: principles & approach," October 15, 2020, available at https://blog.twitter.com/official/en_us/topics/company/2019/worldleaders2019.html.

18 Kate Conger, "Twitter Will Ban All Political Ads, C.E.O. Jack Dorsey Says," The New York Times, October 30, 2019, available at https://www.nytimes.com/2019/10/30/technology/twitter-political-ads-ban.html.

19 Election Integrity Partnership, "Evaluating Platform Election-Related Speech Policies."

20 Sara Fischer, "Scoop: Google to block election ads after Election Day," Axios, September 25, 2020, available at https://www.axios.com/google-to-block-election-ads-after-election-day-4b60650d-b5c2-4fb4-a856-70e30e19af17.html.

21 Reuters, "Reuters to distribute real-time 2020 U.S. presidential election results and exit polls from the National Election Pool in groundbreaking collaboration," July 22, 2019, available at https://www.reuters.com/article/rpb-nepedison/reuters-to-distribute-real-time-2020-u-s-presidential-election-results-and-exit-polls-from-the-national-election-pool-in-groundbreaking-collaboration-idUKKCN1UH1C2.

22 John Lapinski, Katie Primm, and Stephanie Perry, "NBC News Decision Desk: How we call the midterm races on election night," NBC News, November 5, 2018, available at https://www.nbcnews.com/politics/elections/nbc-news-decision-desk-how-we-call-midterm-races-election-n930806.

23 Facebook, "New Steps to Protect the US Elections."

24 Associated Press, "Election Services," available at https://www.ap.org/en-us/topics/politics/elections/election-services (last accessed October 2020).

25 Sara Fischer, "AP to call elections for Alexa and other Big Tech channels," Axios, October 20, 2020, available at https://www.axios.com/ap-to-call-elections-for-alexa-and-other-big-tech-channels-13b35334-a728-42bb-a33e-4b1e242792b6.html?utm_campaign=organic&utm_medium=socialshare&utm_source=twitter.

26 Ben Smith, "Trump Wants to Discredit the Election. This Nerd Could Stop Him.", The New York Times, September 27, 2020, available at https://www.nytimes.com/2020/09/27/business/media/trump-election-fox-news.html.

27 Gadde and Beykpour, "Additional steps we're taking ahead of the 2020 US Election."

28 The New York Times, "What's the Plan if Trump Tweets That He's Won Re-election?", September 27, 2020, available at https://www.nytimes.com/2020/09/27/opinion/social-media-trump-election.html.

29 Evelyn Douek, "The Rise of Content Cartels" (New York: Knight First Amendment Institute at Columbia University, 2020), available at https://knightcolumbia.org/content/the-rise-of-content-cartels.

30 Avaaz, "White Paper: Correcting the Record" (London, Canada: 2020), available at https://fb.avaaz.org/campaign/en/correct_the_record_study/.

31 Soroush Vosoughi, Deb Roy, and Sinan Aral, "The spread of true and false news online," *Science* 359 (6380) (2018): 1146–1151, available at https://doi.org/10.1126/science.aap9559; Brendan Nyhan and Jason Reifler, "Misinformation and Fact-checking: Research Findings from Social Science" (Washington: New America Foundation, 2012), p. 28, available at https://www.dartmouth.edu/~nyhan/Misinformation_and_Fact-checking.pdf; Michelle Ma, "Hold that RT: Much misinformation tweeted after 2013 Boston Marathon bombing," UW News, March 17, 2014, available at https://www.washington.edu/news/2014/03/17/hold-that-rt-much-misinformation-tweeted-after-2013-boston-marathon-bombing/; See quote from Dolores Albarracín in Shannon Bond, "Can Circuit Breakers Stop Viral Rumors On Facebook, Twitter?", NPR, September 22, 2020, available at https://www.npr.org/2020/09/22/915676948/can-circuit-breakers-stop-viral-rumors-on-facebook-twitter.

32 Ellen P. Goodman, "Digital Information Fidelity and Friction" (New York: Knight First Amendment Institute at Columbia University, 2020), available at https://knightcolumbia.org/content/digital-fidelity-and-friction.

33 Erin Simpson and Adam Conner, "Fighting Coronavirus Misinformation and Disinformation: Preventive Product Recommendations for Social Media Platforms" (Washington: Center for American Progress, 2020), available at https://www.americanprogress.org/issues/technology-policy/reports/2020/08/18/488714/fighting-coronavirus-misinformation-disinformation/.

34 Twitter, "Our approach to blocking links," July 2020, available at https://help.twitter.com/en/safety-and-security/phishing-spam-and-malware-links.

35 Change the Terms, "Reducing Hate Online," available at https://www.changetheterms.org/ (last accessed July 2020); Center for American Progress and others, "Recommended Internet Company Corporate Policies and Terms of Service to Reduce Hateful Activities" (Washington: 2018), available at https://assets.website-files.com/5bba6f4828dfc3686095bf6b/5bd0e36186e28d35874f0909_Recommended%20Internet%20Company%20Corporate%20Policies%20%20Terms%20of%20Service_final-10-24.pdf.

36 Change The Terms, @changeterms, August 12, 2020, 11:21 a.m. ET, Twitter, available at https://twitter.com/changeterms/status/1293570827825434624.

37 Accountable Tech, "Election Integrity Roadmap for Social Media Companies."

38 Accountable Tech, "Tell Facebook: Stop All Group Recommendations," 2020, available at https://accountabletech.org/campaign/stop-group-recs/ (last accessed October 2020).

39 Gadde and Beykpour, "Additional steps we're taking ahead of the 2020 US Election."