



DIVISION OF
CORPORATION FINANCE

UNITED STATES
SECURITIES AND EXCHANGE COMMISSION
WASHINGTON, D.C. 20549

April 19, 2019

Pamela L. Marcogliese
Cleary Gottlieb Steen & Hamilton LLP
pmarcogliese@cgsh.com

Re: Alphabet Inc.
Incoming letter dated February 5, 2019

Dear Ms. Marcogliese:

This letter is in response to your correspondence dated February 5, 2019 and March 25, 2019 concerning the shareholder proposal (the "Proposal") submitted to Alphabet Inc. (the "Company") by the New York State Common Retirement Fund et al. (the "Proponents") for inclusion in the Company's proxy materials for its upcoming annual meeting of security holders. We also have received correspondence on the Proponents' behalf dated March 11, 2019 and April 1, 2019. Copies of all of the correspondence on which this response is based will be made available on our website at <http://www.sec.gov/divisions/corpfin/cf-noaction/14a-8.shtml>. For your reference, a brief discussion of the Division's informal procedures regarding shareholder proposals is also available at the same website address.

Sincerely,

M. Hughes Bates
Special Counsel

Enclosure

cc: Sanford J. Lewis
sanfordlewis@strategiccounsel.net

April 19, 2019

Response of the Office of Chief Counsel
Division of Corporation Finance

Re: Alphabet Inc.
Incoming letter dated February 5, 2019

The Proposal requests that the Company issue a report reviewing the efficacy of its enforcement of Google's terms of service related to content policies and assessing the risks posed by content management controversies related to election interference, freedom of expression and the spread of hate speech to the Company's finances, operations and reputation.

We are unable to concur in your view that the Company may exclude the Proposal under rule 14a-8(i)(3). We are unable to conclude that the Proposal, taken as a whole, is so vague or indefinite that it is rendered materially misleading. Accordingly, we do not believe that the Company may omit the Proposal from its proxy materials in reliance on rule 14a-8(i)(3).

We are also unable to concur in your view that the Company may exclude the Proposal under rule 14a-8(i)(10). Based on the information you have presented, it does not appear that the Company's public disclosures substantially implement the Proposal because, among other things, the disclosures do not address how the Company reviews the efficacy of the enforcement of Google's content policies across all its platforms or assesses the risks to the Company posed by the identified content management controversies. Accordingly, we do not believe that the Company may omit the Proposal from its proxy materials in reliance on rule 14a-8(i)(10).

Sincerely,

Courtney Haseley
Special Counsel

DIVISION OF CORPORATION FINANCE
INFORMAL PROCEDURES REGARDING SHAREHOLDER PROPOSALS

The Division of Corporation Finance believes that its responsibility with respect to matters arising under Rule 14a-8 [17 CFR 240.14a-8], as with other matters under the proxy rules, is to aid those who must comply with the rule by offering informal advice and suggestions and to determine, initially, whether or not it may be appropriate in a particular matter to recommend enforcement action to the Commission. In connection with a shareholder proposal under Rule 14a-8, the Division's staff considers the information furnished to it by the company in support of its intention to exclude the proposal from the company's proxy materials, as well as any information furnished by the proponent or the proponent's representative.

Although Rule 14a-8(k) does not require any communications from shareholders to the Commission's staff, the staff will always consider information concerning alleged violations of the statutes and rules administered by the Commission, including arguments as to whether or not activities proposed to be taken would violate the statute or rule involved. The receipt by the staff of such information, however, should not be construed as changing the staff's informal procedures and proxy review into a formal or adversarial procedure.

It is important to note that the staff's no-action responses to Rule 14a-8(j) submissions reflect only informal views. The determinations reached in these no-action letters do not and cannot adjudicate the merits of a company's position with respect to the proposal. Only a court such as a U.S. District Court can decide whether a company is obligated to include shareholder proposals in its proxy materials. Accordingly, a discretionary determination not to recommend or take Commission enforcement action does not preclude a proponent, or any shareholder of a company, from pursuing any rights he or she may have against the company in court, should the company's management omit the proposal from the company's proxy materials.

SANFORD J. LEWIS, ATTORNEY

Via electronic mail

April 1, 2019

Office of Chief Counsel
Division of Corporation Finance
U.S. Securities and Exchange Commission
100 F Street, N.E.
Washington, D.C. 20549

Re: Shareholder Proposal to Alphabet on Behalf of The New York State Common Retirement Fund and Others – Supplemental Reply

Ladies and Gentlemen:

The New York State Common Retirement Fund (the “Proponent”) is beneficial owner of common stock of Alphabet Inc. (the “Company”) and has submitted a shareholder proposal (the “Proposal”) to the Company together with co-lead-filer Natasha Lamb, of Arjuna Capital, on behalf of Lisa Stephanie Myrkalo and Andrea Louise Dixon. We previously responded to the Company’s February 5, 2019 no action request on March 11, 2019. I have been asked by the Proponent to respond to the Company’s supplemental letter dated March 25, 2019 (“Company Letter”) sent to the Securities and Exchange Commission by Pamela L. Marcogliese, of Cleary Gottlieb Steen & Hamilton LLP. A copy of this Supplemental Reply is being emailed concurrently to Ms. Marcogliese.

The record submitted by the Company and in our prior correspondence demonstrates that, regardless of whether email service of the Proposal was misdirected, the Company had received a hard copy of the amended Proposal on a timely basis, from both the Proponent and the co-lead-filer.

The vagueness arguments made in the Company Letter are of little weight or merit. The terms hate speech, freedom of expression, content management, content policies, etc. are all clear in the context used. Google has itself defined hate speech, free expression, deceptive behavior, etc. in its postings online.

The Proposal requests a report to shareholders reviewing the efficacy of the Company’s enforcement of Google’s terms of service related to content policies and assessing the risks posed by content management controversies related to election interference, freedom of expression, and the spread of hate speech, to the company’s finances, operations, and reputation. Reading the Proposal in its entirety the language is clear and would not be difficult for board, management or shareholders to understand.

The Company Letter also attempts to demonstrate that the Company has substantially

implemented the Proposal. However, the Company's existing reporting meets neither the essential purpose nor the guidelines of the Proposal. Missing is an assessment of the efficacy of implementation of policies to control and prevent hate speech and protect free expression.

While the Company Letter points to helpful published information on *methods used* for attempting to intercept hate speech and misinformation, the reported actions by which the Company notes it is "beginning to" address various issues do not provide an assessment of the efficacy of those actions.

Raw statistics on removals of postings and videos are not, by themselves, a dispositive indicator of efficacy. Assessing the timing and impact of removals may be more relevant to efficacy. The recent events surrounding the Christchurch mosque shooting in which hate speech was rapidly disseminated through internet channels (including Google's own YouTube) demonstrate the urgency of understanding whether and how the Company has the capacity to effectively intercept inappropriate speech. Moreover, some removals may raise issues relating to freedom of speech. Detailing how many videos have been taken down or comment opportunities foreclosed on an incidental basis does not adequately assess risks to freedom of expression posed by content management controversies.

An assessment of the efficacy of enforcement of the terms of service related to content management controversies on election interference, freedom of expression and hate speech requires, at a minimum, an examination of the scope of the problem and whether existing tools used by the company are an effective means of control. The Company's reported actions fall far short of providing an *assessment of the efficacy of current efforts* and resulting vulnerabilities "*assessing the risks posed by content management controversies* related to election interference, freedom of expression, and the spread of hate speech, *to the company's finances, operations, and reputation.*"

Accordingly, we believe there is no basis for the Company's claims that the Proposal is excludable pursuant to Rule 14a-8 and we urge the Staff to notify the Company that its request for a no action letter is denied.

Sincerely,

A handwritten signature in black ink, appearing to read "Sanford Lewis". The signature is stylized and cursive, with a long horizontal flourish extending to the right.

Cc: Pamela L. Marcogliese

CLEARY GOTTlieb STEEN & HAMILTON LLP

One Liberty Plaza
New York, NY 10006-1470

T: +1 212 225 2000

F: +1 212 225 3999

clearygottlieb.com

WASHINGTON, D.C. • PARIS • BRUSSELS • LONDON • MOSCOW
FRANKFURT • COLOGNE • ROME • MILAN • HONG KONG
BEIJING • BUENOSAIRES • SÃO PAULO • ABU DHABI • SEOUL

D: +1 (212) 225 2556
pmarcogliese@cgsh.com

VICTOR J. LEWKOW
THOMAS J. MOLONEY
DAVID G. SABEL
JONATHAN I. BLACKMAN
YARON Z. REICH
RICHARD S. LINGER
JAMES A. DUNCAN
STEVEN M. LOEB
CRAIG B. BROD
NICOLAS GRABAR
CHRISTOPHER E. AUSTIN
HOWARD S. ZELBO
DAVID E. BRODSKY
ARTHUR H. KOHN
RICHARD J. COOPER
JEFFREY S. LEWIS
PAUL J. SHIM
STEVEN L. WILNER
ERIKA W. NIJENHUIS
ANDRES DE LA CRUZ
DAVID C. LOPEZ
MICHAEL A. GERSTENZANG
LEWIS J. LIMAN
LEV L. DASSIN
NEIL Q. WHORISKEY
JORGE U. JUANTORENA
MICHAEL D. WEINBERGER
DAVID LEINWAND
DIANA L. WOLLMAN
JEFFREY A. ROSENTHAL
ETHAN A. KLINGSBERG
MICHAEL D. DAYAN
CARMINE D. BOCCUZZI, JR.
JEFFREY D. KARPFF
KIMBERLY BROWN BLACKLOW
ROBERT J. RAYMOND
SUNG K. KANG
LEONARD C. JACOBY
SANDRA L. FLOW

FRANCISCO L. CESTERO
FRANCESCA L. ODELL
WILLIAM L. MCRAE
JASON FACTOR
JOON H. KIM
MARGARET S. PEPONIS
LISA M. SCHWEITZER
JUAN G. GIRALDEZ
DUANE MCLAUGHLIN
BREON S. PEACE
MEREDITH E. KOTLER
CHANTALE E. KORDULA
BENET J. O'REILLY
ADAM E. FLEISHER
SEAN A. O'NEAL
GLENN P. MCGRORY
MATTHEW P. SALERNO
MICHAEL P. ALBANO
VICTOR L. HOU
ROGER A. COOPER
AMY R. SHAPIRO
JENNIFER KENNEDY PARK
ELIZABETH LENAS
LUKE A. BAREFOOT
PAMELA L. MARCOGLIESE
PAUL M. TIGER
JONATHAN S. KOLODNER
DANIEL ILAN
MEYER H. FEDIDA
ADRIAN R. LEIPSIC
ELIZABETH VICENS
ADAM J. BRENNEMAN
ARI D. MACKINNON
JAMES E. LANGSTON
JARED GERBER
COLIN D. LLOYD
COREY M. GOODMAN
RISHI ZUTSHI
JANE VANLARE

DAVID H. HERRINGTON
KIMBERLY R. SPOERRI
AARON J. MEYERS
DANIEL C. REYNOLDS
ABENA A. MAINOO
HUGH C. CONROY, JR.
JOSEPH LANZKRON
MAURICE R. GINDI
KATHERINE R. REAVES
RAHUL MUKHI
RESIDENT PARTNERS
SANDRA M. ROCKS
S. DOUGLAS BORISKY
JUDITH KASSEL
DAVID E. WEBB
PENELOPE L. CHRISTOPHORO
BOAZ S. MORAG
MARY E. ALCOCK
HEIDE H. ILGENFRITZ
KATHLEEN M. EMBERGER
WALLACE L. LARSON, JR.
AVRAM E. LUFT
ANDREW WEAVER
HELENA K. GRANNIS
JOHN V. HARRISON
CAROLINE F. HAYDAY
NEIL R. MARKEL
KENNETH S. BLAZEJEWSKI
ANDREA M. BASHAM
LAURA BAGARELLA
SHIRLEY M. LO
JONATHAN D.W. GIFFORD
SUSANNA E. PARKER
RESIDENT COUNSEL
LOUISE M. PARENT
OF COUNSEL

MARCH 25, 2019

VIA E-MAIL (shareholderproposals@sec.gov)

U.S. Securities and Exchange Commission
Division of Corporation Finance
Office of Chief Counsel
100 F Street, N.E.
Washington, DC 20549

Re: Stockholder Proposal Submitted by Lisa Stephanie Myrkalo, Andrea Louise Dixon, and the New York State Common Retirement Fund

Ladies and Gentlemen:

We refer to our letter dated February 5, 2019 (the “No Action Request”), submitted on behalf of our client, Alphabet Inc., a Delaware corporation (“Alphabet” or the “Company”), pursuant to which we requested that the staff of the Division of Corporation Finance (the “Staff”) of the Securities and Exchange Commission (the “Commission”) concur with the Company’s view that the shareholder proposal and accompanying supporting statement submitted by the New York State Common Retirement Fund (the “Proponent”) and Lisa Stephanie Myrkalo and Andrea Louise Dixon (the “Co-Proponents” and, together with the Proponent, the “Proponents”), may be properly omitted from the proxy materials to be distributed by Alphabet in connection with its 2019 annual meeting of shareholders. Alphabet received a copy of the letter to the Staff dated March 11, 2019, submitted by the Proponent in response to the No Action Request (the “Response Letter”).

This letter is in response to the Response Letter and supplements the No Action Request. In accordance with Rule 14a-8(j), we are simultaneously sending a copy of this letter and its attachments to the Proponents.

BACKGROUND

Pursuant to the instructions in the Company's 2018 Proxy Statement (See Exhibit A), the Proponent submitted a shareholder proposal to the Company on December 6, 2018 (the "Initial Proposal") via email to the address corporatesecretary@abc.xyz, and by mail. However, when the Proponent submitted an amended proposal by mail and via email on December 20, 2018 (the "Amended Proposal"), the Proponent made a typographical error in the email address of the Corporate Secretary, and sent the Amended Proposal to the erroneous address corporatesecretary@abc.xy instead of the correct address corporatesecretary@abc.xyz (See Exhibit B). The Co-Proponents later submitted an identical proposal to the Amended Proposal on December 26, 2018 (the "Co-Proponents' Proposal"). Due to this typographical error, the Corporate Secretary never received an electronic copy of the Amended Proposal and therefore did not receive the Amended Proposal in the same manner as the Initial Proposal, as described in the Company's 2018 Proxy Statement. We would like to draw your attention to the fact that any emails sent to the erroneous address corporatesecretary@abc.xy would have resulted in the sender receiving a "mailer daemon" notice that "Delivery has failed to these recipients or groups," thereby putting the Proponent on notice of the fact that the email address they used was incorrect (See Exhibit C). In spite of this, the Proponent took no action to correct their deficient email delivery. In the resultant confusion caused by the Proponent's uncorrected error, the Company erroneously believed that the Co-Proponents' proposal was being co-filed with the Initial Proposal and not the Amended Proposal, which resulted in the No Action Request being written in response to the Initial Proposal, rather than the Amended Proposal.

Nevertheless, we emphasize that the Company has still substantially implemented the Amended Proposal under Rule 14a-8(i)(10), as had been argued in the original No Action Request. In addition, in light of the new arguments that the Proponent has raised in the Response Letter, we reiterate, as we had already expressed in the original No Action Request, that the Amended Proposal may be omitted under 14a-8(i)(3) because it is vague and indefinite.

ANALYSIS

I. Under Rule 14a-8(i)(3), the Amended Proposal may be omitted because it is vague and indefinite.

Rule 14a-8(i)(3) provides that if a proposal is vague or indefinite, it may be omitted. The Staff has interpreted Rule 14a-8(i)(3) to mean that vague and indefinite shareholder proposals may be excluded because "neither the stockholders voting on the proposal, nor the company in implementing the proposal (if adopted), would be able to determine with any reasonable certainty exactly what actions or measures the proposal requires." SEC Staff Legal Bulletin No. 14B (Sept. 15, 2004). A proposal is sufficiently vague and indefinite to justify exclusion where a company and its shareholders might interpret the proposal differently, such that "any action ultimately taken by the company upon implementation of the proposal could be significantly different from the actions envisioned by the shareholders voting on the proposal." Fuqua Industries, Inc. (avail. Mar. 12, 1991).

In the No Action Request, we had noted that in order for the Initial Proposal not to be impermissibly vague and indefinite under rule 14a-8(i)(3), it needed to be interpreted as

being focused solely on content management in relation to election interference. See No Action Request, at 4. The Amended Proposal creates exactly this issue of impermissible vagueness and indefiniteness, by adding the notoriously difficult to define concepts of “freedom of expression” and “hate speech,” which significantly complicates the task of determining what the Proponent is requesting the Company do or how shareholders would know that the Company has implemented the Amended Proposal.

Furthermore, the text of the Amended Proposal requests that the Company issue a report “reviewing the efficacy of its enforcement of Google’s terms of service related to content policies and assessing the risks posed by content management controversies related to election interference, freedom of expression, and the spread of hate speech, to the company’s finances, operations, and reputation.” “Content policies” and “content management” are extremely broad terms that can potentially encompass the management of content on the Company’s platforms for any reason. There can be an infinite number of different reasons for managing content, including, but not limited to, copyright violation, graphic violence, profanity, defamation, digital scams, privacy, and obscenity. Any one of these issues could be the subject matter of a report all by itself, and none of the terms are susceptible to any one clear interpretation. Scores of books and generations of thinkers have struggled with how to define these terms—as Supreme Court Justice Potter Stewart famously stated in reference to obscenity “I know it when I see it.” Jacobellis v. Ohio, 378 U.S. 184, 197 (1964). Therefore, unless the scope of the Amended Proposal is properly limited, it would be impossible for the Company to determine with any reasonable certainty exactly what the proposed report would require.

Further confusing matters, the focus of the whereas clause of the Amended Proposal is on a series of specific content management events that have already been disclosed and publicly addressed by the Company. This is distinctly at odds with the resolved clause’s request for a report on the enforcement of Google’s terms of service related to content policies.

By way of example, the Proponents themselves do not appear to understand what is being requested by the Amended Proposal. The Response Letter makes clear that they envision the implementation of the Amended Proposal as somehow resulting in a solution to a problem that the Company already dedicates significant resources towards precisely because there is no easy solution. Indeed, the Response Letter states: “The essential purpose of the Proposal is to address the Company’s failure to effectively address content governance concerns, posing a risk to shareholder value.” Response Letter, at 4. However, the concept of effectively addressing content governance concerns does not appear anywhere within the text of the Amended Proposal, nor is there any indication as to what “effective” means in this context. The vague and indefinite nature of the proposal is further demonstrated by the Response Letter’s singular focus on whether content governance concerns exist, rather than the question that the resolved clause of the Amended Proposal actually asks, which is whether the Company’s actions in relation to such issues have been adequately disclosed to shareholders.

As discussed above, it appears that the Proponents themselves cannot articulate what would be required under the Amended Proposal. Therefore, it is clear that if the Amended Proposal were included in the Company’s 2019 Proxy Statement, neither the Company nor its stockholders would be able to determine with any reasonable certainty what actions or measures the proposal requires. For the reasons set forth above, the Company respectfully requests the

Staff's concurrence in the omission of the Amended Proposal under Rule 14a-8(i)(3) for being so vague and indefinite as to be misleading.

II. Under Rule 14a-8(i)(10), the Amended Proposal may be omitted because it has been substantially implemented by the Company.

A company is determined to have substantially implemented a proposal under Rule 14a-8(i)(10) if the company's policies, practices and procedures "compare favorably with the guidelines of the proposal," or the company's actions have satisfactorily addressed the underlying concerns and "essential objectives" of the proposal. See The Talbots, Inc. (avail. Apr. 5, 2002). Differences between a company's actions and a proposal are permitted so long as the company's actions satisfactorily address the proposal's underlying concerns and essential objectives. See Release No. 34-20091. (Aug. 16, 1983).

The Response Letter takes the position that (i) having policies and practices that compare favorably with the guidelines of the proposal and (ii) having satisfactorily addressed the essential objectives of the proposal are two separate requirements that must *both* be satisfied to meet the standard of Rule 14a-8(i)(10). However, the Staff has never viewed these as two separate requirements, but rather different ways of stating the same standard. See The Talbots, Inc. (avail. Apr. 5, 2002). (concurring with the view that "shareholder proposals have been substantially implemented within the scope of Rule 14a-8(i)(10) when the company already has policies and procedures in place relating to the subject matter of the proposal, or has implemented the essential objectives of the proposal." (emphasis added)). Therefore, contrary to the standard stated in the Response Letter, it is enough to show that either (i) the Company's practices and procedures compare favorably with the guidelines of the proposal or (ii) the Company's actions have satisfactorily addressed the underlying concerns and essential objectives of the proposal. Nevertheless, in the current instance, the Company has substantially implemented the Amended Proposal no matter which standard is used.

As discussed in Section I, under the interpretation most legally favorable to the Proponent, the scope of the Amended Proposal must be interpreted to be limited to Alphabet's content management on its social media platforms in relation to election interference, freedom of expression and the spread of hate speech, as any other necessarily broader interpretation would render the Amended Proposal vague and indefinite under Rule 14a-8(i)(3). Although the Company has also been open in its communications regarding other content management issues as well, such as copyright violations and privacy, a discussion about all such content issues would be reaching far beyond a legally supportable interpretation of the scope of the Amended Proposal. In addition, from a content management perspective, YouTube is now the Company's sole major source of user content, as the Company has announced that its other social media platform Google+ is being shut down as of April 2, 2019. Therefore, this section will focus on the Amended Proposal's essential objective, which would seem to be that the Company adequately disclose its actions with respect to content governance regarding election interference, freedom of expression and the spread of hate speech, particularly on YouTube.

The Company shares the Proponents' concerns regarding these issues, and appreciates that in the Response Letter the Proponents appear to concur that the Amended

Proposal has been substantially implemented with respect to election interference issues. However, the Proponents appear to be unaware of the fact that since April of 2018, the Company has been issuing a quarterly report regarding the enforcement of YouTube's Community Guidelines (the "YouTube Enforcement Report") (See Exhibit D), as part of its overarching transparency report on how laws and policies affect the Company's business (the "Transparency Report"). The YouTube Enforcement Report provides substantially similar levels of in-depth disclosure as those found in the election interference transparency reports, which the Proponents appear to agree in the Response Letter compare favorably with the Amended Proposal.

As previously discussed in the No Action Letter, Google, the Company's principal subsidiary, was founded with the mission to "organize the world's information and make it universally accessible and useful." The Company strongly believes that the abuse of its platforms to spread harmful content is antithetical to that mission, and that it has a responsibility to prevent such abuses. The Company works hard, and dedicates significant resources, to maintaining a safe and vibrant online community, and has been open about what type of behavior it will not allow on its platforms. For example, YouTube has Community Guidelines that provide an in-depth look into what content it does not tolerate on its platform (See Exhibit E).

The Amended Proposal requests that the Company issue a report "reviewing the efficacy of its enforcement of Google's terms of service related to content policies." However, the Company already publicly releases an extensive quarterly Transparency Report, which provides a detailed and lengthy disclosure of its content policy enforcement activities. In the YouTube Enforcement Report section, the Company provides the exact number of channels, videos, and comments that have been removed for violating the Community Guidelines each quarter, which are further broken down by their removal reasons and sources of detection (See Exhibit D). The YouTube Enforcement Report also provides details on the videos that have been flagged by users as possible violations, and on the process used by the Company for dealing with these flags. Other sections of the Transparency Report provide similar details on how the Company has dealt with content removals related to copyright, government requests, and privacy laws, in addition to details on security and privacy, political advertising, and traffic disruptions. The Company's quarterly publication of the Transparency Report alone would seem to be a practice that compares favorably with the guidelines of the Amended Proposal, as it provides a detailed review of the enforcement of the Company's content policies, which is what was essentially requested by the Amended Proposal.

In addition to the Transparency Report, the Company has also released a series of blog posts written by senior executives that provide details about how the Company is committed to tackling these issues. For example, Kent Walker, Google's Senior Vice President of Global Affairs & Chief Legal Officer, has written about the Company's efforts to combat harmful content on the Company's official blog on both January 8, 2019, and February 14, 2019 (collectively, the "Blog Posts") (See Exhibit F). The Blog Posts not only provide details about how the Company is combating harmful content throughout its platforms, but they also outline the challenges that the Company faces in the process. The Blog Posts explain how the Company must also be conscious of the importance of protecting legal speech, and how it can be difficult for even highly trained reviewers to tell the difference between legal and illegal content. They also explain in detail how the Company tries to strike a balance between over-regulation and

under-regulation, but stress the importance of collaboration among governments, industry, and civil society to effectively tackle the matter. As another example of the Company's willingness to engage with the public on these matters, YouTube has also released its own blog post about enforcing the YouTube Community Guidelines on December 13, 2018 (See Exhibit G). These communication efforts thus do address the underlying concerns of the Amended Proposal about the Company's enforcement policies on content governance.

The Amended Proposal further requests that the report assess "the risks posed by content management controversies related to election interference, freedom of expression, and the spread of hate speech, to the company's finances, operations, and reputation," but the Company has already made significant disclosures of the risks it faces. In addition to the Blog Posts which outline the challenges that the Company faces, in the Risk Factors section of its Annual Report for the fiscal year ended December 31, 2018, the Company emphasized that its business depends on the strength of its brand, highlighting that "if we fail to appropriately respond to the sharing of objectionable content on our services or objectionable practices by advertisers...our users may lose confidence in our brands. Our brands may also be negatively affected by the use of our products or services to disseminate information that is deemed to be false or misleading." The Company further disclosed that, "if we fail to maintain and enhance equity in our brands, our business, operating results, and financial condition may be materially and adversely affected." As a separate risk factor on advertising, the Company also noted that "Many of our advertisers, companies that distribute our products and services, digital publishers, and content partners can terminate their contracts with us at any time," and that "[i]f we do not provide superior value or deliver advertisements efficiently and competitively, our reputation could be affected, we could see a decrease in revenue from advertisers and/or experience other adverse effects to our business."

A large portion of the whereas clause and Response Letter consists of references to instances in which bad actors have violated the Company's Community Guidelines or terms of service. However, the resolved clause indicates that the Amended Proposal is only requesting transparency on these issues and the steps that the company takes with respect to them. These examples therefore actually serve to demonstrate how the Company has already substantially implemented the Amended Proposal. While the Company cannot prevent every instance of misuse, it has already publicly responded to most of these allegations, and has implemented mechanisms for dealing with such improper content. As already shown in the preceding paragraphs, the Company is actively working to improve upon its content management activities and has been openly engaging the public about its efforts. For example, the Response Letter cites numerous examples of conspiracy theories that were uploaded onto YouTube. On January 25, 2019, YouTube wrote an official blog post that it will begin reducing recommendations of content that could "misinform users in harmful ways," such as "videos promoting a phony miracle cure for a serious illness, claiming the earth is flat, or making blatantly false claims about historic events like 9/11." (See Exhibit H). The blog post emphasized that this change "strikes a balance between maintaining a platform for free speech and living up to [its] responsibility to users." The Company also released a blog post on February 16, 2019 that highlighted its commitment to fighting disinformation (See Exhibit I). In the February 16th blog post, the Company also shared a white paper that it presented at the Munich Security Conference that gives more detail about its work to tackle the intentional spread of misinformation—across

Google Search, Google News, YouTube and its advertising systems (See Exhibit J). The Response Letter also noted that inappropriate messages and links were appearing in the comments of videos of children. However, as the Response Letter noted, YouTube swiftly responded on February 28, 2019 by announcing that it would disable comments on videos featuring minors, launch a new comments classifier, and take action against creators who cause “egregious harm” to the community (See Exhibit K). Although the Response Letter goes on to criticize this action by calling it a “sudden and extreme response,” the Proponents again appear to misunderstand their own proposal, as the very act they criticize is the type of transparent engagement with the larger community that is sought by the Amended Proposal.

The Company shares the Proponents’ concerns and is dedicated to combating harmful content in an effective and fair way. As shown throughout this letter, the Company has been consistently open about its efforts, and its Transparency Report provides a detailed breakdown of the effectiveness of its content management enforcement policies. It is therefore unclear what the Amended Proposal and the Response Letter are asking for, beyond what the Company has already done and is continuing to do. Through its actions, the Company has already addressed the Proposal’s underlying concern and essential objective and rendered the purpose of the Proposal moot. In order to “avoid the possibility of shareholders having to consider matters which already have been favorably acted upon by . . . management,” SEC Release No. 34-12598 (July 7, 1976), the Company respectfully requests the Staff’s concurrence in the omission of the Proposal as having been substantially implemented pursuant to Rule 14a 8(i)(10).

* * * * *

Conclusion

By copy of this letter, the Proponents are being notified that for the reasons set forth herein and in the No Action Request, the Company intends to omit the Proposal from its Proxy Statement. We respectfully request that the Staff confirm that it will not recommend any enforcement action if the Company omits the Proposal from its Proxy Statement. If we can be of assistance in this matter, please do not hesitate to call me.

Sincerely,


Pamela L. Marcogliese

Enclosures

cc: Patrick Doherty, Office of the New York State Comptroller, on behalf of the New York State Common Retirement Fund; and

Natasha Lamb, Managing Partner, Arjuna Capital, on behalf of Lisa Stephanie Myrkalo and Andrea Louise Dixon

EXHIBIT A

25. Who will serve as inspector of elections?

The inspector of elections will be a representative from Computershare.

26. How can I contact Alphabet's transfer agent?

Contact our transfer agent by either writing to Computershare Investor Services, P.O. Box 505000, Louisville, KY 40233-5000 (courier services should be sent to Computershare Investor Services, 462 South 4th Street, Suite 1600, Louisville, KY 40202-3467), by telephoning shareholder services 1-866-298-8535 (toll free within the USA, US territories and Canada) or 1-781-575-2879 or by visiting Investor Centre™ portal at www.computershare.com/investor.

Stockholder Proposals, Director Nominations, and Related Bylaw Provisions

27. What is the deadline to propose actions for consideration at next year's Annual Meeting of Stockholders or to nominate individuals to serve as directors?

Stockholder Proposals: Stockholders may present proper proposals for inclusion in our proxy statement and for consideration at the 2019 Annual Meeting of Stockholders by submitting their proposals in writing to Alphabet's Corporate Secretary in a timely manner. For a stockholder proposal to be considered for inclusion in our proxy statement for our 2019 Annual Meeting of Stockholders, the Corporate Secretary of Alphabet must receive the written proposal at our principal executive offices or at the email address set forth below no later than December 28, 2018. If we hold our 2019 Annual Meeting of Stockholders more than 30 days before or after June 6, 2019 (the one-year anniversary date of the 2018 Annual Meeting of Stockholders), we will disclose the new deadline by which stockholder proposals must be received under Item 5 of Part II of our earliest possible Quarterly Report on Form 10-Q or, if impracticable, by any means reasonably determined to inform stockholders. In addition, stockholder proposals must otherwise comply with the requirements of Rule 14a-8 under the Exchange Act and with the SEC regulations under Rule 14a-8 regarding the inclusion of stockholder proposals in company-sponsored proxy materials. Proposals should be addressed in one of the following two ways:

1. via mail with a copy via email:

 Alphabet Inc.
Attn: Corporate Secretary
1600 Amphitheatre Parkway
Mountain View, California
94043



With a copy via email:
corporatesecretary@abc.xyz

OR



corporatesecretary@abc.xyz

Our bylaws also establish an advance notice procedure for stockholders who wish to present a proposal before an annual meeting of stockholders but do not intend for the proposal to be included in our proxy statement. Our bylaws provide that the only business that may be conducted at an annual meeting is business that is: (1) specified in the notice of a meeting given by or at the direction of our Board of Directors, (2) otherwise properly brought before the meeting by or at the direction of our Board of Directors, or (3) properly brought before the meeting by a stockholder entitled to vote at the annual meeting who has delivered timely written notice to our Corporate Secretary, which notice must contain the information specified in our bylaws. To be timely for our 2019 Annual Meeting of Stockholders, our Corporate Secretary must receive the written notice at our principal executive offices or at the email address set forth above:

- not earlier than the close of business on February 6, 2019, and
- not later than the close of business on March 8, 2019.

If we hold our 2019 Annual Meeting of Stockholders more than 30 days before or after June 6, 2019 (the one-year anniversary date of the 2018 Annual Meeting of Stockholders), the notice of a stockholder proposal that is not intended to be included in our proxy statement must be received not later than the close of business on the earlier of the following two dates:

- the 10th day following the day on which notice of the meeting date is mailed, or
- the 10th day following the day on which public disclosure of the meeting date is made.

EXHIBIT B



Valentina Margulis <valya@google.com>

Shareholder Request

Thu, Dec 20, 2018 at 2:08 PM

Tana Goldsmith <tgoldsmith@osc.ny.gov>
To: "corporatesecretary@abc.xy" <corporatesecretary@abc.xy>, "valya@google.com" <valya@google.com>
Cc: Patrick Doherty <pdoherty@osc.ny.gov>

Mr. David C. Drummond,

Please find attached a copy of the New York State Common Retirement Fund filing letter and shareholder resolution, which was also sent to you today via UPS.

If you have any questions, please feel free to contact me regarding this transmission.

Kind Regards,

Tana

Tana Goldsmith

Special Investment Officer
Pension Investment and Cash Management

Office of the State Comptroller

59 Maiden Lane Fl. 30

New York, NY 10038

tgoldsmith@osc.ny.gov

Direct Line: 212.383.2592

Receptionist: 212.383.3931

Facsimile: 212.383.1331

 **Alphabet, Inc. Shareholder Proposal 12.20.18.pdf**
154K

EXHIBIT C

From: Mail Delivery Subsystem <MAILER-DAEMON@mx0b-00170f01.pphosted.com>
To: corporatesecretary@abc.xy
Sent: Wednesday, March 13, 2019 11:25 AM
Subject: Undeliverable: Test

Delivery has failed to these recipients or groups:

corporatesecretary@abc.xy

The recipient's e-mail address isn't correct. Please check the e-mail address and try to resend the message. If the problem continues, contact your helpdesk.

The following organization rejected your message: abc.xy.

Diagnostic information for administrators:

Generating server: mx0b-00170f01.pphosted.com

corporatesecretary@abc.xy
abc.xy #<abc.xy #5.1.2 SMTP; 550 Host unknown> #SMTP#

Original message headers:

Return-Path: <prvs=5975f22b28=ychang@cgsh.com>
Received: from pps.filterd (m0045993.ppops.net [127.0.0.1]) by
mx0b-00170f01.pphosted.com (8.16.0.27/8.16.0.27) with SMTP id x2DF6uPT027960
for <corporatesecretary@abc.xy>; Wed, 13 Mar 2019 11:23:54 -0400
Received: from amlexchub01.cgsh.com (ammail11.cgsh.com [144.121.68.203]) by
mx0b-00170f01.pphosted.com with ESMTP id 2r6fgahq2j-1 (version=TLSv1
cipher=ECDHE-RSA-AES256-SHA bits=256 verify=NOT) for
<corporatesecretary@abc.xy>; Wed, 13 Mar 2019 11:23:54 -0400
Received: from AM1EXCAPP01.cgsh.com (10.200.40.84) by AM1EXCHUB01.cgsh.com
(10.200.40.22) with Microsoft SMTP Server (TLS) id 14.3.408.0; Wed, 13 Mar
2019 11:23:53 -0400
Received: from AM1EXCMBX04.cgsh.com ([fe80::e9de:4b2:ba3:dfa4]) by
AM1EXCAPP01.cgsh.com ([fe80::84a7:9f31:2e53:a30a%11]) with mapi id
14.03.0415.000; Wed, 13 Mar 2019 11:23:52 -0400

EXHIBIT D

YouTube Community Guidelines enforcement

At YouTube, we work hard to maintain a safe and vibrant community. We have Community Guidelines that set the rules of the road for what we don't allow on YouTube. For example, we do not allow pornography, incitement to violence, harassment, or hate speech. We rely on a combination of people and technology to flag inappropriate content and enforce these guidelines. Flags can come from our automated flagging systems, from members of the Trusted Flagger program (NGOs, government agencies, and individuals) or from users in the broader YouTube community. This report provides data on the flags YouTube receives and how we enforce our policies.

Removed channels by the numbers

OCT 2018 – DEC 2018 ▾

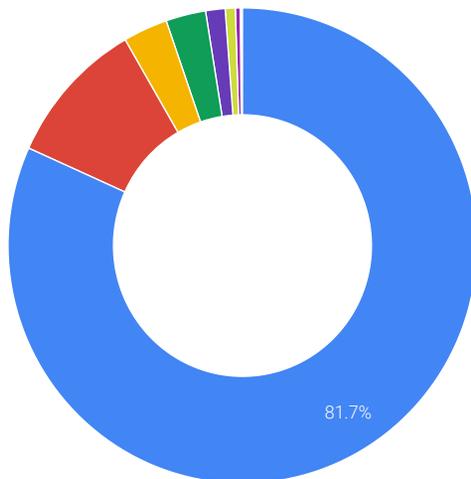
Total channels removed

2,398,961

A YouTube channel is terminated if it accrues three [Community Guidelines strikes](#) in 90 days, has a single case of severe abuse (such as predatory behavior), or is determined to be wholly dedicated to violating our guidelines (as is often the case with spam accounts). When a channel is terminated, all of its videos are removed. In Q4 2018, 76.9 million videos were removed due to a channel-level suspension.

This exhibit shows the number of channels removed by YouTube for violating its Community Guidelines per quarter.

Channels removed, by removal reason



This chart shows the volume of channels removed by YouTube, by the reason a channel was removed. The majority of channel terminations are a result of accounts being dedicated to spam or adult sexual content in violation of our guidelines.

When we terminate a channel for receiving three [Community Guidelines strikes](#) for violating several different policies within a three month period, we categorize it under a separate label - "Multiple policy violations" - because these accounts were not wholly dedicated to one policy violation.

Removed videos by the numbers

OCT 2018 – DEC 2018 ▾ INCLUDE AUTOMATED FLAGGING ▾

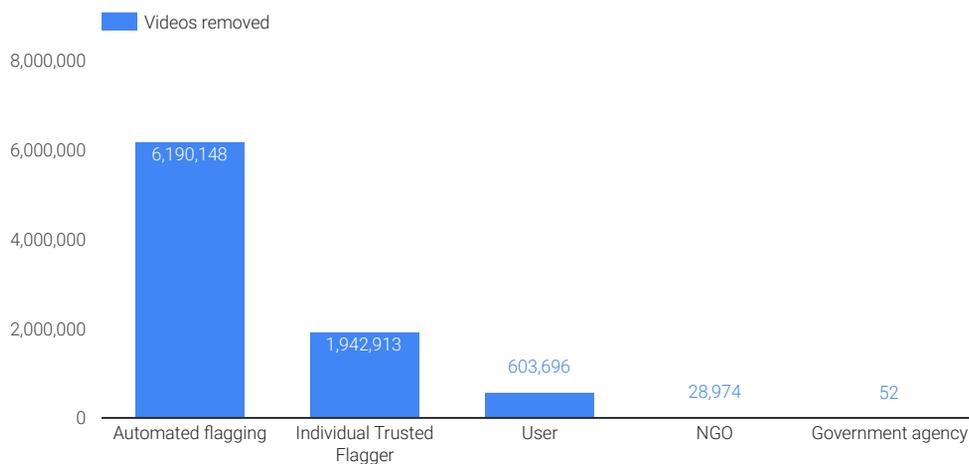
Total videos removed

8,765,783

YouTube relies on teams around the world to review flagged videos and remove content that violates our terms of service; restrict videos (e.g., age-restrict content that may not be appropriate for all audiences); or leave the content live when it doesn't violate our guidelines.

This exhibit shows the number of videos removed by YouTube for violating its Community Guidelines per quarter.

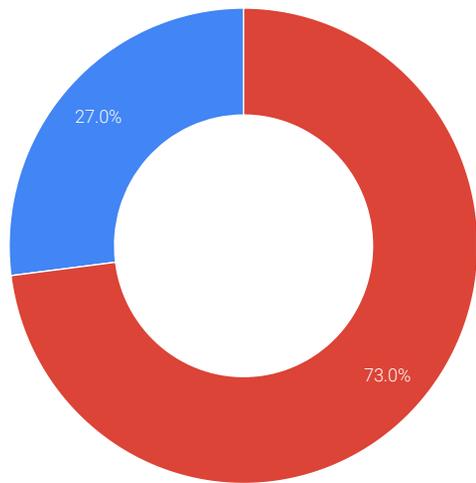
Videos removed, by source of first detection



OCT 2018 – DEC 2018 ▾ INCLUDE AUTOMATED FLAGGING ▾

This chart shows the volume of videos removed by YouTube, by source of first detection (automated flagging or human detection). Flags from human detection can come from a user or a member of YouTube's [Trusted Flagger](#) program. Trusted Flagger program members include individuals, NGOs, and government agencies that are particularly effective at notifying YouTube of content that violates our Community Guidelines.

Removed videos first flagged through automated flagging, with and without views

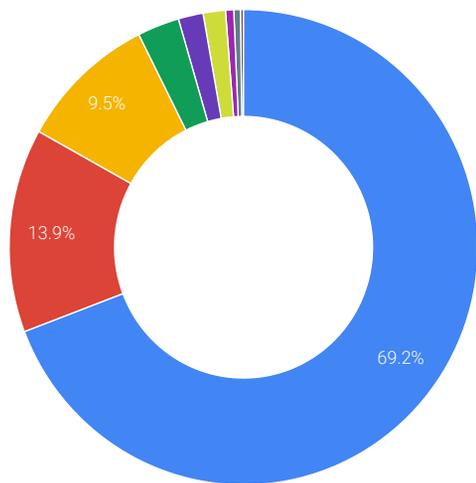


● Removed before any views
 ● Removed after any views

OCT 2018 – DEC 2018 ▾

Automated flagging enables us to act more quickly and accurately to enforce our policies. This chart shows the percentage of video removals (first flagged through our automated flagging systems) that occurred before they received any views versus those that occurred after receiving some views.

Videos removed, by removal reason



● Spam, mislea...
 ● Nudity or sex...
 ● Child safety
 ● Violent or gra...
 ◀ 1/3 ▶

OCT 2018 – DEC 2018 ▾

This chart shows the volume of videos removed by YouTube, by the reason a video was removed. These removal reasons correspond to YouTube's [Community Guidelines](#). Reviewers evaluate flagged videos against all of our Community Guidelines and policies, regardless of the reason the video was originally flagged.

Removed comments by the numbers

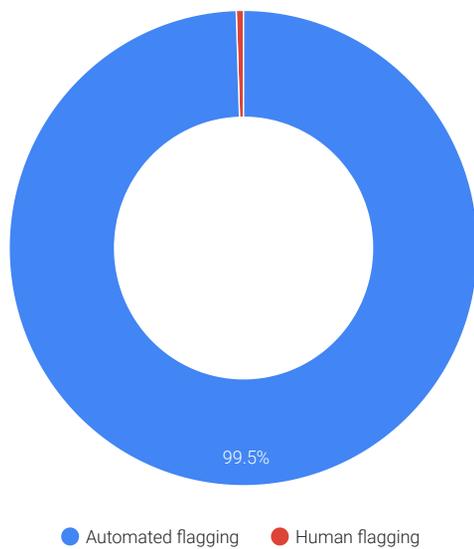
Total comments removed

261,645,574

YouTube is a vibrant community in which millions of people post billions of comments each quarter. Using a combination of people and technology, we remove comments that violate our [Community Guidelines](#). We also filter comments which we have high confidence are spam into a 'Likely spam' folder that creators can review and approve if they choose.

This exhibit shows the volume of comments removed by YouTube for violating our [Community Guidelines](#) and filtered as likely spam which creators did not approve.

Comments removed, by source of first detection



Most removed comments are detected by our automated flagging systems but they can also be flagged by human flaggers. We rely on teams around the world to review flagged comments and remove content that violates our Terms of Service, or leave the content live when it doesn't violate our guidelines.

This chart shows the volume of comments removed by YouTube for violating our [Community Guidelines](#), by source of first detection (automated flagging or human detection). The majority of actions we take on comments is for violating our guidelines against spam.

YouTube Community Guidelines enforcement

YouTube is a community and, over the years, people have used the flagging feature located beneath every video and comment to help report content they believe violates our Community Guidelines. We want to empower the YouTube community to understand how flagging works and to get involved in maintaining our Community Guidelines.

Flags

At YouTube, we work hard to maintain a safe and vibrant community. We have Community Guidelines that set the rules of the road for what we don't allow on YouTube. This section of the report provides data on the flags YouTube receives for possible violations of our Community Guidelines.

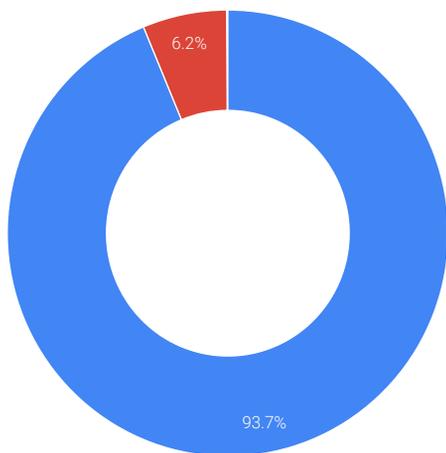
Human flags by flagger type

OCT 2018 – DEC 2018 ▾

Videos flagged by all human flaggers

11,939,958

In addition to our automated flagging systems, [Trusted Flaggers](#) and our broader community of users play an important role in flagging content. This chart shows the breakdown of flags that come from different types of human flaggers. The number above the chart shows the number of unique videos that were flagged. A single video may be flagged multiple times and for different reasons. Flagged content will remain live when it doesn't violate our Community Guidelines.



● User ● Individual Trusted Flagger ● NGO ● Government agency

OCT 2018 – DEC 2018 ▾

INCLUDE USERS ▾

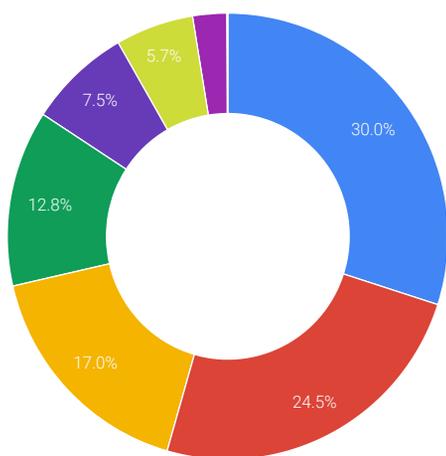


We receive flags for suspected violations of our Community Guidelines from users and Trusted Flaggers all around the world. The chart below shows the countries from which we received the most human flags, ranked by total volume. Flagged content will remain live when it doesn't violate our Community Guidelines.

OCT 2018 – DEC 2018 ▾

Rank	Country
1	India
2	United States
3	Brazil
4	United Kingdom
5	Indonesia
6	Russia
7	South Korea
8	Mexico
9	Turkey
10	Thailand

Human flags by flagging reason



● Spam or misle... ● Sexual ● Hateful or abus... ● Violent or reput... ◀ 1/2 ▶

OCT 2018 – DEC 2018 ▾

When flagging a video, human flaggers can select a reason they are reporting the video and leave comments or video timestamps for YouTube's reviewers. This chart shows the flagging reasons that people selected when reporting YouTube content. A single video may be flagged multiple times and may be flagged for



Community Guidelines enforcement

The YouTube community plays an important role in flagging videos that violate our Community Guidelines. Any logged-in user can flag a video by clicking on the three dots to the bottom right of the video player and selecting "Report." Trained teams evaluate videos before taking action in order to ensure it actually violates our policies and to protect content that has an educational, documentary, scientific, or artistic purpose. The teams carefully evaluate flags 24 hours a day, seven days a week. They remove content that violates our terms, age-restrict content that may not be appropriate for all audiences, and leave content live when it doesn't violate our guidelines.

[LEARN MORE](#) ↗

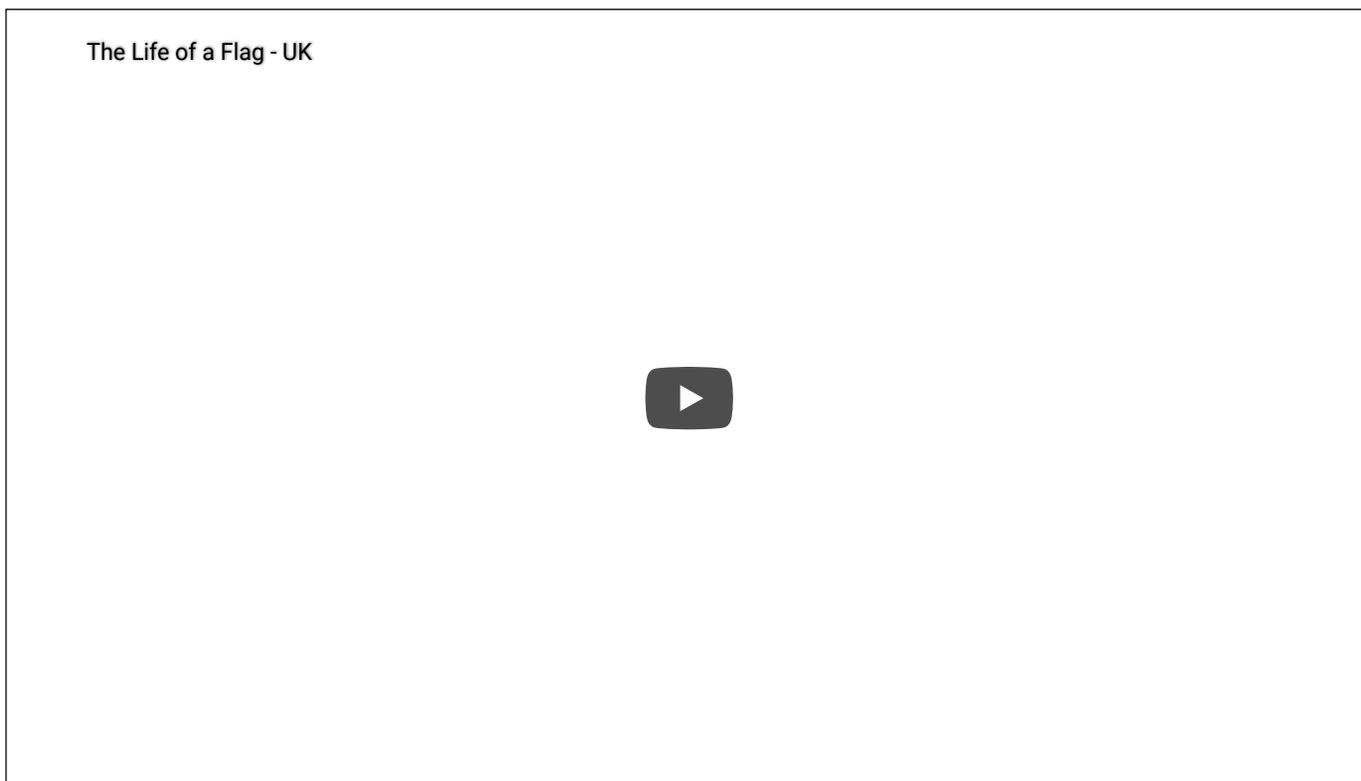
Trusted Flagger program

The Trusted Flagger program was developed to enable highly effective flaggers to alert us to content that violates our Community Guidelines via a bulk reporting tool. Individuals with high flagging accuracy rates, NGOs, and government agencies participate in this program, which provides training in enforcing YouTube's Community Guidelines. Because participants' flags have a higher action rate than the average user, we prioritize them for review. Videos flagged by Trusted Flaggers are subject to the same policies as videos flagged by any other user and are reviewed by our teams who are trained to make decisions on whether content violates our Community Guidelines and should be removed.

[LEARN MORE](#) ↗

The life of a flag

As you've learned in this report, "flags" mark content that may violate our Community Guidelines. This video explains how YouTube receives flags, actions reviewers take on flags, and other processes and policies that help us keep the YouTube community safe.



ALL TIME ▾

ALL REASONS ▾

ALL FLAGGER TYPES ▾

Flagging reason: Hateful or abusive

Flagger type: User

Video description: A publically-broadcast video of a musical performance by a well-known Korean pop group.

Outcome: Content did not violate policy. No action taken.

Flagging reason: Violent or repulsive

Flagger type: User

Video description: A video on an amateur science channel featuring a live grasshopper being microwaved.

Outcome: Despite amateur science intent, video violated our animal abuse policy related to violent or repulsive content. Content was removed.

Flagging reason: Hateful or abusive

Flagger type: User

Video description: A German news station's video of a Member of Parliament from the German AfD Party giving a speech about voting age at the Bundestag.

Outcome: Content did not violate policy. No action taken.

Flagging reason: Child abuse

Flagger type: User

Video description: A news clip from a Russian television broadcast with a military official discussing the situation in Dagestan which did not feature children.

Outcome: Content did not violate policy. No action taken.

Flagging reason: Child abuse

Flagger type: User

Video description: A video from a prominent environmental organization about the life of young lions.

Outcome: Content did not violate policy. No action taken.

EXHIBIT E



Community Guidelines

Safety Tools & Resources

Reporting & Enforcement

Policies and Safety

When you use YouTube, you join a community of people from all over the world. Every cool, new community feature on YouTube involves a certain level of trust. Millions of users respect that trust and we trust you to be responsible too. Following the guidelines below helps to keep YouTube fun and enjoyable for everyone.

You might not like everything you see on YouTube. If you think content is inappropriate, use the [flagging feature](#) to submit it for review by our YouTube staff. Our staff carefully reviews flagged content 24 hours a day, 7 days a week to determine whether there's a violation of our Community Guidelines.



Here are some common-sense rules that'll help you steer clear of trouble. Please take these rules seriously and take them to heart. Don't try to look for loopholes or try to lawyer your way around the guidelines—just understand them and try to respect the spirit in which they were created.



Nudity or sexual content

YouTube is not for pornography or sexually explicit content. If this describes your video, even if it's a video of yourself, don't post it on YouTube. Also, be advised that we work closely with law enforcement and we report child exploitation. [Learn more](#)



Harmful or dangerous content

Don't post videos that encourage others to do things that might cause them to get badly hurt, especially kids. Videos showing such harmful or dangerous acts may get age-restricted or removed depending on their severity. [Learn more](#)



Hateful content

Our products are platforms for free expression. But we don't support content that promotes or condones violence against individuals or groups based on race or ethnic origin, religion, disability, gender, age, nationality, veteran status, or sexual orientation/gender identity, or whose primary purpose is inciting hatred on the basis of these core characteristics. This can be a delicate balancing act, but if the primary purpose is to attack a protected group, the content crosses the line. [Learn more](#)



Violent or graphic content

It's not okay to post violent or gory content that's primarily intended to be shocking, sensational, or gratuitous. If posting graphic content in a news or documentary context, please be mindful to provide enough information to help people understand what's going on in the video. Don't encourage others to commit specific acts of violence. [Learn more](#)



Harassment and cyberbullying

It's not ok to post abusive videos and comments on YouTube. If harassment crosses the line into a malicious attack it can be reported and may be removed. In other cases, users may be mildly annoying or petty and should be ignored. [Learn more](#)



Spam, misleading metadata, and scams

Everyone hates spam. Don't create misleading descriptions, tags, titles, or thumbnails in order to increase views. It's not okay to post large amounts of untargeted, unwanted or repetitive content, including comments and private messages. [Learn more](#)



Threats

Things like predatory behavior, stalking, threats, harassment, intimidation, invading privacy, revealing other people's personal information, and inciting others to commit violent acts or to violate the Terms of Use are taken very seriously. Anyone caught doing these things may be permanently banned from YouTube.

[Learn more](#)



Copyright

Respect copyright. Only upload videos that you made or that you're authorized to use. This means don't upload videos you didn't make, or use content in your videos that someone else owns the copyright to, such as music tracks, snippets of copyrighted programs, or videos made by other users, without necessary authorizations. Visit our Copyright Center for more information. [Learn more](#)



Privacy

If someone has posted your personal information or uploaded a video of you without your consent, you can request removal of content based on our Privacy Guidelines.

[Learn more](#)



Impersonation

Accounts that are established to impersonate another channel or individual may be removed under our impersonation policy. [Learn more](#)



Child Safety

Learn about how we protect minors in the YouTube ecosystem. Also, be advised that we work closely with law enforcement and we report child endangerment. [Learn more](#)



Additional policies

Additional policies on a range of subjects. [Learn more](#)

If a YouTube creator's on- and/or off-platform behavior harms our users, community, employees or

YouTube

About

Partners

Advertise

[Blogs](#)

[Jobs](#)

[YouTube Premium](#)

[YouTube Music](#)

Resources

[Policies & Safety](#)

[Copyright](#)

[Brand Guidelines](#)

[Help Center](#)

[Privacy](#)

[English](#)

© 2019 YouTube

[Terms](#)

[Creators](#)

[Developers](#)

[Press](#)

[Preferred Lineups](#)

Connect

 [Twitter](#)

 [Instagram](#)

 [Facebook](#)

EXHIBIT F

PUBLIC POLICY

How we're supporting smart regulation and policy innovation in 2019



The past year has seen a range of public debates about the roles and responsibilities of technology companies. As 2019 begins, I'd like to share my thoughts on these important discussions and why Google supports smart regulation and other innovative ways to address emerging issues.

We've always been (and still are) fundamentally optimistic about the power of innovative technology. We're proud that Google's products and services empower billions of people, drive economic growth and offer important tools for your everyday life. This takes many forms, whether it's instant access to the world's information, an infinite gallery of sortable photos, tools that let you share documents and calendars with friends, directions that help you avoid traffic jams, or whatever Google tool you find most helpful.

But this optimism doesn't obscure the challenges we face—including those posed by misuse of new technologies. New tools inevitably affect not just the people and businesses who use them, but also cultures, economies and societies as a whole. We've come a long way from our days as a scrappy startup, and with billions of people using our services every day, we recognize the need to confront tough issues regarding technology's broader impacts.

The scrutiny of lawmakers and others often improves our products and the policies that govern them. It's sometimes claimed that the internet is an

unregulated "wild west," but that's just not the case. Many laws and regulations have contributed to the internet's vitality: competition and

The Keyword

consumer protection laws, advertising regulations, and copyright, to name just a few. Existing legal frameworks reflect trade-offs that help everyone reap the benefits of modern technologies, minimize social costs, and respect fundamental rights. As technology evolves, we need to stay attuned to how best to improve those rules.

In some cases, laws do need updates, as we laid out in our [recent post on data protection](#) and our proposal regarding [law enforcement access to data](#). In other cases, collaboration among industry, government, and civil society may lead to complementary approaches, like joint industry efforts to fight online [terrorist content](#), [child sexual abuse material](#), and [copyright piracy](#). Shared concerns can also lead to ways to empower people with new tools and choices, like helping people control and move their data—that's why we have been a leader since 2007 in developing data portability tools and last year helped launch the cross-company [Data Transfer Project](#).

We don't see smart regulation as a singular end state; it must develop and evolve. In an era (and a sector) of rapid change, one-size-fits-all solutions are unlikely to work out well. Instead, it's important to start with a focus on a specific problem and seek well-tailored and well-informed solutions, thinking through the benefits, the second-order impacts, and the potential for unintended side-effects.

Efforts to address illegal and harmful online content illustrate how tech companies can help support this process:

First, to support constructive transparency, we launched our [Transparency Report](#) more than eight years ago, and we have continued to extend our transparency efforts over time, most recently with YouTube's [Community Guidelines enforcement report](#).

Second, to cultivate best practices for responsible content removals, we've supported initiatives like the [Global Internet Forum to Counter Terrorism](#), where tech companies, governments and civil society have worked together to stop exploitation of online services.

Finally, we have participated in government-overseen systems of accountability. For instance, the EU's [Hate Speech Code of Conduct](#) includes an audit process to monitor how platforms are meeting our commitments. And in the recent [EU Code of Practice On Disinformation](#), we agreed to help researchers study this topic and to engage in regular reporting and assessment of [our next steps](#) in this fight.

While the world is no longer at the start of the Information Revolution, the most important and exciting chapters are still to come. Google has pioneered a number of new artificial intelligence (AI) tools, and published a [set of principles](#) to guide our work and inform the larger public debate about the use of these remarkable technologies. We'll have more to say about issues in AI governance in the coming weeks. Of course, every new breakthrough will raise its own set of new issues—and we look forward to hearing from others and sharing our own thoughts and ideas.

POSTED IN:

[PUBLIC POLICY](#) —

[KENT WALKER](#)

PUBLIC POLICY

Smart regulation for combating illegal content

Kent Walker

SVP of Global Affairs

Published Feb 14, 2019



We've written before about how we're working to [support smart regulation](#), and one area of increasing attention is regulation to combat illegal content.

As online platforms have become increasingly popular, there's been a rich debate about the best legal framework for combating illegal content in a way that respects other social values, like free expression, diversity and innovation. Today, various laws provide detailed regulations, including [Section 230 of the Communications Decency Act](#) in the United States and European Union's [e-Commerce Directive](#).

Google [invests millions](#) of dollars in [technology and people](#) to combat illegal content in an effective and fair way. It's a complex task, and—just as in offline contexts—it's not a problem that can be totally solved. Rather, it's a problem that must be managed, and we are constantly refining our practices.

In addressing illegal content, we're also conscious of the importance of protecting legal speech. Context often matters when determining whether content is illegal. Consider a video of military conflict. In one context the footage might be documentary evidence of atrocities in areas where journalists have great difficulty and danger accessing. In another context the footage could be promotional material for an illegal organization. Even a

highly trained reviewer could have a hard time telling the difference, and we need to get those decisions right across many different languages and

The Keyword

cultures, and across the vast scale of audio, video, text, and images uploaded online. We make it easy to easily submit takedown notices; at the same time, we also create checks and balances against misuse of removal processes. And we look to the work of international agencies and principles from leading groups like the [Global Network Initiative](#).

A smart regulatory framework is essential to enabling an appropriate approach to illegal content. We wanted to share four key principles that inform our practices and that (we would suggest) make for an effective regulatory framework:

Shared Responsibility: Tackling illegal content is a societal challenge—in which companies, governments, civil society, and users all have a role to play. Whether a company is alleging copyright infringement, an individual is claiming defamation, or a government is seeking removal of terrorist content, it's essential to provide clear notice about the specific piece of content to an online platform, and then platforms have a responsibility to take appropriate action on the specific content. In some cases, content may not be clearly illegal, either because the facts are uncertain or because the legal outcome depends on a difficult balancing act; in turn, courts have an essential role to play in fact-finding and reaching legal conclusions on which platforms can rely.

Rule of law and creating legal clarity: It's important to clearly define what platforms can do to fulfill their legal responsibilities, including removal obligations. An online platform that takes other voluntary steps to address illegal content should not be penalized. (This is sometimes called “Good Samaritan” protection.)

Flexibility to accommodate new technology: While laws should accommodate relevant differences between platforms, given the fast-evolving nature of the sector, laws should be written in ways that address the underlying issue rather than focusing on existing technologies or mandating specific technological fixes.

Fairness and transparency: Laws should support companies' ability to publish transparency reports about content removals, and provide people with notice and an ability to appeal removal of content. They should also recognize that fairness is a flexible and context-dependent notion—for example, improperly blocking newsworthy content or political expression could cause more harm than mistakenly blocking other types of content.

With these principles in mind, we support refinement of notice-and-takedown regimes, but we have significant concerns about laws that would mandate proactively monitoring or filtering content, impose overly rigid timelines for content removal, or otherwise impose harsh penalties even on those acting in good faith. These types of laws create a risk that platforms won't take a balanced approach to content removals, but instead take a “better safe than sorry” approach—blocking content at upload or implementing a “take down first, ask questions later (or never)” approach. We regularly receive [overly broad](#) removal requests, and [analyses](#) of cease-

and-desist and takedown letters have found that many seek to remove potentially legitimate or protected speech.

The Keyword

There's ample room for debate and nuance on these topics—we discuss them every day—and we'll continue to seek ongoing collaboration among governments, industry, and civil society on this front. Over time, an ecosystem of tools and institutions—like the [Global Internet Forum to Counter Terrorism](#), and the [Internet Watch Foundation](#), which has taken down child sexual abuse material for more than two decades—has evolved to address the issue. Continuing to develop initiatives like these and other multistakeholder efforts remains critical, and we look forward to progressing those discussions.

POSTED IN:
[PUBLIC POLICY](#)

RELATED ARTICLES

Doing our part to share open data responsibly

To help fight the opioid crisis, a new tool from Maps and Search

Investing across the U.S. in 2019

Oracle v. Google and the future of software development

Engaging policy stakeholders on issues in AI governance

Applications are open for the Google North America Public Policy Fellowship

○ ○ ○ ○ ○ ○

EXHIBIT G



Official Blog

Broadcast Yourself

Faster removals and tackling comments – an update on what we're doing to enforce YouTube's Community Guidelines

Thursday, December 13, 2018

We've always used a mix of human reviewers and technology to address violative content on our platform, and in 2017 we started applying more advanced machine learning technology to flag content for review by our teams. This combination of smart detection technology and highly-trained human reviewers has enabled us to consistently enforce our policies with increasing speed.

We are committed to tackling the challenge of quickly removing content that violates our Community Guidelines and reporting on our progress. That's why in April we launched a quarterly [YouTube Community Guidelines Enforcement Report](#). As part of this ongoing commitment to transparency, today we're expanding the report to include additional data like channel removals, the number of comments removed, and the policy reason why a video or channel was removed.

Focus on removing violative content before it is viewed

We [previously shared](#) how technology is helping our human review teams remove content with speed and volume that could not be achieved with people alone. Finding all violative content on YouTube is an immense challenge, but we see this as one of our core responsibilities and are focused on continuously working towards removing this content before it is widely viewed.

- From July to September 2018, we removed 7.8 million videos
- And 81% of these videos were first detected by machines
- Of those detected by machines, 74.5% had never received a single view

When we detect a video that violates our Guidelines, we remove the video and apply a strike to the channel. We terminate entire channels if they are dedicated to posting content prohibited by our Community Guidelines or contain a single egregious violation, like child sexual exploitation. The vast majority of attempted abuse comes from bad actors trying to upload spam or adult content: over 90% of the channels and over 80% of the videos that we removed in September 2018 were removed for violating our policies on spam or adult content.

Looking specifically at the most egregious, but low-volume areas, like violent extremism and child safety, our significant investment in fighting this type of content is having an impact: Well over 90% of the videos uploaded in September 2018 and removed for Violent Extremism or Child Safety had fewer than 10 views.

Each quarter we may see these numbers fluctuate, especially when our teams tighten our policies or enforcement on a certain category to remove more content. For example, over the last year [we've strengthened our child safety enforcement](#), regularly consulting with experts to make sure our policies capture a broad range of content that may be harmful to children, including things like minors fighting or engaging in potentially dangerous dares. Accordingly, we saw that 10.2% of video removals were for child safety, while Child Sexual Abuse Material (CSAM) represents a fraction of a percent of the content we remove.

Making comments safer

As with videos, we use a combination of smart detection technology and human reviewers to flag, review, and remove spam, hate speech, and other abuse in comments.

We've also built [tools](#) that allow creators to moderate comments on their videos. For example, creators can choose to hold all comments for review, or to automatically hold comments that have links or may contain offensive content. Over one million creators now use these tools to moderate their channel's comments.¹

We've also been increasing our enforcement against violative comments:

- From July to September of 2018, our teams removed over 224 million comments for violating our Community Guidelines.
- The majority of removals were for spam and the total number of removals represents a fraction of the billions of comments posted on YouTube each quarter.

- As we have removed more comments, we've seen our comment ecosystem actually grow, not shrink. Daily users are 11% more likely to be commenters than they were last year.

We are committed to making sure that YouTube remains a vibrant community, where creativity flourishes, independent creators make their living, and people connect worldwide over shared passions and interests. That means we will be unwavering in our fight against bad actors on our platform and our efforts to remove egregious content before it is viewed. We know there is more work to do and we are continuing to invest in people and technology to remove violative content quickly. We look forward to providing you with more updates.

YouTube Team

¹ Creator comment removals on their own channels are not included in our reporting as they are based on opt-in creator tools and not a review by our teams to determine a Community Guidelines violation.



EXHIBIT H



Official Blog

Broadcast Yourself

Continuing our work to improve recommendations on YouTube

Friday, January 25, 2019

When recommendations are at their best, they help users find a new song to fall in love with, discover their next favorite creator, or learn that great paella recipe. That's why we update our recommendations system all the time—we want to make sure we're suggesting videos that people actually want to watch.

You might remember that a few years ago, viewers were getting frustrated with clickbait videos with misleading titles and descriptions ("You won't believe what happens next!"). We responded by updating our system to focus on viewer satisfaction instead of views, including measuring likes, dislikes, surveys, and time well spent, all while recommending clickbait videos less often. More recently, people told us they were getting too many similar recommendations, like seeing endless cookie videos after watching just one recipe for snickerdoodles. We now pull in recommendations from a wider set of topics—on any given day, more than 200 million videos are recommended on the homepage alone. In fact, in the last year alone, we've made hundreds of changes to improve the quality of recommendations for users on YouTube.

We'll continue that work this year, including taking a closer look at how we can reduce the spread of content that comes close to—but doesn't quite cross the line of—violating our [Community Guidelines](#). To that end, we'll begin reducing recommendations of borderline content and content that could misinform users in harmful ways—such as videos promoting a phony miracle cure for a serious illness, claiming the earth is flat, or making blatantly false claims about historic events like 9/11.

While this shift will apply to less than one percent of the content on YouTube, we believe that limiting the recommendation of these types of videos will mean a better

experience for the YouTube community. To be clear, this will only affect recommendations of what videos to watch, not whether a video is available on YouTube. As always, people can still access all videos that comply with our Community Guidelines and, when relevant, these videos may appear in recommendations for channel subscribers and in search results. We think this change strikes a balance between maintaining a platform for free speech and living up to our responsibility to users.

This change relies on a combination of machine learning and real people. We work with human evaluators and experts from all over the United States to help train the machine learning systems that generate recommendations. These evaluators are trained using [public guidelines](#) and provide critical input on the quality of a video.

This will be a gradual change and initially will only affect recommendations of a very small set of videos in the United States. Over time, as our systems become more accurate, we'll roll this change out to more countries. It's just another step in an ongoing process, but it reflects our commitment and sense of responsibility to improve the recommendations experience on YouTube.

— The YouTube Team



EXHIBIT I

GOOGLE IN EUROPE

Fighting disinformation across our products

Kristie Canegallo

VP of Trust & Safety, Google

Published Feb 16, 2019



Providing useful and trusted information at the scale that the Internet has reached is enormously complex and an important responsibility. Adding to that complexity, over the last several years we've seen organized campaigns use online platforms to deliberately spread false or misleading information.

We have twenty years of experience in these information challenges and it's what we strive to do better than anyone else. So while we have more work to do, we've [been working hard](#) to [combat](#) this challenge for many years.

Today at the [Munich Security Conference](#), we presented a white paper that gives more detail about our work to tackle the intentional spread of misinformation—across Google Search, Google News, YouTube and our advertising systems. We have a significant effort dedicated to this work throughout the company, based on three foundational pillars:

- Improve our products so they continue to make quality count;
- Counteract malicious actors seeking to spread disinformation;
- Give people context about the information they see.

The [white paper](#) also explains how we work beyond our products to support a healthy journalistic ecosystem, partner with civil society and researchers, and stay one step ahead of future risks.

We hope this paper and increased transparency can lead to more dialogue about what we and others can do better on these issues. We're committed to acting responsibly and thoroughly as we tackle this important challenge.

EXHIBIT J

February 2019

Google

How Google Fights Disinformation

Our work to tackle the intentional spread of misinformation across Google Search, Google News, YouTube, and our advertising platforms.

Introduction	2
What is disinformation?.....	2
Tackling disinformation across Google.....	3
Teaming up with newsrooms and outside experts.....	5
Getting ahead of future risks.....	7
Google Search, Google News & Disinformation	9
Background.....	9
Tackling disinformation on Google Search & Google News.....	10
YouTube & Disinformation	17
Background.....	17
Tackling disinformation on YouTube.....	18
Google Advertising Products & Disinformation	25
Background.....	25
Tackling disinformation on Google’s advertising platforms.....	26
Conclusion	29

Introduction

The open Internet has enabled people to create, connect, and distribute information like never before. It has exposed us to perspectives and experiences that were previously out-of-reach. It has enabled increased access to knowledge for everyone.

Google continues to believe that the Internet is a boon to society – contributing to global education, healthcare, research, and economic development by enabling citizens to become more knowledgeable and involved through access to information at an unprecedented scale.

However, like other communication channels, the open Internet is vulnerable to the organized propagation of false or misleading information. Over the past several years, concerns that we have entered a “post-truth” era have become a controversial subject of political and academic debate.

These concerns directly affect Google and our mission – to organize the world’s information and make it universally accessible and useful. When our services are used to propagate deceptive or misleading information, our mission is undermined.

How companies like Google address these concerns has an impact on society and on the trust users place in our services. We take this responsibility very seriously and believe it begins with providing transparency into our policies, inviting feedback, enabling users, and collaborating with policymakers, civil society, and academics around the world.

This document outlines our perspective on disinformation and misinformation and how we address it throughout Google. It begins with the three strategies that comprise our response across products, and an overview of our efforts beyond the scope of our products. It continues with an in-depth look at how these strategies are applied, and expanded, to Google Search, Google News, YouTube, and our advertising products.

We welcome a dialogue about what works well, what does not, and how we can work with others in academia, civil society, newsrooms, and governments to meet the ever-evolving challenges of disinformation.

What is disinformation?

As we’ve all experienced over the past few years, the words “misinformation”, “disinformation”, and “fake news” mean different things to different people and can become politically charged when they are used to characterize the propagators of a specific ideology or to undermine political adversaries.

However, there is something objectively problematic and harmful to our users when malicious *actors attempt to deceive them*. It is one thing to be wrong about an issue. It is another to purposefully disseminate information one knows to be inaccurate with the hope that others believe it is true or to create discord in society.

We refer to these deliberate efforts to deceive and mislead using the speed, scale, and technologies of the open web as “disinformation”.

The entities that engage in disinformation have a diverse set of goals. Some are financially motivated, engaging in disinformation activities for the purpose of turning a profit. Others are politically motivated, engaging in disinformation to foster specific viewpoints among a population, to exert influence over political processes, or for the sole purpose of polarizing and fracturing societies. Others engage in disinformation for their own entertainment, which often involves bullying, and they are commonly referred to as “trolls”.

Levels of funding and sophistication vary across those entities, ranging from local mom-and-pop operations to well-funded and state-backed campaigns. In addition, propagators of disinformation sometimes end up working together, even unwittingly. For instance, politically motivated actors might emphasize a piece of disinformation that financially motivated groups might latch onto because it is getting enough attention to be a potential revenue source. Sometimes, a successful disinformation narrative is propagated by content creators who are acting in good faith and are unaware of the goals of its originators.

This complexity makes it difficult to gain a full picture of the efforts of actors who engage in disinformation or gauge how effective their efforts may be. Furthermore, because it can be difficult to determine whether a propagator of falsehoods online is acting in good faith, responses to disinformation run the risk of inadvertently harming legitimate expression.

Tackling disinformation in our products and services

We have an important responsibility to our users and to the societies in which we operate to curb the efforts of those who aim to propagate false information on our platforms. At the same time, we respect our users' fundamental human rights (such as free expression) and we try to be clear and predictable in our efforts, letting users and content creators decide for themselves whether we are operating fairly. Of course, this is a delicate balance, as sharing too much of the granular details of how our algorithms and processes work would make it easier for bad actors to exploit them.

We face complex trade-offs and there is no 'silver bullet' that will resolve the issue of disinformation, because:

- It can be extremely difficult (or even impossible) for humans or technology to determine the veracity of, or intent behind, a given piece of content, especially when it relates to current events.
- Reasonable people can have different perspectives on the right balance between risks of harm to good faith, free expression, and the imperative to tackle disinformation.
- The solutions we build have to apply in ways that are understandable and predictable for users and content creators, and compatible with the kind of automation that is required when operating services on the scale of the web. We cannot create standards that require deep deliberation for every individual decision.
- Disinformation manifests differently on different products and surfaces. Solutions that might be relevant in one context might be irrelevant or counter-productive in others. Our products cannot operate in the exact same way in that regard, and this is why they approach disinformation in their own specific ways.

Our approach to tackling disinformation in our products and services is based around a framework of three strategies: make quality count in our ranking systems, counteract malicious actors, and give users more context. We will outline them in this section, as well as the efforts we undertake beyond the scope of our products and services to team up with newsrooms and outside experts, and to get ahead of future risks. It is worth noting that these strategies are also used to address misinformation more broadly, which pertains to the overall trustworthiness of the information we provide users in our products.

In later sections of this paper, we will detail how these strategies are implemented and expanded for Google Search, Google News, YouTube, and our advertising platforms. We adopt slightly different approaches in how we apply these principles to different products given how each service presents its own unique challenges.

1. Make Quality Count

Our products are designed to sort through immense amounts of material and deliver content that best meets our users' needs. This means delivering quality information and trustworthy commercial messages, especially in contexts that are prone to rumors and the propagation of false information (such as breaking news events).

While each product and service implements this differently, they share important principles that ensure our algorithms treat websites and content creators fairly and evenly:

- Information is organized by "ranking algorithms".
- These algorithms are geared toward ensuring the usefulness of our services, as measured by user testing, not fostering the ideological viewpoints of the individuals that build or audit them. When it comes to Google Search, you can find a detailed explanation of how those algorithms operate at google.com/search/howsearchworks.

2. Counteract Malicious Actors

Algorithms cannot determine whether a piece of content on current events is true or false, nor can they assess the intent of its creator just by reading what's on a page. However, there are clear cases of intent to manipulate or deceive users. For instance, a news website that alleges it contains "*Reporting from Bordeaux, France*" but whose account activity indicates that it is operated out of New Jersey in the U.S. is likely not being transparent with users about its operations or what they can trust it to know firsthand.

That's why our policies across Google Search, Google News, YouTube, and our advertising products clearly outline behaviors that are prohibited – such as misrepresentation of one's ownership or primary purpose on Google News and our advertising products, or impersonation of other channels or individuals on YouTube.

Furthermore, since the early days of Google and YouTube, many content creators have tried to deceive our ranking systems to get more visibility – a set of practices we view as a form of 'spam' and that we've invested significant resources to address.

This is relevant to tackling disinformation since many of those who engage in the creation or propagation of content for the purpose to deceive often deploy similar tactics in an effort to achieve more visibility. Over the course of the past two decades, we have invested in systems that can reduce 'spammy' behaviors at scale, and we complement those with human reviews.

3. Give Users More Context

Easy access to context and a diverse set of perspectives are key to providing users with the information they need to form their own views. Our products and services expose users to numerous links or videos in response to their searches, which maximizes the chances that users are exposed to diverse perspectives or viewpoints before deciding what to explore in depth.

Google Search, Google News, YouTube, and our advertising products have all developed additional mechanisms to provide more context and agency to users. Those include:

- **“Knowledge” or “Information” Panels** in Google Search and YouTube, providing high-level facts about a person or issue.
- Making it easier to discover the work of **fact-checkers** on Google Search or Google News, by using labels or snippets making it clear to users that a specific piece of content is a fact-checking article.
- A **“Full Coverage”** function in Google News enabling users to access a non-personalized, in-depth view of a news cycle at the tap of a finger.
- **“Breaking News”** and **“Top News”** shelves, and **“Developing News”** information panels on YouTube, making sure that users are exposed to news content from authoritative sources when looking for information about ongoing news events.
- Information panels providing **“Topical Context”** and **“Publisher Context”** on YouTube, providing users with contextual information from trusted sources to help them be more informed consumers of content on the platform. These panels provide authoritative information on well-established historical and scientific topics that have often been subject to misinformation online and on the sources of news content, respectively.
- **“Why this ad”** labels enabling users to understand why they’re presented with a specific ad and how to change their preferences so as to alter the personalization of the ads they are shown, or to opt out of personalized ads altogether.
- **In-ad disclosures and transparency reports on election advertising**, which are rolling out during elections in the US, Europe, and India as a starting point.

We also empower users to let us know when we’re getting it wrong by using **feedback buttons** across Search, YouTube, and our advertising products to flag content that might be violating our policies.

Teaming up with newsrooms and outside experts

Our work to address disinformation is not limited to the scope of our products and services. Indeed, other organizations play a fundamental role in addressing this societal challenge, such as newsrooms, fact-checkers, civil society organizations, or researchers. While we all address different aspects of this issue, it is only by coming together that we can succeed. That is why we dedicate significant resources to supporting quality journalism, and to weaving together partnerships with many other organizations in this space.

Supporting quality journalism

People come to Google looking for information they can trust and that information often comes from the reporting of journalists and news organizations around the world.

A thriving news ecosystem matters deeply to Google and directly impacts our efforts to combat disinformation. When quality journalism struggles to reach wide audiences, malicious actors have more room to propagate false information.

Over the years, we've worked closely with the news industry to address these challenges and launched products and programs to help improve the business model of online journalism. These include the Accelerated Mobile Pages Project¹ to improve the mobile web, YouTube Player for Publishers² to simplify video distribution and reduce costs, and many more.

In March 2018, we launched the Google News Initiative (GNI)³ to help journalism thrive in the digital age. With a \$300 million commitment over 3 years, the initiative aims to elevate and strengthen quality journalism, evolve business models to drive sustainable growth, and empower news organizations through technological innovation. \$25M of this broader investment was earmarked as innovation grants for YouTube to support news organizations in building sustainable video operations.

One of the programs supported by the Google News Initiative is Subscribe with Google⁴, a way for people to easily subscribe to various news outlets, helping publishers engage readers across Google and the web. Another is News Consumer Insights, a new dashboard built on top of Google Analytics, which will help news organizations of all sizes understand and segment their audiences with a subscriptions strategy in mind. More details on these projects and others can be found at g.co/newsinitiative.

Partnering with outside experts

Addressing disinformation is not something we can do on our own. The Google News Initiative also houses our products, partnerships, and programs dedicated to supporting **news organizations** in their efforts create quality reporting that displaces disinformation. This includes:

- Helping to launch the **First Draft Coalition** (<https://firstdraftnews.org/>), a nonprofit that convenes news organizations and technology companies to tackle the challenges around combating disinformation online – especially in the run-up to elections.
- Participating in and providing financial support to the **Trust Project** (<http://thetrustproject.org/>), of which Google is a founding member and which explores how journalism can signal its trustworthiness online. The Trust Project has developed eight indicators of trust that publishers can use to better convey why their content should be seen as credible, with promising results for the publishers who have trialed them.
- Partnering with Poynter's **International Fact-Checking Network** (IFCN)⁵, a nonpartisan organization gathering fact-checking organizations from the United States, Germany, Brazil, Argentina, South Africa, India, and more.

In addition, we support the work of **researchers** who explore the issues of disinformation and trust in journalism by funding research at organizations like First Draft, the Oxford University's Reuters News Institute, Michigan University's Quello Center for Telecommunication Management law, and more.

Finally, in March 2018, Google.org (Google's philanthropic arm) launched a \$10 million global initiative to support **media literacy around the world** in the footsteps of programs we have already supported in the UK, Brazil, Canada, Indonesia, and more.

We will continue to explore more ways to partner with others on these issues, whether by building new products that might benefit the work of journalists and fact-checkers, supporting more independent initiatives that help curb disinformation, or developing self-regulatory practices to demonstrate our responsibility.

Getting ahead of future risks

Creators of disinformation will never stop trying to find new ways to deceive users. It is our responsibility to make sure we stay ahead of the game. Many of the product strategies and external partnerships mentioned earlier help us reach that goal. In addition, we dedicate specific focus to bolstering our defenses in the run-up to elections and invest in research and development efforts to stay ahead of new technologies or tactics that could be used by malicious actors, such as synthetic media (also known as 'deep fakes').

Protecting elections

Fair elections are critical to the health of democracy and we take our work to protect elections very seriously. Our products can help make sure users have access to accurate information about elections. For example, we often partner with election commissions, or other official sources, to make sure key information like the location of polling booths or the dates of the votes are easily available to users.

We also work to protect elections from attacks and interference, including focusing on combating political influence operations, improving account and website security, and increasing transparency.

To prevent political influence operations, working with our partners at Jigsaw, we have multiple internal teams that identify malicious actors wherever they originate, disables their accounts, and shares threat information with other companies and law enforcement officials. We routinely provide public updates about these operations.⁷

There is more we can do beyond protecting our own platforms. Over the past several years, we have taken steps to help protect accounts, campaigns, candidates, and officials against digital attacks. Our Protect Your Election project⁸ offers a suite of extra security to protect against malicious or insecure apps and guards against phishing. To protect election and campaign websites, we also offer Project Shield⁹, which can mitigate the risk of Distributed Denial of Service (DDoS) attacks.

In the run-up to elections, we provide free training to ensure that campaign professionals and political parties are up-to-speed on the means to protect themselves from attack. For instance, in 2018, we trained more than 1,000 campaign professionals and the eight major U.S. Republican and Democratic committees on email and campaign website security.

Furthermore, as a part of our security efforts, for the past eight years, we have displayed warnings to Gmail users who are at risk of phishing by potentially state-sponsored actors (even though, in most cases, the specific phishing attempt never reaches the user's inbox).

Finally, in order to help understand the context for the election-related ads they see online, we require additional verification for advertisers who wish to purchase political ads in the United States, provide transparency about the advertiser to the user, and have established an online transparency report and creative repository on US federal elections.¹⁰

We look forward to expanding these tools, trainings, and strategies to more elections in 2019, starting with efforts focused on two of the world's largest upcoming elections, which are in Europe¹¹ and in India.¹²

Expecting the unexpected

Creators of disinformation are constantly exploring new ways to bypass the defenses set by online services in an effort to spread their messages to a wider audience.

To stay ahead of the curve, we continuously invest resources to stay abreast of the next tools, tactics, or technologies that creators of disinformation may attempt to use. We convene with experts all around the world to understand what concerns them. We also invest in research, product, and policy developments to anticipate threat vectors that we might not be equipped to tackle at this point.

One example is the rise of new forms of AI-generated, photo-realistic, synthetic audio or video content known as "synthetic media" (often referred to as "deep fakes"). While this technology has useful applications (for instance, by opening new possibilities to those affected by speech or reading impairments, or new creative grounds for artists and movie studios around the world), it raises concerns when used in disinformation campaigns and for other malicious purposes.

The field of synthetic media is fast-moving and it is hard to predict what might happen in the near future. To help prepare for this issue, Google and YouTube are investing in research to understand how AI might help detect such synthetic content as it emerges, working with leading experts in this field from around the world.

Finally, because no detector can be perfect, we are engaging with civil society, academia, newsrooms, and governments to share our best understanding of this challenge and work together on what other steps societies can take to improve their preparedness. This includes exploring ways to help others come up with their own detection tools. One example may involve releasing datasets of synthesized content that others can use to train AI-based detectors.¹³

Google Search, Google News & Disinformation

Background

Google created its search engine in 1998, with a mission to organize the world's information and make it universally accessible and useful. At the time, the web consisted of just 25 million pages.

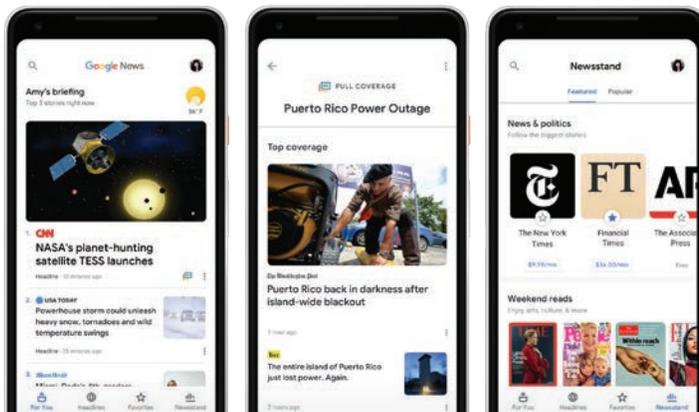
Today, we index hundreds of billions of pages – more information than all the libraries in the world could hold – and serve people all over the world. Search is offered in more than 150 languages and over 190 countries.

We continue to improve on Search every day. In 2017 alone, Google conducted more than 200,000 experiments that resulted in about 2,400 changes to Search. Each of those changes is tested to make sure it aligns with our publicly available Search Quality Rater Guidelines,¹⁴ which define the goals of our ranking systems and guide the external evaluators who provide ongoing assessments of our algorithms.

Over the past 20 years, we have grappled with the tension between the open access to information and expression that the web enables and the need to ensure trust in authoritative information. Our work on disinformation continues to be informed by these dual goals, as we attempt to strike the right balance in tackling this challenge.

Different types of content may require different approaches to ranking and presentation in order to meet our users' needs. Google News arose from such a realization and was one of the first products Google launched beyond Search. Former Googler Krishna Bharat observed that when people searched for news after the tragic 9/11 attacks in New York, Google responded with old news stories about New York rather than the latest events. He set about to fix that, and on September 22, 2002, Google News was born.

Over time, Google News has improved, including how we present content related to current events in Google Search. In 2018, we launched a reimagined Google News that uses a new set of AI techniques to take a constant flow of information as it hits the web, analyze it in real-time, and organize it around breaking news events.¹⁵



Google News App

Through all of this, we've remained grounded in our mission and the importance of providing greater access to information, helping users navigate the open web. We continue to believe that this access is fundamental to helping people make sense of the world around them, exercise their own critical thinking, and make informed decisions as citizens.

Tackling disinformation on Google Search & Google News

Since Google's early days, malicious actors have attempted to harm or deceive Search users through a wide range of actions, including tricking our systems in order to promote their own content (via a set of practices we refer to as "spam"), propagating malware, and engaging in illegal acts online. The creators and purveyors of disinformation employ many of the same tactics.

Disinformation poses a unique challenge. Google is not in a situation to assess objectively, and at scale, the veracity of a piece of content or the intent of its creators. Further, a considerable percentage of content contains information that cannot be objectively verified as fact. This is because it either lacks necessary context, because it is delivered through an ideological lens others may disagree with, or because it is constructed from contested datapoints.

Disinformation also raises broader concerns of harm. In the worst cases, the impacts of disinformation campaigns can affect an entire society. The stakes of accurately identifying disinformation are higher because disinformation often concerns issues at the core of political society for which the free exchange of ideas and information among genuine voices is of the greatest importance.

To deal with this issue, Google Search and Google News take a pragmatic approach that reinforces the product strategies we have highlighted in the opening section of this paper:

- **Make Quality Count**
 - We use ranking algorithms to elevate authoritative, high-quality information in our products.
 - We take additional steps to improve the quality of our results for contexts and topics that our users expect us to handle with particular care.
- **Counteract Malicious Actors**
 - We look for and take action against attempts to deceive our ranking systems or circumvent our policies.
- **Give Users More Context**
 - We provide users with tools to access the context and diversity of perspectives they need to form their own views.

Do Google News and Google Search combat disinformation in the same ways?

Google News' focus – coverage of current events – is narrower than that of Google Search. However, their goals are closely related. Both products present users with trustworthy results that meet their information needs about the issues they care about.

For that reason, both products have a lot in common when it comes to the way they operate. For instance, ranking in Google News is built on the basis of Google Search ranking and they share the same defenses against “spam” (attempts at gaming our ranking systems).

In addition, both products share some fundamental principles:

- They use **algorithms, not humans**, to determine the ranking of the content they show to users. No individual at Google ever makes determinations about the position of an individual webpage link on a Google Search or Google News results page.
- Our algorithms are geared toward **ensuring the usefulness of our services, as measured by user testing**, not fostering the ideological viewpoints of the individuals who build or audit them.
- The systems **do not make subjective determinations about the truthfulness of webpages**, but rather focus on **measurable signals** that correlate with how users and other websites value the expertise, trustworthiness, or authoritativeness of a webpage on the topics it covers.

That said, because Google News' purposes are explicitly narrower than those of Google Search and solely focused on coverage of current events, it builds its own *ranking systems* and *content policies* on top of those of Google Search.

When it comes to **ranking**, this means that the systems we use in Google News and in places that are focused on News in Google Search (e.g. our “Top Stories” Carousel or our “News” Tab) make special efforts to understand things like the prominence of a news story in the media landscape of the day, which articles most relate to that story, or which sources are most trusted for specific news topics. It also means that Google News might give additional importance toward factors that indicate a webpage's newsworthiness or journalistic value for users, such as its freshness or (for specific tabs within Google News).

When it comes to **content policies**:

- Google Search aims to make information from the web available to all our users. That's why we do not remove content from results in Google Search, except in very limited circumstances. These include legal removals, violations of our webmaster guidelines, or a request from the webmaster responsible for the page.
- Google Search contains some features that are distinct from its general results, such as Autocomplete. For features where Google specifically promotes or highlights content, we may remove content that violates their specific policies.¹⁶
- Because Google News does not attempt to be a comprehensive reflection of the web, but instead to focus on journalistic accounts of current events, it has more restrictive content policies than Google Search. Google News explicitly prohibits content that incites, promotes, or glorifies violence, harassment, or dangerous activities. Similarly, Google News does not allow sites or accounts that impersonate any person or organization, that misrepresent or conceal their ownership or primary purpose, or that engage in coordinated activity to mislead users.¹⁷

With those nuances in mind, it is still safe to think of Google News and Google Search's approaches to disinformation and misinformation are mostly similar, and the content of the following sections apply to both products. Where there is a difference, it will be outlined explicitly in the body of the text or in a dedicated callout box.

We use ranking algorithms to elevate high-quality information in our products

Ranking algorithms are an important tool in our fight against disinformation. Ranking elevates the relevant information that our algorithms determine is the most authoritative and trustworthy above information that may be less reliable. These assessments may vary for each webpage on a website and are directly related to our users' searches. For instance, a national news outlet's articles might be deemed authoritative in response to searches relating to current events, but less reliable for searches related to gardening.

For most searches that could potentially surface misleading information, there is high-quality information that our ranking algorithms can detect and elevate. When we succeed in surfacing high-quality results, lower quality or outright malicious results (such as disinformation or otherwise deceptive pages) are relegated to less visible positions in Search or News, letting users begin their journey by browsing more reliable sources.

Our ranking system does not identify the intent or factual accuracy of any given piece of content. However, it is specifically designed to identify sites with high indicia of expertise, authority, and trustworthiness.

How do Google's algorithms assess expertise, authority, and trustworthiness?

- Google's algorithms identify signals about pages that correlate with trustworthiness and authoritativeness. The best known of these signals is PageRank, which uses links on the web to understand authoritativeness.
- We are constantly evolving these algorithms to improve results – not least because the web itself keeps changing. For instance, in 2017 alone, we ran over 200,000 experiments with trained external Search Evaluators and live user tests, resulting in more than 2,400 updates to Google Search algorithms.
- To perform these evaluations, we work with Search Quality Evaluators who help us measure the quality of Search results on an ongoing basis. Evaluators assess whether a website provides users who click on it with the content they were looking for, and they evaluate the quality of results based on the expertise, authoritativeness, and trustworthiness of the content.
- The resulting ratings do not affect the ranking of any individual website, but they do help us benchmark the quality of our results, which in turn allows us to build algorithms that globally recognize results that meet high-quality criteria. To ensure a consistent approach, our evaluators use the Search Quality Rater Guidelines (publicly available online)¹⁸ which provide guidance and examples for appropriate ratings. To ensure the consistency of the rating program, Search Quality evaluators must pass a comprehensive exam and are audited on a regular basis.
- These evaluators also perform evaluations of each improvement to Search we roll out: in side-by-side experiments, we show evaluators two different sets of Search results, one with the proposed change already implemented and one without. We ask them which results they prefer and why. This feedback is central to our launch decisions.

For more information about how our ranking work, please visit:
www.google.com/search/howsearchworks

We take additional steps to improve the trustworthiness of our results for contexts and topics that our users expect us to handle with particular care.

Our Search Quality Raters Guidelines acknowledge that some types of pages could potentially impact the future happiness, health, financial stability, or safety of users. We call those “Your Money or Your Life” pages or YMYL. We introduced the YMYL category in 2014. They include financial transaction or information pages, medical and legal information pages, as well as news articles, and public and/or official information pages that are important for having an informed citizenry. This last category can comprise anything from information about local, state, or national government processes or policies, news about important topics in a given country, or disaster response services.

General Guidelines		July 20, 2018
General Guidelines Overview		
0.0 Introduction to Search Quality Rating		5
0.1 The Purpose of Search Quality Rating		6
0.2 Raters Must Represent the User		6
0.3 Browser Requirements		6
0.4 Ad Blocking Extensions		6
0.5 Internet Safety Information		6
Part 1: Page Quality Rating Guideline		
1.0 Introduction to Page Quality Rating		7
2.0 Understanding Webpages and Websites		7
2.1 Important Definitions		7
2.2 What is the Purpose of a Webpage?		8
2.3 Your Money or Your Life (YMYL) Pages		9
2.4 Understanding Webpage Content		9
2.4.1 Identifying the Main Content (MC)		9
2.4.2 Identifying the Supplementary Content (SC)		10
2.4.3 Identifying Advertisements/Monetization (Ads)		10
2.4.4 Summary of the Parts of the Page		11
2.5 Understanding the Website		11
2.5.1 Finding the Homepage		11
2.5.2 Finding Who is Responsible for the Website and Who Created the Content on the Page		13
2.5.3 Finding About Us, Contact Information, and Customer Service Information		13
2.6 Reputation of the Website or Creator of the Main Content		14

Our Search Quality Raters Guidelines

For these “YMYL” pages, we assume that users expect us to operate with our strictest standards of trustworthiness and safety. As such, where our algorithms detect that a user’s query relates to a “YMYL” topic, we will give more weight in our ranking systems to factors like our understanding of the authoritativeness, expertise, or trustworthiness of the pages we present in response. Similarly, we direct our Google Search evaluators to be more demanding in their assessment of the quality and trustworthiness of these page than they would otherwise. Specifically, in 2016, we added additional guidance to our Search Quality Rater Guidelines advising evaluators to give lower quality ratings to informational pages that contain demonstrably inaccurate content or debunked conspiracy theories. While their ratings don’t determine individual page rankings, they are used to help us gather data on the quality of our results and identify areas where we need to improve. This data from Search Evaluators also plays a significant role in determining which changes we roll out to our ranking systems.

Beyond specific types of content that are more sensitive to users, we realize that some contexts are more prone to the propagation of disinformation than others. For instance, breaking news events, and the heightened level of interest that they elicit, are magnets for bad behavior by malicious players. Speculation can outrun facts as legitimate news outlets on the ground are still investigating. At the same time, malicious actors are publishing content on forums and social media with the intent to mislead and capture people’s attention as they rush to find trusted information. To reduce the visibility of this type of content, we have designed our systems to prefer authority over factors like recency or exact word matches while a crisis is developing.

In addition, we are particularly attentive to the integrity of our systems in the run-up to significant societal moments in the countries where we operate, such as elections.¹⁹

We actively look for and take action against attempts to deceive our ranking systems or circumvent our policies.

Google is designed to help users easily discover and access the webpages that contain the information they are looking for. Our goals are aligned with those of site owners who publish high-quality content online because they want it to be discovered by users who might be interested. That’s why we provide extensive tools and tips

to help webmasters and developers manage their Search presence and succeed in having their content, sites, and apps found. We provide interactive websites, videos, starter guides, frequent blog posts, user forums, and live expert support to inform webmasters. Our publicly available webmaster guidelines complement these resources by outlining some of the tips and behaviors that we recommend webmasters adopt to make it easiest for our systems to crawl and index their websites.²⁰



Google Webmaster Guidelines

Not all site owners act in good faith. Since the early days of Google, many have attempted to manipulate their way to the top of Search results through deceptive or manipulative behavior, using any insights into the functioning of our systems they can get to try to circumvent them. The earliest example of such attempts dates back to 1999, when Google's founders published a seminal paper on PageRank, a key innovation in Google's algorithm.²¹ The paper described how our algorithms use links between websites as an indicator of authority. Once that paper was published, spammers tried to game Google by paying each other for links.

These manipulative behaviors aim to elevate websites to users not because they are the best response to a query, but because a webmaster has deceived our systems. As such, they are considered "spam" and run afoul of our core mission. Our webmaster guidelines clearly spell out actions that are prohibited and state that we will take action against websites engaging in such behaviors.

While not all spammers engage in disinformation, many of the malicious actors who try to distribute disinformation (at all levels of sophistication or funding) engage in some form of spam. The tactics they use are similar to those of other spammers. Therefore, our work against spam goes hand-in-hand with our work against disinformation.

Our algorithms can detect the majority of spam and demote or remove it automatically. The remaining spam is tackled manually by our spam removal team, which reviews pages (often based on user feedback) and flags them if they violate the Webmaster Guidelines. In 2017, we took action on 90,000 user reports of search spam and algorithmically detected many more times that number.

As our tactics improve and evolve, so does spam. One of the trends we observed in 2017 was an increase in website hacking, both for spamming search ranking and for spreading malware. We focused on reducing this threat and were able to detect and remove from Search results more than 80 percent of these sites over the following year.

We continue to be vigilant regarding techniques used by spammers and remain conscientious of what we share about the ways our ranking systems work so as not to create vulnerabilities they can exploit.

Google News policies against deceptive content

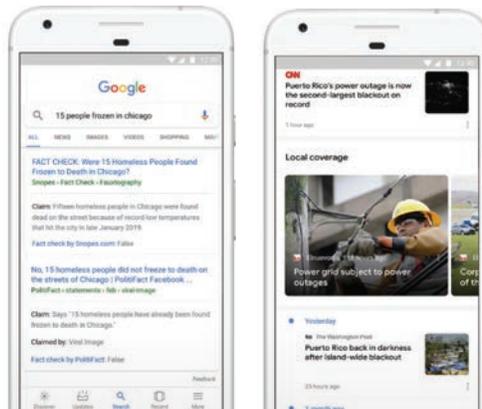
In addition to other efforts to fight spam, Google News’ content policies prohibit:

- Sites or accounts that impersonate any person or organization;
- Sites that misrepresent or conceal their ownership or primary purpose;
- Sites or accounts that engage in coordinated activity to mislead users – including, but not limited to, sites or accounts that misrepresent or conceal their country of origin or that direct content at users in another country under false pretenses.

In addition to algorithmic signals that might indicate such behavior, where there is an indication that a publisher may be violating our policies, such as through a user report or suspicious account activity, our Trust and Safety team will investigate and, where appropriate, take action against that site and related sites that can be confirmed to be operating in concert.

We provide users with the context and diversity of perspectives they need to form their own views.

From the very beginning of Google Search, the very nature of Google Search results pages has ensured that when looking for information on news or public interest topics they care about, users are presented with links to multiple websites and perspectives.



Fact-Check on Google Search

Local Coverage & Timeline in Google News

This remains true today. When users search for news on Google, they are always presented with multiple links. In many cases, they are also presented with additional elements that help them get more context about their search. For instance, “Knowledge Panels” might show in Search results to provide context and basic information about people, places, or things that Google knows about. Fact-check tags or snippets might show below links in Google Search and Google News, outlining that a specific piece of content purports to fact-check a claim made by a third party.²² Or we might call out related searches or questions that users tend to ask about the topic of a search query.

In Google News, additional cues may help users pick up on points of context that are particularly relevant to News stories, such as “Opinion” or “User-Generated Content” tags under articles that news publishers want to signal as such; or algorithmically generated story timelines that let users explore at-a-glance the milestones of a news story over the weeks or months that led to the day’s events.

Does Google personalize the content that shows in Google Search and Google News so that users only see news consistent with their views, sometimes known as “filter bubbles”?

We try to make sure that our users continue to have access to a diversity of websites and perspectives. Google Search and Google News take different approaches toward that goal.

Google Search: contrary to popular belief, there is very little personalization in Search based on users' inferred interests or Search history before their current session. It doesn't take place often and generally doesn't significantly change Search results from one person to another. Most differences that users see between their Search results and those of another user typing the same Search query are better explained by other factors such as a user's location, the language used in the search, the distribution of Search index updates throughout our data centers, and more.²³ Furthermore, the Top Stories carousel that often shows in Search results in response to news-seeking searches is never personalized.

Google News: To meet the needs of users who seek information on topics they care about, Google News aims to strike a balance between providing access to the same content and perspectives as other users and presenting content that relates to news topics one cares about. To do this, Google News offers three interconnected ways to discover information:

- **'Headlines' and 'Top Stories':** To help users stay on top of the top trending news in their country, the "Headlines" and "Top Stories" tabs shows the major stories and issues that news sources are covering at any point in time and shows them to everyone in a non-personalized manner.
- **For You:** To help users stay on top of the news that matter to them, the "For You" tab lets them specify the topics, publications, and locations they are interested in so they can see the news that relates to those selections. Additionally, and depending on their permission settings, the "For You" tab may show them news they may be interested in in light of their past activity on Google products.
- **Full Coverage:** To help users access context and diverse perspectives about the news stories they read, the "Full Coverage" feature in Google News lets users explore articles and videos from a variety of publishers related to an article or news story of their choice. The "Full Coverage" feature is not personalized and is accessible in one click or tap from most articles in the "For You" and "Headlines" tabs.

Importantly, for both services, **we never personalize content based on signals relating to point of view on issues and/or political leanings** – our systems do not collect such signals, nor do they have an understanding of political ideologies.

We constantly improve our algorithms, policies, and partnerships, and are open about issues we have yet to address.

Because the malicious actors who propagate disinformation have the incentive to keep doing so, they continue to probe for new ways to game our systems, and it is incumbent on us to stay ahead of this technological arms race. A compounding factor in that challenge is that our systems are constantly confronted with searches they have never seen before. Every day, 15% of the queries that our users type in the Google Search bar are new.

For these reasons, we regularly evolve our ranking algorithms, our content policies, and the partnerships we enter into as part of our efforts to curb disinformation.

We are aware that many issues remain unsolved at this point. For example, a known strategy of propagators of disinformation is to publish a lot of content targeted on "data voids", a term popularized by the U.S. based think tank Data and Society to describe Search queries where little high-quality content exists on the web for Google to display due to the fact that few trustworthy organizations cover them.²⁴ This often applies, for instance, to niche conspiracy theories, which most serious newsrooms or academic organizations won't make the effort to debunk. As a result, when users enter Search terms that specifically refer to these theories, ranking algorithms can only elevate links to the content that is actually available on the open web – potentially including disinformation.

We are actively exploring ways to address this issue, and others, and welcome the thoughts and feedback of researchers, policymakers, civil society, and journalists around the world.

YouTube & Disinformation

Background

YouTube started in 2005 as a video-sharing website and quickly evolved into one of the world's most vibrant online communities. Thousands, then millions, then billions of people connected through content that educated, excited, or inspired them. YouTube is one of the world's largest exporters of online cultural and learning content and is a significant driver of economic activity, providing many of its creators with the ability to make a livelihood by using its services.

Disinformation is not unique to YouTube. It is a global problem afflicting many platforms and publishers. When a platform fosters openness, as we do at YouTube, there's a risk that unreliable information will be presented. While disinformation has been a problem as long as there's been news to report, the Internet has made it possible for disinformation to spread further and faster than ever before. We take our responsibility to combat disinformation in this domain seriously. To be effective at our size, we invest in a combination of technological solutions with a large and growing base of human talent. Technology provides scale and speed, while human talent provides the contextual knowledge needed to fine-tune and improve our systems every step of the way.

YouTube has developed a comprehensive approach to tackling controversial content on our platform. This approach is guided by three principles:

1. Keep content on the platform unless it is in violation of our Community Guidelines
2. Set a high bar for recommendations
3. Monetization is a privilege

From these principles, we create robust systems to responsibly manage all types of controversial content, including disinformation.

Given how broad the spectrum of disinformation is, we implement the three product strategies mentioned in our opening section in ways that are relevant to YouTube's specific products, community, and challenges:

- **Make Quality Count**
 - We deploy effective product and ranking systems that demote low-quality disinformation and elevate more authoritative content
- **Counteract Malicious Actors**
 - We rigorously develop and enforce our content policies
 - We protect the integrity of information tied to elections through effective ranking algorithms, and tough policies against users that misrepresent themselves or who engage in other deceptive practices
 - We remove monetary incentives through heightened standards for accounts that seek to utilize any of YouTube's monetization products
- **Give Users Context**
 - We provide context to users via information panels on YouTube

We take our responsibilities as a platform seriously. We believe a responsible YouTube will continue to embrace the democratization of access to information while providing a reliable and trustworthy service to our users.

Tackling disinformation on YouTube

Given the spectrum of content and intent, it is necessary to have a nuanced approach that strikes the right balance between managing our users' expectations to express themselves freely on the platform with the need to preserve the health of the broader community of the creator, user, and advertiser ecosystem.

Let's take a closer look at the three guiding principles on which we base our approach for YouTube:

1. Keep content on the platform unless it is in violation of our Community Guidelines

YouTube's Community Guidelines²⁵ prohibit certain categories of material, including sexually explicit content, spam, hate speech, harassment and incitement to violence. We aim to balance free expression with preventing harmful content in order to maintain a vibrant community. Striking this balance is never easy, especially for a global platform. YouTube has always had Community Guidelines, but we revise them as user behavior changes and as the world evolves.

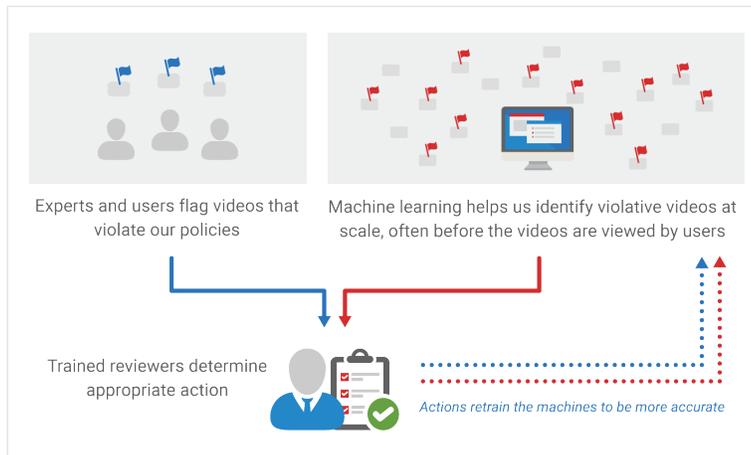
YouTube also maintains a more detailed and living set of enforcement guidelines that provide internal guidance on the enforcement of the public Community Guidelines. These enforcement guidelines are extensive and dynamic to ensure that the policies apply to changing trends and new patterns of controversial content online. YouTube does not typically disclose these updates to the public because doing so would make it easier for unscrupulous users to evade detection.

To help formulate rules that are consistent, unbiased, well-informed, and broad enough to apply to a wide scope of content, YouTube often relies on external subject matter experts and NGOs to consult on various issues. YouTube has also worked with independent experts as a member of the Global Network Initiative (GNI),²⁶ to establish key principles to guide content review efforts and systems, including notifying users if a video is removed and allowing for appeals. To honor YouTube's commitment to human rights, we also make exceptions to the Community Guidelines for material that is educational, documentary, scientific, and/or artistic.

Consistent enforcement

With hundreds of hours of new content uploaded to YouTube every minute, clear policies and enforcement guidelines are only part of what matters. To maintain a site where abuse is minimized, the systems used to curtail abuse must scale. YouTube has always relied on a mix of humans and technology to enforce its guidelines and will continue to do so.

YouTube has thousands of reviewers who operate 24/7 to address content that may violate our policies and the team is constantly expanding to meet evolving enforcement needs. Our review teams are diverse and global. Linguistic and cultural knowledge is needed to interpret the context of a flagged video and decide whether it violates our guidelines. Reviewers go through a comprehensive training program to ensure that they have a full understanding of YouTube's Community Guidelines. We use frequent tests as part of the training process to ensure quality and knowledge retention. Human reviewers are essential to evaluating context and to ensuring that educational, documentary, scientific, and artistic content is protected.



YouTube flagging and review decisions continuously improves our systems

We strive to be as transparent as possible when it comes to actions we take on content on our platform. That is why we release a Community Guidelines Enforcement Report²⁷ where we give insight into the scale and nature of our extensive policy enforcement efforts. It shows that YouTube's 'crime rate' is low – only a fraction of YouTube's total views are on videos that violate company policies.

Application to disinformation

There are several policies in the Community Guidelines that are directly applicable in some form to disinformation. These include policies against spam, deceptive practices, scams,²⁸ impersonation,²⁹ hate,³⁰ and harassment.³¹

The policy against spam, deceptive practices, and scams prohibits posting large amounts of untargeted, unwanted, or repetitive content in videos, comments, private messages, especially if the main purpose of the content is to drive people off to another site. Similarly, activity that seeks to artificially increase the number of views, likes, dislikes, comments, or other metrics either through the use of automated systems or by serving up videos to unsuspecting viewers is against our terms. Additionally, content that exists solely to incentivize viewers for engagement (views, likes, comments, etc.), or by coordinating at scale with other users to drive up views for the primary purpose of interfering with our systems, is prohibited.

One of the abuses this policy covers is content that deliberately seeks to spread disinformation that could suppress voting or otherwise interfere with democratic or civic processes. For example, demonstrably false content that claims one demographic votes on one day while another votes on a separate day would be in violation of our policies.

Another applicable policy regards impersonation. Accounts seeking to spread disinformation by misrepresenting who they are via impersonation are clearly against our policies and the account will be removed. For example, if a user copies a channel's profile, background, or text, and writes comments to make it look like somebody else's channel posted the comments, we remove the channel. Impersonation can also occur if a user creates a channel or video using another individual's real name, image, or other personal information to deceive people into thinking they are someone else on YouTube.

YouTube has clear policies against hate and harassment. Hate speech refers to content that promotes violence against, or has the primary purpose of inciting hatred against, individuals or groups based on certain attributes, such as race or ethnic origin, religion, disability, gender, age, veteran status, or sexual orientation/gender identity. Harassment may include abusive videos, comments, messages, revealing someone's personal information, unwanted sexualization, or incitement to harass other users or creators. Users that spread disinformation that runs afoul of either our hate or harassment policies will be removed and further appropriate action taken.

2. Set a high bar for recommendations

Our primary objective when it comes to our search and discovery systems is to help people find content that they will enjoy watching, whether through their Homepage, Watch Next, or Search results. We aim to provide content that lets users dive into topics they care about, broaden their perspective, and connect them to the current zeitgeist. When a user is seeking content with high intent – subscribing to a channel or searching for the video – it is our responsibility to help users find and watch the video. On the other hand, in the absence of strong or specific intent for a particular video, we believe it is our responsibility to not proactively recommend content that may be deemed low quality.

How our approach to recommendations has evolved

When YouTube's recommendation systems first launched, it sought to optimize for content that got users to click. We noticed that this system incentivized creators to publish misleading and sensationalist clickbait, and users would click on the video but very quickly realize the content was not something they liked. The system was failing to meet our user-centric goals.

To provide a better service for our users, we began to look at the amount of time a video was watched, and whether it was watched to completion, rather than just whether it was clicked. Additionally, we began to demote clickbait. We realized that watchtime was a better signal to determine whether the content we were surfacing to users was connecting them to engaging content they'd enjoy watching. But we learned that just because a user might be watching content longer does not mean that they are having a positive experience. So we introduced surveys to ask users if they were satisfied with particular recommendations. With this direct feedback, we started fine-tuning and improving these systems based on this high-fidelity notion of satisfaction.

The efforts to improve YouTube's recommendation systems did not end there. We set out to prevent our systems from serving up content that could misinform users in a harmful way, particularly in domains that rely on veracity, such as science, medicine, news, or historical events.

To that end, we introduced a higher bar for videos that are promoted through the YouTube homepage or that are surfaced to users through the "watch next" recommendations. Just because content is available on the site, it does not mean that it will display as prominently throughout the recommendation engine.

As has been mentioned previously, our business depends on the trust users place in our services to provide reliable, high-quality information. The primary goal of our recommendation systems today is to create a trusted and positive experience for our users. Ensuring these recommendation systems less frequently provide fringe or low-quality disinformation content is a top priority for the company. The YouTube company-wide goal is framed not just as "Growth", but as "Responsible Growth".

Beyond removal of content that violates our community guidelines, our teams have three explicit tactics to support responsible content consumption. They are:

- **Where possible and relevant, elevate authoritative content from trusted sources.** In areas such as music or entertainment, relevance, newness, or popularity might be better signals to tilt our systems to achieve the user's desired intent and connect them to quality content they'd enjoy. But as we describe in our Search section, in verticals where veracity and credibility are key, including news, politics, medical, and scientific domains, we work hard to ensure our search and recommendation systems provide content from more authoritative sources.
- **Provide users with more context (often text-based information) to make them more informed users on the content they consume.** On certain types of content, including content produced by organizations that receive state or public funding or topical content that tends to be accompanied by disinformation online, we have started to provide information panels that contain additional contextual information and links to authoritative third-party sites so that our users can make educated decisions about the content they watch on our platform.
- **Reduce recommendations of low-quality content.** We aim to design a system that recommends quality content while less frequently recommending content that may be close to the line created by our Community Guidelines, content that could misinform users in harmful ways, or low-quality content that may result in a poor experience for our users, like clickbait. For example, content that claims that the Earth is flat or promises a "miracle cure" for a serious disease might not necessarily violate our Community Guidelines, but we don't want to proactively recommend it to users. Our incentives are to get this right for our users, so we use everyday people as evaluators to provide input on what constitutes disinformation or borderline content under our policies, which in turn informs our ranking systems.

Case Study: Applying Our Principles to YouTube News & Politics

Disinformation in news and politics is a priority given its importance to society and the outsized impact disinformation can have during fast-moving news events. Although news content only generates a small fraction of overall YouTube watchtime, it's a specific use case that is especially important to us.

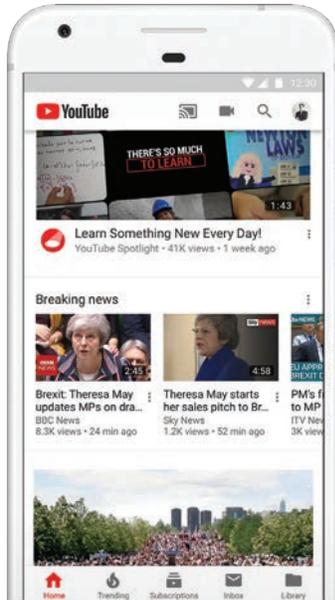
In July 2018, YouTube announced product and partnership tactics that directly apply our guiding content management principles to news.³²

The first solution included **making authoritative sources readily available**. To help meet this goal, we created a system to elevate authoritative sources for people coming to YouTube for news and politics. For example, if a user is watching content from a trusted news source, the "watch next" panel should similarly display content from other trusted news sources. Assumed within this principle is the demotion of disinformation content that we outlined earlier.

The team has also built and launched two cornerstone products – the Top News shelf and the Breaking News shelf – to prominently display authoritative political news information. The Top News shelf triggers in response to search queries that have political news-seeking intent, and provides content from verified news channels. These systems rely on a variety of signals that are derived from Google News and from our internal systems when a user might be seeking information on a given topic.



Top News Shelf on YouTube Search



Breaking News Shelf on YouTube Home

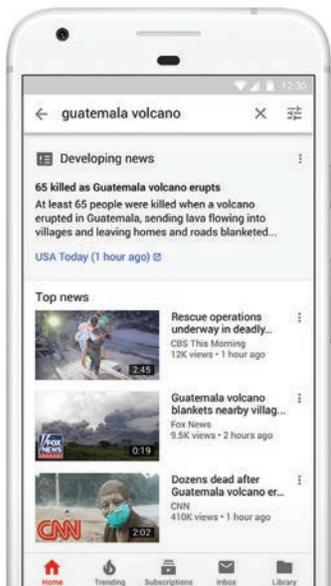
The Breaking News shelf triggers on the YouTube homepage automatically when there is a significant news event happening in a specific country.

Similar to the Top News shelf, only content from authoritative sources is eligible to be displayed in the Breaking News shelf.

More recently, YouTube has been developing products that directly address a core vulnerability involving the spread of disinformation in the immediate aftermath of a breaking news event. After a breaking news

event, it takes some time for journalists to create authoritative fact-based video content and upload it to YouTube, while unscrupulous uploaders can more quickly upload bizarre conspiracy theories.

In these events, YouTube's systems historically delivered the most relevant content that matched the typed query, and without appropriate guardrails would display content from these users seeking to exploit this vulnerability.



Information Panel Providing Breaking News Context

The first step toward a resolution involved creating systems that determine when a breaking news event might be happening, which tilts results tied to that event toward authority and away from strict relevance, popularity, or recency. This helped display content from credible sources.

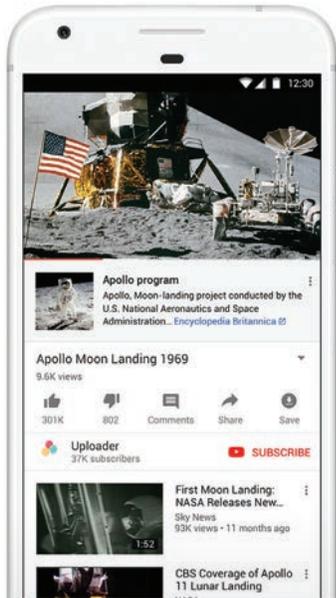
Furthermore, while authoritative video content takes time, credible text-based reporting is much quicker. As a result, YouTube launched a product that displays an information panel providing text-based breaking news content from an authoritative news source while a significant news event is developing. The information panel also links directly to the article's website so that viewers can easily access and read the full article about the news event.

Once a critical mass of authoritative news videos has been published on the topic, Breaking News and Top News shelves begin to take over as the primary news consumption experiences on the platform.



Information Panel Providing Publisher Context

Another information panel product that was launched in 2018 aims to provide additional factual information from outside sources around topics that tend to be accompanied by disinformation online, particularly on YouTube. Users may see panels alongside videos on topics such as “moon landing hoax” that link to information from credible sources, including Encyclopaedia Britannica and Wikipedia.

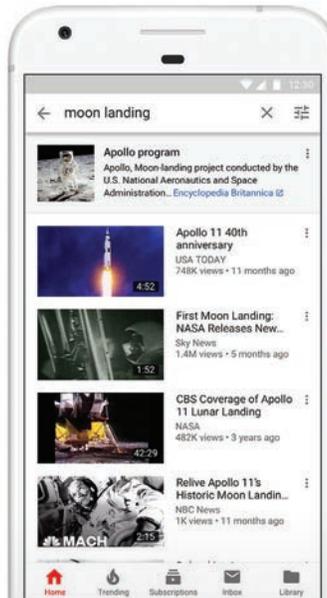


Information Panel Providing Topical Context on YouTube Search and Watch Pages

The second solution focuses on **providing context to help people make their own decisions**. There are certain instances where YouTube provides viewers with additional information to help them better understand the sources of news content they watch. For example, if a channel is owned by a news publisher that is funded by a government, or is otherwise publicly funded, an information panel providing publisher context is displayed on the watch page of the channel's videos. This information panel indicates how the publisher is funded and provides a link to the publisher's Wikipedia page.

Another information panel product that was launched in 2018 aims to provide additional factual information from outside

sources around topics that tend to be accompanied by disinformation online, particularly on YouTube. Users may see panels alongside videos on topics such as “moon landing hoax” that link to information from credible sources, including Encyclopaedia Britannica and Wikipedia.



An information panel providing topical context may appear in Search results or on the watch page of videos. It will include basic, independent information about a given topic and will link to a third-party website to allow viewers to learn more about the topic. This information panel appears alongside all videos related to the topic, regardless of the opinions or perspectives expressed in the videos. Information panels do not affect any video features or monetization eligibility.

The third solution involves **supporting journalism with technology that allows news to thrive**.

These important societal issues go beyond any single platform and involve shared values across our society. The fight against disinformation is only as good as the quality of news information available in the ecosystem, and are intent on doing our part to support the industry if we are to truly address the broader problem of disinformation. We believe quality journalism requires sustainable revenue streams and that we have a responsibility to support innovation in products and funding for news.

Several years ago, YouTube developed a solution in partnership with key publishers to help newsrooms improve and maximize their video capabilities. The program is called Player for Publishers, and it allows publishers to use YouTube to power the videos on their sites and applications. The program is free

and 100% of the advertising revenue for ads sold by the publisher and served on their own properties goes to the publisher.

In addition, last year, YouTube committed \$25M in funding, as part of a broader \$300M investment by the Google News Initiative,³³ to support news organizations in building sustainable video operations. YouTube announced³⁴ the winners of our first ever innovation funding program. These partners hail from 23 countries across the Americas, Europe, Africa, and Asia-Pacific, representing a diverse mix of broadcasters, traditional publishers, digital publishers, news agencies, local media, and creators. Best practices gained from this program will be shared publicly via case studies, providing all newsrooms the opportunity to learn and apply insights as we work together to support the development of long-term, sustainable news video businesses.



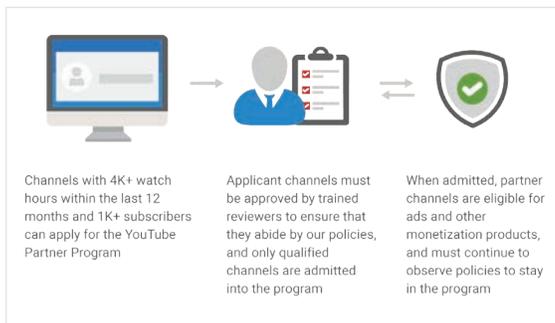
Player for Publishers

In conjunction with this investment, YouTube created a news working group – a quarterly convening of news industry leaders with whom we are collaborating to shape the future of news on YouTube. The news working group consists of top broadcasters, publishers, creators, and academics around the world and they have been providing feedback on items such as how to better quantify authoritativeness, what additional types of information might be useful to provide to users in our information panel projects, and what more we can do to support online video operations in newsrooms. Given how complex these issues are, we know we can't work in a silo. We must partner with industry and civil society to come together around solutions that work.

3. We view monetization on our platform as a privilege

Many people use YouTube simply to share their content with the world. Creators who meet the eligibility criteria can apply to join the YouTube Partner Program, which makes their videos eligible to run advertising and earn money through Google's advertising products. Monetizing creators must comply with advertiser-friendly content guidelines. Advertising will be disabled from running on videos that violate these policies.

Over the last few years, YouTube and Google's advertising products have taken steps to strengthen the requirements for monetization so that spammers, impersonators, and other bad actors can't hurt the ecosystem or take advantage of good creators. To apply for membership in the YouTube Partner Program, the thresholds were increased for channels to be deemed eligible. Channels must have generated 4,000 watch hours in the previous 12 months and have more than 1,000 subscribers.



Process for Admittance to YouTube Partner Program

Following application, YouTube's review team ensures the channel has not run afoul of monetization, content, and copyright policies prior to admitting them to the program. Only creators with sufficient history and demonstrated advertiser safety will receive access to ads and other monetization products. In changing these thresholds, YouTube has significantly improved the protections in place against impersonating creators.

More information on how Google protects its monetization services from abuse is discussed in the next section, "Google Ads & Disinformation".

Google Advertising Products & Disinformation

Background

Google provides multiple products to help content creators – website builders, video makers, and app developers – make money from doing what they love. Our advertising products enable creators to place ads within the content they make and manage the process by which they sell space in that content. In addition, we have products for advertisers to purchase that inventory across various content creators.

Google Ads and DV 360 help businesses of all sizes get their messages to the audiences they need to grow. These services are “front doors” for advertisers of all sizes to buy ads across Google’s monetization products and platforms – connecting them with billions of people finding answers on Search, watching videos on YouTube, exploring new places on Google Maps, discovering apps on Google Play, and more.

AdSense, AdMob, and Ad Manager, support content creators’ and publishers’ efforts to create and distribute their creations. We launched AdSense in 2003 to help publishers fund their content by placing relevant ads on their website. Over time, it has become a core part of our advertising products, serving more than 2 million website owners around the world.

Our ads and monetization products enable businesses of all sizes from around the world to promote a wide variety of products, services, applications, and websites on Google and across our partner sites and apps, making it possible for Internet users to discover more content they care about.

We also understand that the content of both ads and publisher sites needs to be safe and provide a positive experience for users. We aim to protect users and ensure a positive ad experience across our partner sites and apps as well as owned and operated properties like Maps and Gmail, by creating clear policies that govern what content can and cannot be monetized. When we create these policies, we think about our values and culture as a company, as well as operational, technical, and business considerations. We regularly review changes in online trends and practices, industry norms, and regulations to keep our policies up-to-date. And, we listen to our users’ feedback and concerns about the types of ads they see.

As we create new policies and update existing ones, we strive to ensure a safe and positive experience for our users. We also consider the impact that certain kinds of content will have on our advertisers and publishers. For example, some advertisers do not want their ads shown next to particular types of publisher content, and vice versa.

At the same time, we are mindful that the advertisers and publishers who use our services represent a broad range of experiences and viewpoints and we don’t want to be in the position of limiting those viewpoints or their ability to reach new audiences.

Oftentimes, these goals are in tension with one another. We aim for a balanced approach that prevents harm to our users by putting limits on the types of content we allow to be monetized without being overly restrictive, while also creating clear, enforceable, and predictable policies for advertisers and publishers.

We have a responsibility to balance the imperatives of making sure we leave room for a variety of opinions to be expressed, while preventing harmful or misrepresentative content on our advertising platforms.

Tackling disinformation on Google’s advertising products

The considerations described above influence the policies we create for advertisers and publishers, and those policies are the primary way by which our advertising platforms implement the strategies to counter disinformation that we mention in the opening section of this paper, including:

- **Counteract Malicious Actors**
 - We look for, and take action against, attempts to circumvent our policies.
- **Give Users More Context**
 - “Why this ad” labels enabling users to understand why they’re presented with a specific ad and how to change their preferences so as to alter the personalization of the ads they are shown, or to opt out of personalized ads altogether.
 - In-ad disclosures and transparency reports on election advertising, which are rolling out during elections in the U.S., Europe, and India as our starting point.

Google’s policies to tackle disinformation on our advertising platforms favor an approach that focuses on misrepresentative or harmful behavior by advertisers or publishers while avoiding judgments on the veracity of statements made about politics or current events. To that end, we have developed a number of policies designed to catch bad behaviors, including many that can be associated with disinformation campaigns.

While we do not classify content specifically as “disinformation”, we do have a number of long-standing content policies aimed at preventing deceptive or low-quality content on our platforms. These policies complement, and build on, the strategies we outline in the opening section of this paper.

Each of these policies reflect a behavior-based approach to fighting deceptive content. Rather than make a judgment on specific claims, we enforce policies against advertiser and publisher behavior that is associated with misrepresentative or harmful content.

The policies described in this document are current as of the publication of this paper but are also subject to continual refinement and improvement to take into account emerging trends and threats to ensure the integrity of our platforms and the information we are providing to partners and users.

Managing “scraped” or unoriginal content

In order to ensure a good experience for users and advertisers, we have policies for publishers that limit or disable ad serving on pages with little to no value and/or excessive advertising.³⁵ This results in a significant number of policy violations. In 2017, we blocked over 12,000 websites for “scraping,” duplicating and copying content from other sites, up from 10,000 in 2016.³⁶

Also, Google Ads does not allow advertisements that point users to landing pages with insufficient original content. This includes content that is replicated from another source without providing any additional value in the form of added content or functionality. For example, a site that consists of news articles that are scraped from other sources without adding additional commentary or value to the user would not be allowed to advertise with us.³⁷

Misrepresentation

We have long prevented ads that intend to deceive users by excluding relevant information or giving misleading information about products, services, or businesses. This includes making false statements about the advertiser's identity or qualifications, or making false claims that entice a user with an improbable result.

Our policies on misrepresentation were extended to content that is available via our monetization products (AdSense, AdMob, and Ad Manager) in 2016, and are publicly available online.³⁸

We made an additional update to our Google Ads and AdSense policies in 2018 to specifically state that it's not acceptable to direct content about politics, social issues, or matters of public concern to users in a country other than your own, if you misrepresent or conceal your country of origin or other material details about yourself or your organization.

Inappropriate content

We also have long-standing policies to disallow monetization of shocking, dangerous, or inappropriate content on our advertising platforms, the details of which are publicly available online.³⁹ This includes derogatory content, shocking or violent content, or ads that lack reasonable sensitivity toward a tragic event.

Political influence operations

As discussed in a blog post in August 2018, we have also conducted investigations into foreign influence operations on our advertising platforms.⁴⁰ To complement the work of our internal teams, we engage independent cybersecurity experts and top security consultants to provide us with intelligence on these operations. Actors engaged in these types of influence operations violate our policies and we swiftly remove such content from our services and terminate these actors' accounts.

Election integrity

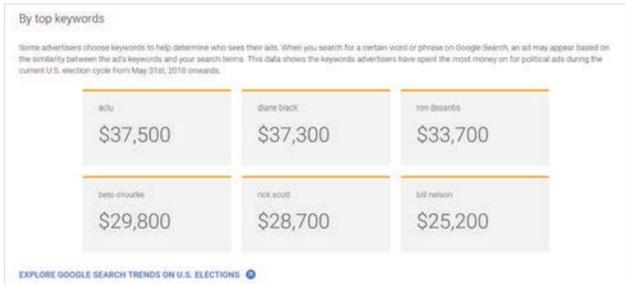
When it comes to elections, we recognize that it is critical to support democratic processes by helping users get important voting information, including insights into who is responsible for the political advertising content they see on our platforms.

Beginning with the 2018 U.S. Congressional midterm election, we require additional verification for anyone who wants to purchase an election ad on Google in the U.S., and require that advertisers confirm they are a U.S. citizen or lawful permanent resident.⁴¹ In an effort to provide transparency around who is paying for an election ad, we also require that ad creatives incorporate a clear disclosure of who is paying for it. Additionally, we released a Transparency Report specifically focused on election ads.⁴² This Report describes who is buying election-related ads on our platforms and how much money is being spent. We have also built a searchable library for election ads where anyone can find election ads purchased on Google and who paid for them.⁴³ In parallel, we updated our personalized ads policies to require verification for all advertisers who use our limited political affiliation options to target ads to users or to promote advertisers' products and services in the United States.⁴⁴

As we look ahead to 2019, we are also planning to extend these election integrity efforts to other elections around the globe. Similar to our approach for U.S. federal elections, we will be requiring verification and disclosure for election ads in the European Union Parliamentary elections⁴⁵ and the Indian Lok Sabha election.⁴⁶ Ads that mention a political party, candidate, or current office holder will be verified by Google and be required to disclose to voters who is paying for the ad. We will also be introducing respective Political Ads Transparency Reports and searchable ad libraries for each of these elections that will provide more information about who is purchasing election ads, who is being targeted, and how much money is being spent.



The new political advertising section in our Transparency Report shows how much money is spent across states and congressional districts for U.S. federal elections



The political advertising section in our U.S. Transparency Report also shows which ads had the highest views, the latest election ads running on our platform, and explores specific advertisers' campaigns

In addition to these specific efforts, we're thinking hard about elections and how we continue to support democratic processes around the world, including by bringing more transparency to political advertising online, by helping connect people to useful and relevant election-related information, and by working to protect election information online. We will continue to invest in initiatives that build further on our commitment to election transparency.

Consistent Enforcement

Our enforcement teams use a variety of robust methods to ensure content on our advertising platforms adheres to our policies, including machine learning, human review, and other technological methods. This approach is very similar to the one used by YouTube and described earlier in this paper. We have always relied on a combination of humans and technology to enforce our policies and will continue to do so.

When we find policy violations we take action to enforce our policies. Depending on the policy violation, this can include blocking a particular ad from appearing and removing ads from a publisher page or site. In cases of repeated or egregious violations, we may disable an account altogether.^{47/48} Every year we publish a report on our efforts to remove bad actors from our advertising ecosystem.⁴⁹

We also know that some content, even if it complies with our policies, may not be something that all advertisers want to be associated with. That's why, in addition to these policies, we provide advertisers with additional controls, and help them exclude certain types of content that, while in compliance with our policies, may not fit their brand or business. These controls let advertisers exclude certain types of content or terms from their video, display and search advertising campaigns. Advertisers can exclude whole categories of content such as politics, news, sports, beauty, fashion, and many others. Similarly, publishers can also review and block certain ads from showing on their pages, including by specific advertiser URL, general ad category like "apparel" or "vehicles", and sensitive ad category like "religion" or "politics".

Conclusion

Tackling the propagation of false or misleading information is core to Google's mission and to ensuring our products remain useful to the billions of users and partners who utilize our services every day. While we have always fought against malicious actors' efforts to manipulate our systems and deceive our users, it's never been more important to thwart them and to ensure we provide our users with information worthy of the trust they have in our services.

As we have outlined in this paper, this is not a straightforward endeavor. Disinformation and misinformation can take many shapes, manifest differently in different products, and raise significant challenges when it comes to balancing risks of harm to good faith, free expression, with the imperative to serve users with information they can trust.

We believe that we're at our best when we improve our products so they continue to make quality count, to counteract malicious actors, and to give users context, as well as working beyond our products to support a healthy journalistic ecosystem, partner with civil society and researchers, and stay one step ahead of future risks.

We are constantly striving to make progress on these issues. This is by no means a solved problem, and we know that we have room to make progress. We welcome a constructive dialogue with governments, civil society, academia, and newsrooms on what more can be done to address the challenges of misinformation and disinformation and hope that this paper will be useful in sparking these conversations.

References

- 1 <https://europe.googleblog.com/2015/10/introducing-accelerated-mobile-pages.html>
- 2 <https://www.blog.google/outreach-initiatives/google-news-initiative/digital-news-initiative-introducing/>
- 3 <https://g.co/newsinitiative>
- 4 <https://blog.google/topics/google-news-initiative/introducing-subscribe-google/>
- 5 <https://www.poynter.org/channels/fact-checking>
- 6 <https://jigsaw.google.com/>
- 7 <https://www.blog.google/technology/safety-security/update-state-sponsored-activity/>
- 8 <https://protectyourelection.withgoogle.com/>
- 9 <https://projectshield.withgoogle.com/>
- 10 <https://transparencyreport.google.com/political-ad>
- 11 <https://www.blog.google/around-the-globe/google-europe/supporting-european-union-parliamentary-elections/>
- 12 <https://india.googleblog.com/2019/01/bringing-more-transparency-to-indian.html>
- 13 On January 31st 2019, we have made a dataset of synthetic speech available to all participants in the third-party and independent 2019 ASVspoof challenge, which invites researchers from all around the world to test countermeasures against fake (or 'spoofed') speech. Blog post: <https://www.blog.google/outreach-initiatives/google-news-initiative/advancing-research-fake-audio-detection/>
- 14 More on the guidelines: <https://www.google.com/search/howsearchworks/mission/web-users/>
- 15 <https://www.blog.google/products/news/new-google-news-ai-meets-human-intelligence/>
- 16 E.g. for Autocomplete: <https://support.google.com/websearch/answer/7368877>
- 17 Link to Google News content policies: <https://support.google.com/news/producer/answer/6204050>
- 18 Link to our Search Quality Raters Guidelines: <https://static.googleusercontent.com/media/www.google.com/en//insidesearch/howsearchworks/assets/searchqualityevaluatorguidelines.pdf>
- 19 More information on our work to protect elections can be found in our opening section, [p.7-8](#)
- 20 <https://support.google.com/webmasters/answer/35769>
- 21 <http://infolab.stanford.edu/~backrub/google.html>
- 22 Webmasters can signal fact-check content to Google Search and Google News using dedicated HTML Code – more information about the technical and content criteria applying to these fact-checks here: <https://developers.google.com/search/docs/data-types/factcheck>
- 23 For more information: <https://twitter.com/searchliaison/status/1070027261376491520>
- 24 <https://datasociety.net/output/data-voids-where-missing-data-can-easily-be-exploited/>
- 25 In addition to these Community Guidelines, we have guidelines related to copyright, privacy, and impersonation that are not discussed in this paper. See our list of Community Guidelines here: <https://www.youtube.com/yt/about/policies/#community-guidelines>
- 26 <https://www.globalnetworkinitiative.org/>
- 27 <https://transparencyreport.google.com/youtube-policy/overview>
- 28 <https://support.google.com/youtube/answer/2801973?hl=en>
- 29 https://support.google.com/youtube/answer/2801947?hl=en&ref_topic=2803176
- 30 https://support.google.com/youtube/answer/2801939?hl=en&ref_topic=2803176
- 31 https://support.google.com/youtube/answer/2802268?hl=en&ref_topic=2803176
- 32 <https://youtube.googleblog.com/2018/07/building-better-news-experience-on.html>
- 33 <https://newsinitiative.withgoogle.com/>
- 34 <https://youtube.googleblog.com/2018/07/building-better-news-experience-on.html>
- 35 Read more: <https://support.google.com/adsense/answer/1346295?hl=en>
- 36 Read more: https://support.google.com/adsense/answer/1348737?hl=en&ref_topic=1261918
- 37 Read more: https://support.google.com/adspolicy/answer/6368661?hl=en&ref_topic=1626336
- 38 Google Ads Misrepresentation Policy: <https://support.google.com/adspolicy/answer/6020955> - AdSense misrepresentative content policy: <https://support.google.com/adsense/answer/1348688>
- 39 Google Ads Dangerous or Derogatory Policy: <https://support.google.com/adspolicy/answer/6015406>; AdSense Dangerous or Derogatory Policy: https://support.google.com/adsense/answer/1348688?hl=en&rd=1#Dangerous_or_derogatory_content; AdSense Shocking Content policy: https://support.google.com/adsense/answer/1348688?hl=en&rd=1#Shocking_content
- 40 Read the post: <https://www.blog.google/technology/safety-security/update-state-sponsored-activity/>
- 41 Read more: <https://support.google.com/adspolicy/answer/9002729>
- 42 Read more: <https://transparencyreport.google.com/political-ads/overview>
- 43 Search the library: <https://transparencyreport.google.com/political-ads/library>
- 44 Read more: <https://support.google.com/adspolicy/answer/143465?#533>
- 45 Read more: <https://www.blog.google/around-the-globe/google-europe/update-our-work-prevent-abuse-ahead-eu-elections/>
- 46 <https://india.googleblog.com/2019/01/bringing-more-transparency-to-indian.html>
- 47 More information regarding Google Ads enforcement: https://support.google.com/adspolicy/answer/7187501?hl=en&ref_topic=1308266
- 48 More information regarding AdSense enforcement: <https://support.google.com/adspolicy/answer/7003627?hl=en>
- 49 Our most recent report from March 2018: <https://blog.google/technology/ads/advertising-ecosystem-works-everyone>

Google

EXHIBIT K



Creator Blog

More updates on our actions related to the safety of minors on YouTube

Thursday, February 28, 2019

Dear Creators,

We know that many of you have been closely following the actions we're taking to protect young people on YouTube and are as deeply concerned as we are that we get this right. We want to update you on some additional changes we're making, particularly in regards to comments, building on the efforts [we shared last week](#).

We recognize that comments are a core part of the YouTube experience and how you connect with and grow your audience. At the same time, the important steps we're sharing today are critical for keeping young people safe. Thank you for your understanding and feedback as we continue our work to protect the YouTube community.

Below is a summary of the main steps we've taken to improve child safety on YouTube since our [update](#) last Friday:

Disabling comments on videos featuring minors

Over the past week, we disabled comments from tens of millions of videos that could be subject to predatory behavior. These efforts are focused on videos featuring young minors and we will continue to identify videos at risk over the next few months. Over the next few months, we will be broadening this action to suspend comments on videos featuring young minors and videos featuring older minors that could be at risk of

attracting predatory behavior.

A small number of creators will be able to keep comments enabled on these types of videos. These channels will be required to actively moderate their comments, beyond just using our moderation tools, and demonstrate a low risk of predatory behavior. We will work with them directly and our goal is to grow this number over time as our ability to catch violative comments continues to improve.

Launching a new comments classifier

While we have been removing hundreds of millions of comments for violating our policies, we had been working on an even more effective classifier, that will identify and remove predatory comments. This classifier does not affect the monetization of your video. We accelerated its launch and now have a new comments classifier in place that is more sweeping in scope, and will detect and remove 2X more individual comments.

Taking action on creators who cause egregious harm to the community

No form of content that endangers minors is acceptable on YouTube, which is why we have terminated certain channels that attempt to endanger children in any way. Videos encouraging harmful and dangerous challenges targeting any audience are also clearly against our policies. We will continue to take action when creators violate our policies in ways that blatantly harm the broader user and creator community. Please continue to [flag](#) these to us.

Thank you for your understanding as we make these changes,

TeamYouTube



SANFORD J. LEWIS, ATTORNEY

Via electronic mail

March 11, 2019

Office of Chief Counsel
Division of Corporation Finance
U.S. Securities and Exchange Commission
100 F Street, N.E.
Washington, D.C. 20549

Re: Shareholder Proposal to Alphabet on Behalf of The New York State Common Retirement Fund and Others

Ladies and Gentlemen:

The New York State Common Retirement Fund (the “Proponent”) is beneficial owner of common stock of Alphabet Inc. (the “Company”) and has submitted a shareholder proposal (the “Proposal”) to the Company together with co-lead-filer Natasha Lamb, of Arjuna Capital, on behalf of Lisa Stephanie Myrkalo and Andrea Louise Dixon. I have been asked by the Proponent to respond to the letter dated February 5, 2019 (“Company Letter”) sent to the Securities and Exchange Commission by Pamela L. Marcogliese, of Clearly Gottlieb Steen & Hamilton LLP. In that letter, the Company contends that the Proposal may be excluded from the Company’s 2019 proxy statement. A copy of this letter is being emailed concurrently to Pamela L. Marcogliese.

I have reviewed the Proposal, as well as the letter sent by the Company, and based upon the foregoing, as well as the relevant rules, it is my opinion that the Proposal must be included in the Company’s 2019 proxy materials and that it is not excludable under Rule 14a-8. A copy of this letter is being emailed concurrently to Ms. Marcogliese.

SUMMARY

The Company Letter asserts that the Proposal is excludable under Rule 14a-8(i)(10) as substantially implemented. However, the Company Letter asserts this argument against the wrong proposal, focusing on an outdated proposal which focused solely on election interference, rather than the duly amended content management-related Proposal, that focuses on hate speech and freedom of expression in addition to election interference.

The Company Letter documents increasing transparency and accountability of the Company on the issue of election interference. However, the Company has not fulfilled the request of the Proposal to report to shareholders on the efficacy of its enforcement of Google’s terms of service

related to content policies and assessing the risks posed by content management controversies related to election interference, freedom of expression, and the spread of hate speech, to the company's finances, operations, and reputation.

In particular, the whereas clauses of the Proposal demonstrate a significant focus on issues of *hate speech and freedom of expression*, including how "Google's YouTube continues to provide a home for extremist content" and how the Company has allowed racist, anti-Semitic, and conspiratorial content to remain on its platforms. The Company's existing disclosures do not assess the efficacy of its content policies to address these issues.

Therefore, the Proposal is not substantially implemented and is not excludable pursuant to Rule 14a-8(i)(10).

THE PROPOSAL

Report on Content Governance

WHEREAS, shareholders are concerned that Alphabet's Google is failing to effectively address content governance concerns, posing risks to shareholder value. These concerns extend globally across multiple Google platforms and products.

Google's attempts thus far to address content governance have misfired. For example, in October 2017, Google acknowledged its automated system incorrectly flagged Google Docs content as violating its terms of service, blocking user-generated content and inconveniencing users. In June 2018, YouTube apologized for repeatedly filtering or demonetizing content created by LGBTQ users.

Google's YouTube continues to provide a home for extremist content. In June 2018, U.S. news outlets reported that white supremacist and white nationalist content "found a home" on social network Google Plus, violating the user policy. Data & Society research institute says: "The [YouTube] platform, and its parent company, have allowed racist, misogynist, and harassing content to remain online — and in many cases, to generate advertising revenue — as long as it does not explicitly include slurs." The Network Contagion Research Institute, a group tracking the spread of hate speech, found that the man charged in a mass shooting at a Pittsburgh synagogue linked to racist and anti-Semitic YouTube videos 71 times.

In December 2018, the Washington Post reported: "A year after YouTube's chief executive promised to curb 'problematic' videos, it continues to harbor and even recommend hateful, conspiratorial videos, allowing racists, anti-Semites and proponents of other extremist views to use the platform as an online library for spreading their ideas.... The struggle to control the spread of such content poses ethical and political challenges to YouTube and its embattled parent company, Google. "These controversies have drawn regulatory scrutiny. The European Union,

for example, announced measures intended to pressure Google and other companies to combat disinformation ahead of EU parliament elections in May 2019. Security Commissioner Julian King said: “No excuses, no more foot-dragging, because the risks are real.... They’ve got to get serious about this stuff.” CEO Sundar Pichai was summoned to Congress to respond to questions about the spread of conspiracy theories on YouTube, during which he said: “We are constantly undertaking efforts to deal with misinformation ... we are looking to do more ... it’s an area we acknowledge there’s more work to be done.”

Shareholders are concerned that Google’s inability to address these issues proactively poses substantial regulatory, legal, and reputational risks to long-term value.

RESOLVED: Shareholders request Alphabet Inc. issue a report to shareholders at reasonable cost, omitting proprietary or legally privileged information, reviewing the efficacy of its enforcement of Google’s terms of service related to content policies and assessing the risks posed by content management controversies related to election interference, freedom of expression, and the spread of hate speech, to the company’s finances, operations, and reputation.

SUPPORTING STATEMENT: Proponents recommend the report include assessment of the scope of platform abuses and address related ethical concerns.

ANALYSIS

The Company has not substantially implemented the Proposal, and therefore the Proposal is not excludable under Rule 14a-8(i)(10).

In order for a Company to meet its burden of proving substantial implementation pursuant to Rule 14a-8(i)(10), the actions in question must compare favorably with the guidelines and essential purpose of the Proposal. The Staff has noted that a determination that a company has substantially implemented a proposal depends upon whether a company’s particular policies, practices, and procedures compare favorably with the guidelines of the proposal. *Texaco, Inc.* (Mar. 28, 1991). Substantial implementation under Rule 14a-8(i)(10) requires a company’s actions to have satisfactorily addressed *both* the proposal’s guidelines and its essential objective. See, e.g., *Exelon Corp.* (Feb. 26, 2010). Thus, when a company can demonstrate that it has already taken actions that meet most of the guidelines of a proposal and meet the proposal’s essential purpose, the Staff has concurred that the proposal has been “substantially implemented.” In the current instance, the Company has substantially fulfilled *neither* the guidelines nor the essential purpose of the Proposal, and therefore the Proposal cannot be excluded.

Identifying the correct proposal

The Proponent, the New York State Common Retirement Fund, initially submitted a proposal on

December 6, 2018. The Proponent then sent an amended version of the Proposal on December 20. (Appendix A) The co-lead-filer, Arjuna Capital on behalf of Lisa Stephanie Myrkalo and Andrea Louise Dixon submitted the same amended Proposal on December 21, 2018. (Appendix B) The amended Proposal was an updated version reflecting developments during the year and included discussion of freedom of expression and hate speech as well as election interference. The amended Proposal is included in the no-action request submitted by the Company under the cover letter from the co-lead-filer, but the Company Letter's arguments refer to the Proponent's original Proposal – filed on December 6, 2018 ignoring the fact that it was subsequently amended when the later version was filed on December 20, 2019. Enclosed as Appendix A is documentation that the amended Proposal was received by the Company on December 26, 2018, two days before the filing deadline of December 28, 2019. Accordingly, it is the amended Proposal that is at issue in the current no-action request.

Assessing the Proposal's purpose and guidelines

The guidelines of the amended Proposal request that Alphabet Inc. issue a report to shareholders at reasonable cost, omitting proprietary or legally privileged information, reviewing the efficacy of its enforcement of Google's terms of service related to content policies and assessing the risks posed by content governance controversies related to election interference, freedom of expression, and the spread of hate speech, to the company's finances, operations, and reputation. The supporting statement recommends the report include assessment of the scope of platform abuses and address related ethical concerns.

The essential purpose of the Proposal is to address the Company's failure to effectively address content governance concerns, posing a risk to shareholder value.

In this regard, as documented in the Company Letter, the Company has been confronted publicly for its poor content governance in relation to election interference, but has been far less transparent about its other major content governance issues — hate speech and freedom of expression — that are featured in the current Proposal.

Due to the exploitation of social media platforms by a range of malign actors, from fake conspiracy theorists and pedophiles to white supremacists and foreign terrorist organizations, a company like Google, a unit of Alphabet, whose core business involves operating social media platforms, faces substantial risk associated with unresolved issues regarding content governance.

Among internet service companies and social media sites, the Google search engine and YouTube stand out as two leading global platforms. As a result, both of these platforms are facing substantial challenges balancing users' freedom of expression with the need not to amplify harmful content: hate speech, online harassment and other manipulative and dangerous behavior. Over time, the ability of malign users to exploit the unconstrained nature of social media has intensified the need to address public safety and ensure the sustainability of the business model.

Under Staff precedents,¹ the background section of a proposal as well as the resolved clause is relevant to assessing substantial implementation — defining the essential purpose of the proposal as well as the guidelines. In the present instance, the majority of the background sections of the Proposal discuss issues of extremist content on the Company’s platform.

Background on the Company

Alphabet is the parent company of Google, whose segment revenues make up 98.9% of the Company’s total revenue, per Alphabet’s Form 10-K for 2017.² Google’s dominance within the Company and in the wider technology industry through its search engine and social media platforms amplifies the challenges posed by these issues.

Google essentially dominates the process of searching on the internet, as the top search engine in the U.S. It generates 63.1% of all core search queries, and 93% market share for mobile searches.³ As such, malign actors frequently try to game the searching process, to have their messages elevated in searches.

Yet the Company continues to struggle with managing the oversight and disclosure of violations of its terms of service, and balancing freedom of expression with terms of service abuses, including hate speech, speech related to election interference, or content suggestive of pedophilia. Recent developments, some in the last few months, demonstrate the severity of the types of speech under the Company’s purview, and the degree to which oversight and disclosure is lacking.

Evidence of violence stemming from use of the Company’s online platform

The Company’s content governance challenges go far beyond election interference, and now include preventing violence. Many social media platforms are staffing up significantly and striving to rapidly develop algorithms that can detect and intercept malign actors on their platforms in order to prevent harm to users or the public. Meanwhile, Google struggles to govern content that may contribute to real-world violence.

Evidence of violence stemming from the promotion of propaganda on YouTube includes the Pizzagate conspiracy. Pizzagate, a conspiracy that circulated in YouTube videos before the 2016 U.S. presidential election, was “predicated on the idea that emails leaked by Wikileaks from Clinton campaign chairman John Podesta reveal a secret pedophile cabal,” which ultimately led to a shooting in Washington, D.C., restaurant Comet Ping Pong. The conspiracy theory falsely

¹ For instance, see *Lowe’s Companies, Inc.* (March 21, 2006).

² <https://www.sec.gov/Archives/edgar/data/1652044/000165204418000007/goog10-kq42017.htm>.

³ “Share of search queries handled by leading U.S. search engine providers as of October 2018”.
<https://www.statista.com/statistics/267161/market-share-of-search-engines-in-the-united-states/>.

claimed that a celebrity- and Democrat-run sex ring operated out of its nonexistent basement. Edgar M. Welch, a conspiracy theorist, armed himself with an assault rifle and demanded to see the child sex ring, firing his rifle several times.⁴

A similar conspiracy theory has led a YouTube user allegedly set fire to a restaurant in an attempt to burn it down. The tie from this alleged arson to YouTube is clear:⁵

A video posted to [this user's] parents' YouTube account the night of the fire seems to provide a possible link between the alleged arson and the disturbing conspiracy theory that became popular among a fringe group of Trump supporters during the 2016 election, and inspired a shooting in the restaurant that same year.

...One hour later, at 9:17 p.m., federal prosecutors allege [this user] attempted to set fire to Comet Ping Pong, where pizzagate conspiracy theorists falsely believe the global pedophile ring is partially located.

These examples show that although YouTube's conspiracy theory videos seem unbelievable to most people, they can lead some individuals to cause real-world harm. This extraordinary problem facing social media companies only threatens to get worse as the users finding ways of exploiting the platforms play cat and mouse with the Company.

Other conspiracy videos on YouTube

Developments revealing the multitude of conspiracy videos proliferating on YouTube demonstrate the Company's struggle to curb abuses of its terms of use while allowing for freedom of expression. Indeed, the *Atlantic* magazine observed:⁶

The deeper argument that YouTube is making is that conspiracy videos on the platform are just a kind of mistake. But the conspiratorial mind-set is threaded through the social fabric of YouTube. In fact, it's intrinsic to the production economy of the site.

The Atlantic also reported that despite YouTube stating it would stop recommending "content that could misinform users in harmful ways—such as videos promoting a phony miracle cure for a serious illness, claiming the earth is flat, or making blatantly false claims about historic events like 9/11":

⁴ Brandy Zadrozny and Ben Collins, "'Pizzagate' video was posted to YouTube account of alleged arsonist's parents before fire," NBC News, February 14, 2019. <https://www.nbcnews.com/tech/social-media/pizzagate-conspiracy-video-posted-youtube-account-alleged-arsonist-s-parents-n971891>.

⁵ Ibid.

⁶ Alexis C. Madrigal, "The Reason Conspiracy Videos Work So Well on YouTube," *The Atlantic*, February 21, 2019. <https://www.theatlantic.com/technology/archive/2019/02/reason-conspiracy-videos-work-so-well-youtube/583282/>.

... **the conspiracy videos continue to burble up in the great seething mass of moving pictures.** Earlier this week, in a report on the continued success of conspiracy videos on the platform, The New York Times' Kevin Roose observed, "Many young people have absorbed a YouTube-centric worldview, including rejecting mainstream information sources in favor of platform-native creators bearing 'secret histories' and faux-authoritative explanations."

An example of one of the conspiracies which used YouTube as a platform to gain more exposure involved false information regarding the state of Supreme Court Justice Ruth Bader Ginsburg's health. In January 2019, the Washington Post reported:⁷

Conspiracy theories about the health of Supreme Court Justice Ruth Bader Ginsburg have dominated YouTube this week, illustrating how the world's most popular video site is failing to prevent its algorithm from helping popularize viral hoaxes and misinformation....

The circulation of these conspiracy videos is not victimless. The BBC reports "Conspiracy theories around mass shootings have had real life consequences":⁸

David Neiwert, the author of *Alt-America: The rise of the radical right in the age of Trump*, believes that conspiracy theories are a key component of the online radicalisation process, and that they can even lead to more murders and more mass shootings.

"The person who most embodies this process on the alt right is Dylann Roof, the young man who walked into that black church in June of 2015 and murdered nine people," Neiwert says.

The spread of radicalizing information through Google's websites is not limited to YouTube. Teaching Tolerance Magazine describes how Dylann Roof's radicalization began with a Google search:⁹

... the initial step that led Roof to murder nine people was not extreme: He took a curiosity and turned it into a Google search.

⁷ Tony Romm, "Searching for news on RBG? YouTube offered conspiracy theories about the Supreme Court justice instead," *The Washington Post*, January 11, 2019.

https://www.washingtonpost.com/technology/2019/01/11/searching-news-rbg-youtube-offered-conspiracy-theories-about-supreme-court-justice-instead/?utm_term=.649e3bb9dd67.

⁸ Chris Bell, "People who think mass shootings are staged," *BBC News*, February 2, 2018.

<https://www.bbc.com/news/blogs-trending-42187105>.

⁹ Cory Collins, "The Miseducation of Dylann Roof," *Teaching Tolerance*, Issue 57, Fall 2017.

<https://www.tolerance.org/magazine/fall-2017/the-miseducation-of-dylann-roof>.

Roof's curiosity was first piqued by the trial of Trayvon Martin's killer, George Zimmerman. He searched for the case he kept hearing about on the news and, after reading a Wikipedia article, determined that Zimmerman was "in the right" to see Martin as a threat.

Roof then typed "black on White crime" into the search engine, hit enter and fell into a wormhole. Top results sent him to the website for the Council of Conservative Citizens, which offered page after page featuring what Roof referred to as "brutal black on White murders."

Google presented Roof with well-packaged propaganda—misinformation published by a group with a respectable-sounding name and a history of racist messaging, a group that once referred to black people as a "retrograde species of humanity."

In his manifesto, Roof claimed he has "never been the same since that day."

From that point on, he immersed himself in white supremacist websites, as both reader and participant, honing a philosophy far removed from his upbringing—one that would inform his manifesto and fuel a mass murder.

YouTube: abused as a platform for racist and antisemitic messaging

Beyond conspiracy theories, YouTube channels are being used to proliferate extremist views, including, for example, racist white nationalist views, often involving hate-filled speech in seeming violation of YouTube's terms of service. Recent events involving the most popular users on YouTube have demonstrated the Company is not managing content effectively.

The Daily Beast reported on a study conducted to examine this issue more broadly:¹⁰

The study tracked 65 YouTubers—some of them openly alt-right or white nationalist, others who claim to be simply libertarians, and most of whom have voiced anti-progressive views—as they collaborated across YouTube channels.

The result, the study found, is an ecosystem in which a person searching for video game reviews can quickly find themselves watching a four-hour conversation with white nationalist Richard Spencer.

On YouTube, the most popular user has nearly 80 million subscribers — a majority of which are younger than 24 years old, and are known for their aggressive loyalty. This influential user

¹⁰ Kelly Weill, "Nesting Dolls of Hate: Inside YouTube's Far-Right Radicalization Factory," *The Daily Beast*, September 18, 2018. <https://www.thedailybeast.com/inside-youtubes-far-right-radicalization-factory>. Original study at <https://datasociety.net/output/alternative-influence/>

recommended a YouTube channel called “E;R”, whose creator refers to his reputation as a racist. Vox reports that “a more-than-cursory dive into the channel would have revealed several instances of disturbing imagery, racial slurs, and white supremacist messaging”, under the guise of pop culture commentary.¹¹ In addition, this most popular user was reported to have:

... referred to several past incidents that sparked a similar outcry: a video in which he performed a Nazi “heil” salute, and one in which he hired a pair of performers from a freelancer website to hold up a sign reading “Death to all Jews.” He said these examples were satirical, but many observers condemned them as antisemitic.

Pedophilia related content

After reports surfaced that pedophiles were leaving pedophilic messages and links in the comments section of innocuous videos of children, YouTube announced on February 28, it would turn off the ability to comment on all videos featuring kids. Paul Verna, an eMarketer analyst observed:¹²

YouTube, like Facebook, Twitter and other sites that allow user publishing, have faced increasing calls to monitor what appears on their sites and get rid of unsuitable content. The companies all say they have taken action to protect users. But issues keep popping up.

Concerns about YouTube comments weren’t even a top priority for advertisers and viewers a couple weeks ago, Verna said.

“It just makes you wonder, what’s the next thing that going to happen?”

The Company struggles with how to manage its content recommendation algorithms

The Company’s use of algorithms to provide users with recommendations, is both the heart of its functionality and the heart of the problem. Algorithms are used to suggest content to users, and are susceptible to exploitation by those seeking to gain viewership, including malign actors. The algorithms being used by Google search and YouTube therefore reveal an ongoing vulnerability to the Company resulting from the ever-changing search results generated and accessed by users, as well as an inability to properly manage this content.

Verge, a leading tech industry publication, observed:¹³

¹¹ Aja Romano, “YouTubes most popular user amplified anti-Semitic rhetoric. Again,” *Vox*, December 13, 2018. <https://www.vox.com/2018/12/13/18136253/pewdiepie-vs-tseries-links-to-white-supremacist-alt-right-redpill>

¹² Rachel Lerman, “YouTube suspends comments on videos of kids,” *Daily Hampshire Gazette*, February 28, 2019. <https://www.gazettenet.com/YouTube-suspends-comments-on-videos-of-kids-23823547>.

¹³ Julia Alexander, “YouTube still can’t stop child predators in its comments,” *The Verge*, February 19, 2019.

...(the) heart of the problem is YouTube’s recommendation algorithm, a system that has been widely criticized in the past. It only took two clicks...to venture away from a video of a woman showcasing bikinis she’s purchased to a video of a young girl playing. Although the video is innocent, the comments below — which include timestamps calling out certain angles in the video and predatory responses to the images — certainly aren’t.

...While individual videos are removed, the problematic users are rarely banned, leaving them free to upload more videos in the future.

In reporting on the same incidents, Newsweek concluded:¹⁴

Having content that’s algorithmically decided based on your personal habits is still a new industry, with new problems. YouTube is becoming a place that can be exploited by those with bad intentions, even as the company actively takes measures against them.

Seemingly innocuous Google searches such as those related to shopping, or the news, can easily direct users to hate speech content, revealing that hate speech remains a serious, persistent and dangerous problem on YouTube. In January 2019, BuzzFeed reported:¹⁵

Despite year-old promises to fix its “Up Next” content recommendation system, YouTube is still suggesting conspiracy videos, ... misogynist videos, pirated videos, and content from hate groups following common news-related searches.

...The test revealed that only NINE clicks through the “Up Next” recommendations would take users “from an anodyne PBS clip about the 116th United States Congress to an anti-immigrant video from a designated hate organization.”

Another report found:

The far-right has used YouTube’s algorithms to its advantage, as

<https://www.theverge.com/2019/2/19/18229938/youtube-child-exploitation-recommendation-algorithm-predators>.

¹⁴ Steven Asarch, “Youtube Struggles With Child Porn Moderation: Bans Channels by Mistake While Letting Offenders Slide,” *Newsweek*, February 18, 2019. <https://www.newsweek.com/youtube-cp-algorithm-ban-mattwhatitis-algorithm-1334873>.

¹⁵ Caroline O’Donovan, Charlie Warzel, Logan McDonald, Brian Clifton, and Max Wolf, “We Followed YouTube’s Recommendation Algorithm Down the Rabbit Hole,” *BuzzFeed*, January 24, 2019. <https://www.buzzfeednews.com/article/carolineodonovan/down-youtubes-recommendation-rabbithole>.

demonstrated by a report on a related study. Becca Lewis, the researcher behind the report, described the relatively quick succession of clicks involved for a YouTube follower to “find themselves clicking through a series of increasingly extremist videos.”¹⁶

Meanwhile, the Company has not demonstrated competence in managing the problems that have been identified regarding its algorithms. USA Today observed how the algorithms used by the Company are falling short of preventing hateful messages:¹⁷

Measures such as hiring thousands of moderators and training artificial intelligence software to root out online hate and abuse have not yet solved the problem. Algorithms still struggle to accurately interpret the meaning and intent of social media posts while moderators, when reviewing posts, frequently stumble, too, missing important cultural cues and context.

In attempting to manage issues as they arise, the Company has mistakenly allowed problem users to continue to use its service, while it terminates others’ use by mistake. Newsweek reported:¹⁸

Nearly a dozen YouTubers’ channels were wrongly terminated because every one of them has included in the title or tags (YouTube keywords) the term “CP,” which the numbers and code inside YouTube’s algorithm seemed to have mistaken for “child porn.”

In a statement to Newsweek, a YouTube spokesperson said “with the massive volume of videos on our site, sometimes we make the wrong call. When it’s brought to our attention that a video or channel has been removed mistakenly, we act quickly to reinstate it. We give uploaders the ability to appeal these decisions and we will re-review the videos.”

While these channels struggled with the algorithm, others managed to abuse it to share what YouTube has been trying so hard to combat.

The above demonstrates the larger issue regarding how the Company continues to grapple with creating a platform for its users when those users are composed of many who follow its terms of service, and a disturbing number who don’t. This is illustrated by the enormous presence of one

¹⁶ Kelly Weill, “Nesting Dolls of Hate: Inside YouTube’s Far-Right Radicalization Factory,” *The Daily Beast*, September 18, 2018. <https://www.thedailybeast.com/inside-youtubes-far-right-radicalization-factory>.

¹⁷ Jessica Guynn, “If you’ve been harassed online, you’re not alone. More than half of Americans say they’ve experienced hate,” *USA Today*, February 13, 2019. <https://www.usatoday.com/story/news/2019/02/13/study-most-americans-have-been-targeted-hateful-speech-online/2846987002/>.

¹⁸ Steven Asarch, “YouTube Struggles With Child Porn Moderation: Bans Channels By Mistake While Letting Offenders Slide,” *Newsweek*, February 18, 2019. <https://www.newsweek.com/youtube-cp-algorithm-ban-mattswatit-is-algorithm-1334873>.

particular channel on YouTube.¹⁹

[This user's] massive popularity has given him considerable influence over the future of You Tube. In fact, his channel currently sits directly at the center of what seems to be a growing divide between two very different directions for an increasingly polarized platform.

On one side lies many overlapping subcultures that make up huge swaths of the You Tube population: its tremendous gaming communities, and its increasingly insidious alt-right presence.

On the other side lie many, many YouTube users who visit the site for other reasons and other forms of entertainment, and who arguably aren't interested in supporting the cult of personalities that might be said to represent "old-school" YouTube. Instead, they come to the site for music, memes, narrative media, instructional videos, and more general forms of content consumption and entertainment.

The Company's sudden and extreme response, operating from a position of crisis management in deciding to suspend its YouTube comments further demonstrates how the Company is grappling with an increasingly difficult balancing act. On the one hand, as YouTube CEO Susan Wojcicki tweeted, "Nothing is more important to us than ensuring the safety of young people on the platform."²⁰ But on the other, turning off comments involves a level of control over speech, affecting the experience of the Company's users and video creators.

The Company's challenges in governing content on its service is affecting its brand and appeal to advertisers

Reactive decisions such as the Company's move to suspend YouTube comments demonstrate its weakness in assessing content, and generate a negative impact on the Company's reputation, as well as its appeal to advertisers. The Company recognizes the importance of maintaining its reputation and attracting advertisers to its platform. Under the heading "How we make money" in its 10K, the Company explains it generates revenues "primarily by delivering both performance advertising and brand advertising." At the same time, the Company states:²¹

Our business depends on strong brands, and failing to maintain and enhance our brands would hurt our ability to expand our base of users,

¹⁹ Aja Romano, "YouTubes most popular user amplified anti-Semitic rhetoric. Again," *Vox*, December 13, 2018. <https://www.vox.com/2018/12/13/18136253/pewdiepie-vs-tseries-links-to-white-supremacist-alt-right-redpill>.

²⁰ Rachel Lerman, "YouTube suspends comments on videos of kids," *Daily Hampshire Gazette*, February 28, 2019. <https://www.gazettenet.com/YouTube-suspends-comments-on-videos-of-kids-23823547>.

²¹ <https://www.sec.gov/Archives/edgar/data/1652044/000165204418000007/goog10-kq42017.htm>.

advertisers, content providers, and other partners.

For example, if we fail to appropriately respond to the sharing of objectionable content on our services or objectionable practices by advertisers, or to otherwise adequately address user concerns, our users may lose confidence in our brands. Our brands may also be negatively affected by the use of our products or services to disseminate information that is deemed to be misleading.

Furthermore, if we fail to maintain and enhance equity in the Google brand, our business, operating results, and financial condition may be materially and adversely affected.²²

In recent weeks, even before the Company announced the suspension of its comments section on YouTube videos, YouTube was already confronting suspensions of advertising budgets by major global brands. These resulted from a media report that:

...demonstrated how a search for something like “bikini haul,” a subgenre of video where women show various bikinis they’ve purchased, can lead to disturbing and exploitative videos of children. The videos aren’t pornographic in nature, but the comment sections are full of people time stamping specific scenes that sexualize the child or children in the video. Comments about how beautiful young girls are also litter the comment section.

As a result of that report, the Company lost advertising-based revenue, as large companies, including AT&T, Walt Disney Company, and Nestle, suspended their advertising with the Company. According to Reuters, AT&T Inc. pulled all its advertising from YouTube for the second time in two years.²³

“Until Google can protect our brand from offensive content of any kind, we are removing all advertising from YouTube,” an AT&T spokesman said in a statement on Thursday.

The move comes just one month after the U.S. wireless carrier announced it would resume buying advertising on YouTube, after a nearly two-year boycott of the platform. The previous boycott was also due to concerns that its ads could run on videos featuring hate speech or other disturbing material.

²² <https://www.sec.gov/Archives/edgar/data/1652044/000165204418000007/goog10-kq42017.htm>.

²³ Sheila Dang, “AT&T pulls ads from YouTube over videos exploiting children,” *Reuters*, February 21, 2019. <https://www.reuters.com/article/us-at-t-advertising-youtube/att-pulls-ads-from-youtube-over-videos-exploiting-children-idUSKCN1QA2L8>.

In addition, Bloomberg reported that Walt Disney Company, all the U.S. units of Nestle, and multiple other advertisers had also suspended advertising on YouTube due to reports that “comments on Google’s video site were being used to facilitate a ‘soft-core pedophilia ring.’”²⁴

Evaluation of Google’s content governance by independent researchers questions adequacy of oversight

Independent researchers have similarly found that there is no evidence of the Company’s oversight of issues raised in the Proposal, namely freedom of expression and user’s rights in how they use the Company’s service. Researchers from Ranking Digital Rights, an organization that monitors disclosures by leading tech companies regarding issues such as privacy and freedom of expression, concluded in their most recent Corporate Accountability Index:²⁵

While it articulated a clear commitment to uphold users’ freedom of expression and privacy rights, Google did not disclose evidence of board- or executive-level oversight over these issues. The company committed to conduct human rights due diligence when entering new markets, but researchers were not able to locate evidence that it conducts assessments of risks associated with the processes and mechanisms used to enforce its terms of service.

Regulatory and legislative concerns raised regarding Google’s content governance

The unending stream of content-management events, of which some are discussed above, demonstrate the Company’s challenge with governing content while maintaining a balance between freedom of expression and controlling certain types of content. This challenge has thrust Google in the spotlight of the legislative arena. Google CEO Sundar Pichai was summoned to testify before Congress in December 2018. According to one media report:²⁶

From Pizzagate to QAnon, YouTube has a serious problem with conspiracy theories. The basic moderation problem has splintered into a number of different scandals over the past two years, including disturbing children’s content, terrorism videos, white supremacy dog whistling, and radicalization via YouTube’s algorithm. ***But when confronted on those issues at a House Judiciary hearing today, Pichai offered the same***

²⁴ Mark Bergen, Gerrit De Vynck, and Christopher Palmeri, “Nestle, Disney Pull YouTube Ads, Joining Furor Over Child Videos,” *Bloomberg*, February 20, 2019. <https://rankingdigitalrights.org/index2018/companies/google/ps://www.bloomberg.com/news/articles/2019-02-20/disney-pulls-youtube-ads-amid-concerns-over-child-video-voyeurs>.

²⁵ <https://rankingdigitalrights.org/index2018/companies/google/>.

²⁶ Julia Alexander, “Google still has no answers for YouTube’s biggest problem,” *The Verge*, December 11, 2018. <https://www.theverge.com/2018/12/11/18136525/google-ceo-sundar-pichai-youtube-moderation-hearing-house-judiciary-committee>.

response that YouTube CEO Susan Wojcicki has offered in the past: there is no immediate cure.

The most vigorous questions came from Rep. Jamie Raskin (D-MD), who confronted Pichai over a Washington Post report on conspiracy videos that plague YouTube. These videos, which he summarized as “videos claiming politicians, celebrities, and other elite figures were sexually abusing or consuming the remains of children,” are part of a conspiracy theory that suggests Hillary Clinton is killing young girls in satanic rituals.

“Is your basic position that this is something you want to try to do something about,” Raskin asked, “but there’s just an avalanche of such material, and there’s nothing really to be done, so it should just be buyer beware when you go on YouTube?”

Pichai didn’t endorse that position exactly, but he didn’t give much reason to expect improvement either. “This is an area we acknowledge there’s more work to be done,” the Google CEO told Raskin. “We have to look at it on a video-by-video basis, and we have clearly stated policies, so we’d have to look at whether a specific video violates those policies.”

Pichai added that YouTube and Google have “clear policies against hate speech,” which includes “things that could inspire harm or violence,” but he added that it isn’t enough.

CONCLUSION

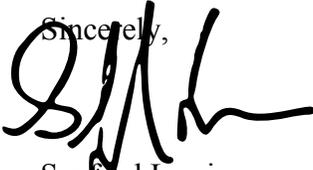
It is apparent that no responsive actions by the Company substantially implement the Proposal. The Company has not reviewed the efficacy of its enforcement of Google’s terms of service related to content policies and assessing the risks posed by content management controversies related to freedom of expression, and the spread of hate speech — the issues highlighted in the amended Proposal — to the Company’s finances, operations, and reputation.

The quandary associated with managing the larger tensions between users’ freedom of expression and the prevention of harm needs a clearly articulated strategy and analysis, from the perspective of investors. The Company’s new transparency on the issue of election interference was a first step but by no means encompasses implementation of the issues raised in this resolution. In the absence of the disclosures requested by the Proposal, it has not been substantially implemented. We are concerned that the Company could fall behind on these issues, and therefore are calling for robust discussion and analysis by board and management.

Office of Chief Counsel
March 11, 2019
Page 16

Accordingly, the Company has not substantially implemented the Proposal, and we urge the Staff to instruct the Company that the Proposal must appear on the 2019 proxy statement.

Sincerely,

A handwritten signature in black ink, appearing to be 'SML', written over the word 'Sincerely,'.

Sanford Lewis

Cc: Pamela L. Marcogliese

APPENDIX A
Amended Proposal as Submitted
by the Proponent

THOMAS P. DINAPOLI
STATE COMPTROLLER



STATE OF NEW YORK
OFFICE OF THE STATE COMPTROLLER

DIVISION OF CORPORATE GOVERNANCE
59 Maiden Lane-30th Floor
New York, NY 10038
Tel: (212) 383-3931
Fax: (212) 681-4468

December 20, 2018

Mr. David C. Drummond
Corporate Secretary
Alphabet Inc.
1600 Amphitheatre Parkway
Mountain View, California 94043

Dear Mr. Drummond:

The Comptroller of the State of New York, Thomas P. DiNapoli, is the trustee of the New York State Common Retirement Fund (the "Fund") and the administrative head of the New York State and Local Retirement System. The Comptroller has authorized me to inform of his intention to offer the enclosed shareholder proposal for consideration of stockholders at the next annual meeting. The enclosed proposal amends the Fund's December 6, 2018, submission.

I submit the enclosed proposal to you in accordance with rule 14a-8 of the Securities Exchange Act of 1934 and ask that it be included in your proxy statement.

A letter from J.P. Morgan Chase, the Fund's custodial bank verifying the Fund's ownership of Alphabet Inc. shares, continually for over one year, is enclosed. The Fund intends to continue to hold at least \$2,000 worth of these securities through the date of the annual meeting.

We would be happy to discuss this initiative with you. Should Alphabet decide to endorse its provisions as company policy, the Comptroller will ask that the proposal be withdrawn from consideration at the annual meeting. Please feel free to contact me at (212) 383-1428 and/or email at pdoherty@osc.ny.gov should you have any further questions on this matter.

Very truly yours,

A handwritten signature in black ink, appearing to read "Patrick Doherty", written over a light blue circular stamp.

Patrick Doherty
Director of Corporate Governance

Enclosures

Report on Content Governance

WHEREAS, shareholders are concerned that Alphabet's Google is failing to effectively address content governance concerns, posing risks to shareholder value. These concerns extend globally across multiple Google platforms and products.

Google's attempts thus far to address content governance have misfired. For example, in October 2017, Google acknowledged its automated system incorrectly flagged Google Docs content as violating its terms of service, blocking user-generated content and inconveniencing users. In June 2018, YouTube apologized for repeatedly filtering or demonetizing content created by LGBTQ users.

Google's YouTube continues to provide a home for extremist content. In June 2018, U.S. news outlets reported that white supremacist and white nationalist content "found a home" on social network Google Plus, violating the user policy. Data & Society research institute says: "The [YouTube] platform, and its parent company, have allowed racist, misogynist, and harassing content to remain online — and in many cases, to generate advertising revenue — as long as it does not explicitly include slurs." The Network Contagion Research Institute, a group tracking the spread of hate speech, found that the man charged in a mass shooting at a Pittsburgh synagogue linked to racist and antisemitic YouTube videos 71 times.

In December 2018, the *Washington Post* reported: "A year after YouTube's chief executive promised to curb 'problematic' videos, it continues to harbor and even recommend hateful, conspiratorial videos, allowing racists, anti-Semites and proponents of other extremist views to use the platform as an online library for spreading their ideas. . . . The struggle to control the spread of such content poses ethical and political challenges to YouTube and its embattled parent company, Google."

These controversies have drawn regulatory scrutiny. The European Union, for example, announced measures intended to pressure Google and other companies to combat disinformation ahead of EU parliament elections in May 2019. Security Commissioner Julian King said: "No excuses, no more foot-dragging, because the risks are real.... They've got to get serious about this stuff." CEO Sundar Pichai was summoned to Congress to respond to questions about the spread of conspiracy theories on YouTube, during which he said: "We are constantly undertaking efforts to deal with misinformation ... we are looking to do more ... it's an area we acknowledge there's more work to be done."

Shareholders are concerned that Google's inability to address these issues proactively poses substantial regulatory, legal, and reputational risks to long-term value.

RESOLVED: Shareholders request Alphabet Inc. issue a report to shareholders at reasonable cost, omitting proprietary or legally privileged information, reviewing the efficacy of its enforcement of Google's terms of service related to content policies and assessing the risks posed by content management controversies related to election interference, freedom of expression, and the spread of hate speech, to the company's finances, operations, and reputation.

SUPPORTING STATEMENT: Proponents recommend the report include assessment of the scope of platform abuses and address related ethical concerns.

J.P.Morgan

Miriam G. Awad
Vice President
CIB Client Service Americas

December 20, 2018

Mr. David C. Drummond
Corporate Secretary
Alphabet Inc.
1600 Amphitheatre Parkway
Mountain View, California 94043

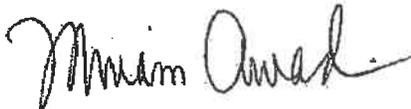
Dear Mr. Drummond,

This letter is in response to a request by The Honorable Thomas P. DiNapoli, New York State Comptroller, regarding confirmation from JP Morgan Chase that the New York State Common Retirement Fund has been a beneficial owner of Alphabet, Inc. continuously for at least one year as of and including December 20, 2018.

Please note that J.P. Morgan Chase, as custodian for the New York State Common Retirement Fund, held a total of 813,797 shares of common stock as of December 20, 2018 and continues to hold shares in the company. The value of the ownership stake continuously held by the New York State Common Retirement Fund had a market value of at least \$2,000.00 for at least twelve months prior to, and including, said date.

If there are any questions, please contact me at (212) 623-8481.

Regards,



Miriam Awad

cc: Patrick Doherty – NYSCRF
Tana Goldsmith – NYSCRF
Kyle Seeley - NYSCRF



UPS Next Day Air

Saver®

Shipping Document

See instructions on back. Visit UPS.com® or call 1-800-PICK-UPS® (800-742-5877) for additional information and Tariff/Terms and Conditions.

TRACKING NUMBER 1Z W49 033 23 1003 290 5

SHIPMENT FROM

SHIPPER'S UPS ACCOUNT NO. W 4 9 0 3 3

REFERENCE NUMBER

NAME Patrick Doherty TELEPHONE 212 383-1428

COMPANY OFC OF THE STATE COMPTRROLLER

STREET ADDRESS 59 MAIDEN LN FLR 30

CITY AND STATE NEW YORK NY ZIP CODE 10038

EXTREMELY URGENT DELIVERY TO

NAME David Drummond TELEPHONE 650 253-3393

COMPANY Alphabet, Inc.

STREET ADDRESS 1600 Amphitheatre Pkwy DEPT./FLR. Residential Delivery

CITY AND STATE Mountain View, CA ZIP CODE 94043



3 WEIGHT	LTR	WEIGHT	DIMENSIONAL WEIGHT If Applicable	LARGE PACKAGE	4 SHIPPER RELEASE
	<input type="checkbox"/>	<input checked="" type="checkbox"/>		<input type="checkbox"/>	

1P
1P

5 NEXT DAY AIR SAVER CHARGE		CHARGES
		\$

6 OPTIONAL SERVICES	<input type="checkbox"/> SATURDAY PICKUP See instructions.	\$
	<input type="checkbox"/> DECLARED VALUE FOR CARRIAGE For declared value over \$100, see instructions.	\$ AMOUNT
	<input type="checkbox"/> C.O.D. If C.O.D. enter amount to be collected and attach completed UPS C.O.D. tag to package.	\$ AMOUNT

7 ADDITIONAL HANDLING CHARGE	<input type="checkbox"/> An Additional Handling Charge applies for certain items. See instructions.	\$
------------------------------	---	----

8 TOTAL CHARGES		\$
-----------------	--	----

8 METHOD OF PAYMENT	BILL SHIPPER'S ACCOUNT NUMBER <input checked="" type="checkbox"/> In Section 1	BILL RECEIVER <input type="checkbox"/>	BILL THIRD PARTY <input type="checkbox"/>	CREDIT CARD <input type="checkbox"/>	American Express Diner's Club MasterCard Visa	CHECK <input type="checkbox"/>
	Record Account No. In Section 9					

9 RECEIVER'S / THIRD PARTY'S UPS ACCT. NO. OR MAJOR CREDIT CARD NO.	EXPIRATION DATE

THIRD PARTY'S COMPANY NAME AND ADDRESS

STREET ADDRESS

CITY AND STATE

ZIP CODE

Shipper authorizes UPS to act as forwarding agent for export control and customs purposes. Shipper certifies that these commodities, technology or software, if exported from the United States, were exported in accordance with the Export Administration Regulations. Diversion contrary to law is prohibited.

10 SHIPPER'S SIGNATURE	DATE OF SHIPMENT
X	12/26/18

All shipments are subject to the terms contained in the UPS Tariff/Terms and Conditions of Service, which are available at UPS.com and local UPS offices.

0201952542609 6/12 RRD

SHIPPER'S COPY

Unless a greater value is recorded in writing on this receipt, the shipper agrees that the declared value of each package covered by this receipt is \$100, which is a maximum liability under the circumstances surrounding the shipment. The shipper's liability for loss or damage to the contents of a package shipped by air is limited to the actual value of the contents, up to a maximum of \$500,000, except for packages shipped via the UPS Direct Mail, in which case the maximum liability per package assumed by UPS shall not exceed \$500, regardless of the actual value of the contents. For complete details of the limitations on UPS's liability, including the conditions of service of insurance, see the UPS Tariff/Terms and Conditions of Service.

This form not needed with UPS Internet Shipping at UPS.com

Proof of Delivery

Dear Customer,

This notice serves as proof of delivery for the shipment listed below.

Tracking Number

1ZW490332310032905

Service

UPS Next Day Air Saver

Delivered On

12/26/2018 8:37 A.M.

Delivered To

MOUNTAIN VIEW, CA, US

Received By

JULIAN

Left At

Receiver

Thank you for giving us this opportunity to serve you. Details are only available for shipments delivered within the last 120 days. Please print for your records if you require this information after 120 days.

Sincerely,

UPS

Tracking results provided by UPS: 02/12/2019 10:00 A.M. EST

APPENDIX B
Amended Proposal as Submitted
by the Co-lead Filer
and Included with
No Action Request

Arrived in mail slot
12/26/2018

ARJUNA  CAPITAL
ENLIGHTENED INVESTING

December 21, 2018

VIA OVERNIGHT MAIL

Alphabet Inc.
Attn: Corporate Secretary
1600 Amphitheatre Parkway
Mountain View, Ca 94043

To whom it may concern:

Arjuna Capital is an investment firm focused on sustainable and impact investing.

I am hereby authorized to notify you of our intention to co-lead file the enclosed shareholder resolution with Alphabet Incorporated on behalf of our clients Lisa Stephanie Myrkalo and Andrea Louise Dixon. Please expect an additional co-lead filing from the New York State Common Retirement Fund. Arjuna Capital submits this shareholder proposal for inclusion in the 2019 proxy statement, in accordance with Rule 14a-8 of the General Rules and Regulations of the Securities and Exchange Act of 1934 (17 C.F.R. § 240.14a-8). Per Rule 14a-8, Lisa Stephanie Myrkalo and Andrea Louise Dixon holds more than \$2,000 of GOOGL common stock, acquired more than one year prior to today's date and held continuously for that time. Our clients will remain invested in this position continuously through the date of the 2019 annual meeting.

Enclosed please find verification of this position and letter from Lisa Stephanie Myrkalo and Andrea Louise Dixon authorizing Arjuna Capital to undertake this filing on their behalf. We will send a representative to the stockholders' meeting to move the shareholder proposal as required by the SEC rules.

We would welcome discussion with Alphabet Incorporated about the contents of our proposal.

Please direct any written communications to me at the address below or to natasha@arjuna-capital.com. Please also confirm receipt of this letter via email.

Sincerely,



Natasha Lamb
Managing Partner
Arjuna Capital
1 Elm Street
Manchester, MA 01944

Enclosures

Report on Content Governance

WHEREAS, shareholders are concerned that Alphabet's Google is failing to effectively address content governance concerns, posing risks to shareholder value. These concerns extend globally across multiple Google platforms and products.

Google's attempts thus far to address content governance have misfired. For example, in October 2017, Google acknowledged its automated system incorrectly flagged Google Docs content as violating its terms of service, blocking user-generated content and inconveniencing users. In June 2018, YouTube apologized for repeatedly filtering or demonetizing content created by LGBTQ users.

Google's YouTube continues to provide a home for extremist content. In June 2018, U.S. news outlets reported that white supremacist and white nationalist content "found a home" on social network Google Plus, violating the user policy. Data & Society research institute says: "The [YouTube] platform, and its parent company, have allowed racist, misogynist, and harassing content to remain online — and in many cases, to generate advertising revenue — as long as it does not explicitly include slurs." The Network Contagion Research Institute, a group tracking the spread of hate speech, found that the man charged in a mass shooting at a Pittsburgh synagogue linked to racist and antisemitic YouTube videos 71 times.

In December 2018, the *Washington Post* reported: "A year after YouTube's chief executive promised to curb 'problematic' videos, it continues to harbor and even recommend hateful, conspiratorial videos, allowing racists, anti-Semites and proponents of other extremist views to use the platform as an online library for spreading their ideas. . . . The struggle to control the spread of such content poses ethical and political challenges to YouTube and its embattled parent company, Google."

These controversies have drawn regulatory scrutiny. The European Union, for example, announced measures intended to pressure Google and other companies to combat disinformation ahead of EU parliament elections in May 2019. Security Commissioner Julian King said: "No excuses, no more foot-dragging, because the risks are real.... They've got to get serious about this stuff." CEO Sundar Pichai was summoned to Congress to respond to questions about the spread of conspiracy theories on YouTube, during which he said: "We are constantly undertaking efforts to deal with misinformation ... we are looking to do more ... it's an area we acknowledge there's more work to be done." Shareholders are concerned that Google's inability to address these issues proactively poses substantial regulatory, legal, and reputational risks to long-term value.

RESOLVED: Shareholders request Alphabet Inc. issue a report to shareholders at reasonable cost, omitting proprietary or legally privileged information, reviewing the efficacy of its enforcement of Google's terms of service related to content policies and assessing the risks posed by content management controversies related to election interference, freedom of expression, and the spread of hate speech, to the company's finances, operations, and reputation.

SUPPORTING STATEMENT: Proponents recommend the report include assessment of the scope of platform abuses and address related ethical concerns.

CLEARY GOTTlieb STEEN & HAMILTON LLP

One Liberty Plaza
New York, NY 10006-1470

T: +1 212 225 2000

F: +1 212 225 3999

clearygottlieb.com

WASHINGTON, D.C. • PARIS • BRUSSELS • LONDON • MOSCOW
FRANKFURT • COLOGNE • ROME • MILAN • HONG KONG
BEIJING • BUENOSAIRES • SÃO PAULO • ABU DHABI • SEOUL

D: +1 (212) 225 2556
pmarcoqliese@csh.com

VICTOR J. LEWKOW
LEO C. BUCHHEIT
THOMAS J. MOLONEY
DAVID G. SABEL
JONATHAN I. BLACKMAN
YARON Z. REICH
RICHARD S. LINGER
JAMES A. DUNCAN
STEVEN M. LOEB
CRAIG B. BROD
NICOLAS GRABAR
CHRISTOPHER E. AUSTIN
HOWARD S. ZELBO
DAVID E. BRODSKY
ARTHUR H. KOHN
RICHARD J. COOPER
JEFFREY S. LEWIS
PAUL J. SHIM
STEVEN L. WILNER
ERIKA W. NIENHUIS
ANDRES DE LA CRUZ
DAVID C. LOPEZ
JAMES L. BROMLEY
MICHAEL A. GERSTENZANG
LEWIS J. LIMAN
LEV L. DASIN
NEIL O. WHORISKEY
JORGE U. JUANTORENA
MICHAEL D. WEINBERGER
DAVID LEINWAND
DIANA L. WOLLMAN
JEFFREY A. ROSENTHAL
ETHAN A. KLINGSBERG
MICHAEL D. DAYAN
CARMINE D. BOCCUZZI, JR.
JEFFREY D. KARPFF
KIMBERLY BROWN BLACKLOW
ROBERT J. RAYMOND
SUNG K. KANG
LEONARD C. JACOBY

SANDRA L. FLOW
FRANCISCO L. CESTERO
FRANCESCA L. ODELL
WILLIAM L. MCGRAE
JASON FACTOR
JOON H. KIM
MARGARET S. PEPONIS
LISA M. SCHWEITZER
JUAN G. GIRÁLDEZ
DUANE MCLAUGHLIN
BREON S. PEACE
MEREDITH E. KOTLER
GHANTALE KORDULA
BENET J. O'REILLY
ADAM E. FLEISHER
SEAN A. O'NEAL
GLENN P. MCGRORY
MATTHEW P. SALERNO
MICHAEL J. ALBANO
VICTOR L. HOU
ROGER A. COOPER
AMY R. SHAPIRO
JENNIFER KENNEDY PARK
ELIZABETH LENAS
LUKE A. BAREFOOT
PAMELA L. MARCOGLIESE
PAUL M. TIGER
JONATHAN S. KOLODNER
DANIEL ILAN
MEYER H. FEDIDA
ADRIAN R. LEIPSIK
ELIZABETH VICENS
ADAM J. BRENNEMAN
ARI D. MACKINNON
JAMES E. LANGSTON
JARED GERBER
COLIN D. LLOYD
COREY M. GOODMAN
RISHI ZUTSHI
JANE VANLARE

DAVID H. HERRINGTON
KIMBERLY R. SPOERRI
AARON J. MEYERS
DANIEL C. REYNOLDS
ABENA A. MAINOO
HUGH C. CONROY, JR.
JOSEPH LANZKRON
MAURICE R. GINDI
KATHERINE R. REAVES
RAHUL MUKHI
RESIDENT PARTNERS
SANDRA M. ROCKS
S. DOUGLAS BORISKY
JUDITH KASSEL
DAVID E. WEBB
PENELOPE L. CHRISTOPHORO
BOAZ S. MORAG
MARY E. ALCOCK
HEIDE H. LIGENFRITZ
KATHLEEN M. EMBERGER
WALLACE L. LARSON, JR.
AVRAM E. LUFT
ANDREW WEAVER
HELENA K. GRANNIS
JOHN V. HARRISON
CAROLINE F. HAYDAY
NEIL R. MARKEL
HUMAYUN KHALID
KENNETH S. BLAZEJEWSKI
ANDREA M. BASHAM
LAURA BAGARELLA
SHIRLEY M. LO
JONATHAN D.W. GIFFORD
SUSANNA E. PARKER
RESIDENT COUNSEL
LOUISE M. PARENT
OF COUNSEL

FEBRUARY 5, 2019

VIA E-MAIL (shareholderproposals@sec.gov)

U.S. Securities and Exchange Commission
Division of Corporation Finance
Office of Chief Counsel
100 F Street, N.E.
Washington, DC 20549

Re: Stockholder Proposal Submitted by Lisa Stephanie Myrkalo, Andrea Louise Dixon, and the New York State Common Retirement Fund

Ladies and Gentlemen:

We are writing on behalf of our client, Alphabet Inc., a Delaware corporation (“Alphabet” or the “Company”), pursuant to Rule 14a-8(j) under the Securities Exchange Act of 1934, as amended (the “Exchange Act”), to notify the staff of the Division of Corporation Finance (the “Staff”) of the Securities and Exchange Commission (the “Commission”) of the Company’s intention to exclude the shareholder proposal and accompanying supporting statement (the “Proposal”) submitted by Lisa Stephanie Myrkalo, Andrea Louise Dixon, and the New York State Common Retirement Fund (the “Proponents”), by a letter dated December 6, 2018, from the Company’s proxy statement for its 2019 annual meeting of shareholders (the “Proxy Statement”).

In accordance with Section C of SEC Staff Legal Bulletin No. 14D (Nov. 7, 2008) (“SLB 14D”), we are emailing this letter and its attachments to the Staff at shareholderproposals@sec.gov. In accordance with Rule 14a-8(j), we are simultaneously sending a copy of this letter and its attachments to the Proponents as notice of the Company’s intent to omit the Proposal from the Proxy Statement. The Company expects to file its definitive Proxy Statement with the Commission on or about April 26, 2019, and this letter is being filed

Cleary Gottlieb Steen & Hamilton LLP or an affiliated entity has an office in each of the cities listed above.

with the Commission no later than 80 calendar days before that date in accordance with Rule 14a-8(j). Rule 14a-8(k) and Section E of SLB 14D provide that shareholder proponents are required to send companies a copy of any correspondence that the shareholder proponent elects to submit to the Commission or the Staff. Accordingly, we are taking this opportunity to remind the Proponents that if the Proponents submit correspondence to the Commission or the Staff with respect to the Proposal, a copy of that correspondence should concurrently be furnished to the undersigned on behalf of the Company.

THE PROPOSAL

The Proposal is attached hereto as Exhibit A. The Proposal states in full:

WHEREAS: With an estimated 1.2 trillion searches per year worldwide – and billions of users – Alphabet Inc.’s Google sits at the center of global controversy regarding its role in Russia’s reported election interference during the 2016 United States presidential election and what experts say is an ongoing threat to the democratic process.

Shareholders are concerned that Google's failure to have proactively addressed this issue poses substantial regulatory, legal, and reputational risk to shareholder value.

In October 2017, *Bloomberg* reported Google found evidence Russian agents bought Google ads to interfere with the 2016 presidential campaign, using YouTube and Google’s main search advertising systems.

Richard Clark, cybersecurity adviser to President George W. Bush, and Robert Knake, cybersecurity adviser to President Barack Obama, wrote: “Russia could well interfere in the 2020 presidential vote, or the 2018 midterm elections... They will be back. And when they are, we better be ready with a plan that’s suited to our current moment.”

We believe Google has an obligation to demonstrate how it manages content to prevent violations of its terms of service. Yet, disclosures have been inadequate. Content policies appear reactive, not proactive.

Congressional committees have launched multiple investigations into Russian interference. Lawmakers plan to introduce legislation to require internet companies to disclose more information about political ad purchases. A United States Senator stated, “If Vladimir Putin is using Facebook or Google or Twitter to, in effect, destroy our democracy, the American people should know about it.”

The *New York Times* reported, “Despite Google’s insistence that its search algorithm undergoes a rigorous testing process to ensure that its results do not reflect political, gender, racial or ethnic bias, there is growing political support for regulating Google and other tech giants like public utilities and forcing it to disclose how exactly it arrives at search results.”

Foreign ministers of the Group of Seven countries, including the United States, said, “we are increasingly concerned about cyber-enabled interference in democratic political processes.” Germany enacted a law with fines of up to 50 million Euros if social media platforms don’t promptly remove posts containing unlawful content. The U.K. government is considering regulating Google as a news organization.

Advertisers have raised alarms about fake user accounts. Some companies have reduced expenditures on digital advertising. Nomura Securities estimated YouTube has lost up to 750 million dollars in revenue due to advertiser fear of being associated with objectionable content.

RESOLVED: Shareholders request Alphabet Inc. issue a report to shareholders at reasonable cost, omitting proprietary or legally privileged information, reviewing the efficacy of its enforcement of Google's terms of service related to content policies and assessing the risks posed by content management related to election interference, to the company's finances, operations, and reputation.

SUPPORTING STATEMENT: Proponents recommend the report include assessment of the scope of platform abuses and address related ethical concerns.

BASIS FOR EXCLUSION

In accordance with Rule 14a-8(i)(10), we hereby respectfully request that the Staff confirm that no enforcement action will be recommended against the Company if the Proposal is omitted from the Proxy Statement because the Company has substantially implemented the Proposal.

ANALYSIS

Under Rule 14a-8(i)(10), the Proposal may be omitted because it has been substantially implemented by the Company.

- A. *The Company’s actions have satisfactorily addressed the underlying concerns and “essential objectives” of the Proposal.*

Rule 14a-8(i)(10) permits a company to exclude a proposal from its proxy materials if the company “has already substantially implemented the proposal.” The purpose behind this exclusion has been described as follows:

“A company may exclude a proposal if the company is already doing or substantially doing what the proposal seeks to achieve. In that case, there is no reason to confuse shareholders or waste corporate resources in having shareholders vote on a matter that is moot. In the [Commission’s] words, the exclusion ‘is designed to avoid the possibility of shareholders having to consider matters

which have already been favorably acted upon by the management
.....”

Broc Romanek and Beth Young (W. Morley, editor), Shareholder Proposal Handbook, Sec. 23.01(8) at p. 23-4 (Aspen Law & Business 2003 ed.) (quoting SEC Release No. 34-12598 (July 7, 1976)). The determination that a company has substantially implemented the proposal depends upon whether the company’s policies, practices and procedures “compare favorably with the guidelines of the proposal.” Texaco, Inc. (avail. Mar. 28, 1991). See also, e.g., Albertson’s Inc. (avail. Mar. 23, 2005); The Talbots, Inc. (avail. Apr. 5, 2002); and Cisco Systems, Inc. (avail. Aug. 11, 2003). In other words, substantial implementation under Rule 14a-8(i)(10) requires a company’s actions to have satisfactorily addressed the underlying concerns and “essential objectives” of the proposal. See The Talbots, Inc. (avail. Apr. 5, 2002) (permitting omission of a proposal that required the establishment of a code of corporate conduct regarding human rights because the company had an existing Standard for Business Practice and Code of Conduct). Differences between a company’s actions and a proposal are permitted so long as the company’s actions satisfactorily address the proposal’s underlying concerns and essential objectives. See Release No. 34-20091. (Aug. 16, 1983).

The Proposal and its preamble make clear that the underlying concern and essential objective of the Proponents is for the Company to publish a report on election interference, regarding incidents such as the 2016 Russian election interference (the “Russian Interference”). The preamble to the Proposal leads with the statement “Google sits at the center of global controversy regarding its role in Russia’s reported election interference during the 2016 United States presidential election,” and continues discussing the Russian Interference almost exclusively, noting that “[s]hareholders are concerned that Google’s failure to have proactively addressed [the election interference] issue poses substantial regulatory, legal, and reputational risk to shareholder value.” This conclusion is further supported by the second half of the resolved clause, which clarifies that the Proponents mean only “the risks posed by content management *related to election interference*, to the company’s finances, operations, and reputation” (emphasis added). Any other interpretation that the Proposal is making a broad request that the Company review “the efficacy of its enforcement of Google’s terms of service related to content policies,” would be impermissibly vague and indefinite under 14a-8(i)(3).

Google, Alphabet’s principal subsidiary, was founded with the mission to “organize the world’s information and make it universally accessible and useful.” Alphabet strongly believes that the abuse of its platforms to spread misinformation is antithetical to that mission, and that it has a responsibility to prevent such abuses. Alphabet is committed to working with Congress, law enforcement, others in the industry, and the NGO community to strengthen protections around elections, ensure the security of users, and help combat disinformation. Since before the 2016 election, the Company has worked to detect and minimize opportunities for manipulation and abuse, and has built industry-leading security systems that are constantly evolving to stay ahead of ever-changing threats.

In response to similar concerns about the Russian Interference raised by Congress in 2017, the Company conducted investigations into whether individuals connected to government-backed entities were using Google’s platforms to disseminate information with the purpose of interfering with the 2016 U.S. election. The Company reported its findings through

publicly available written testimony by Kent Walker, then Senior Vice President and General Counsel of Google, on November 1, 2017 (the “2017 Testimony”). (See Exhibit B). The testimony gave a detailed overview of how the Company is committed to preventing any future interference incidents, and announced several initiatives to provide increased transparency in election advertising. Mr. Walker also reported that the actual amount of state-sponsored interference was limited because of the Company’s various safeguards and the fact that the Company’s products are not particularly suitable for the kind of micro-targeting or viral dissemination that these state-sponsored entities preferred. Nevertheless, the Company has been vigilant in tracking down and disabling accounts connected to such activities. The Company has continued to investigate state-sponsored election influence attempts, and on August 23, 2018, it issued a status update report through its official Company blog (the “Blog Update”). (See Exhibit C). In the Blog Update, Alphabet published a thorough status report on its findings, indicating which state actors it believes are behind certain attacks and misuses of its services. The Blog Update noted that “[a]ctors engaged in this type of influence operation violate our policies, and we swiftly remove such content from our services and terminate these actors’ accounts.” It also provided additional detail on the different accounts it had terminated. The Company has continued to revise the Blog Update with new information, with an update on November 20, 2018 to reflect changes since the original publication date.

On September 5, 2018, Mr. Walker provided another written testimony to Congress, providing further updates on the many actions the Company has undertaken to prevent attempts to use its services to interfere with elections (the “2018 Testimony”) (See Exhibit D). The 2018 Testimony announced that it had fulfilled all of the commitments to increase transparency in election advertising that it had announced in the 2017 Testimony. Specifically, the Company announced that it had implemented the following actions since the 2017 Testimony was published:

- The Company rolled out its Verification Program, which requires anyone who wants to purchase a federal election ad on Google in the U.S. to provide government-issued identification and other key information to confirm they are a U.S. citizen or lawful permanent resident or a U.S.-based organization.
- The Company has incorporated In-ad Disclosures to help people better understand who is paying for an election ad. The Company now identifies by name any advertiser running election-related campaigns on Search, YouTube, Display and Video 360, and the Google Display Network.
- The Company launched a “Political advertising on Google” Transparency Report for election ads, which provides data about the entities buying election-related ads on its platforms, how much money is spent across states and congressional districts on such ads, and who the top advertisers are overall. The report also shows the keywords advertisers have spent the most money on for political ads since the start of the 2018 U.S. midterm elections (from May 31, 2018 onwards).
- The Company now offers a searchable election “Ad Library,” updated weekly, within the public Transparency Report, which discloses relevant items such as which ads had

the most views, what the latest election ads running on the Company's platforms are, and provides deep dives into specific advertisers' campaigns. The disclosed data also includes the overall amount spent and number of ads run by each political advertiser, whether the advertiser targeted its ad campaigns geographically, by age or by gender, the approximate amount spent on each individual ad, the approximate impressions generated by each ad, and the dates each ad ran on the Company's platforms.

See Exhibit D.

In addition to these transparency efforts, the 2018 Testimony also highlighted some of the initiatives that the Company has undertaken to improve the cybersecurity of election infrastructure, candidates, and political campaigns. These measures show that the Company is committed to a proactive and comprehensive approach to preventing foreign election interference. These initiatives include the following:

- In October 2017, the Company unveiled the Advanced Protection Program, the strongest level of account protection offered by the Company. It is designed for users that may need extra protection against targeted spearphishing attacks, such as elected officials and candidates for public office.
- In May 2018, Jigsaw, the Company's technology incubator for next generation security solutions, announced that it would make available to U.S. political organizations a free service designed to protect against DDoS attacks called Project Shield.
- Since June 2012, the Company has continued to issue warnings to users when it believes that they are at risk of state-sponsored efforts to hijack their accounts.

The Company also partners with state actors and NGOs in order to promote election security. For example, Alphabet has supported significant outreach to increase security for candidates and campaigns across the United States, France, Germany, and other countries. The Company has also partnered with the National Cyber Security Alliance to help promote better account security, which includes security training programs that focus specifically on elected officials and staff members. Furthermore, the Company continues to support the bipartisan Defending Digital Democracy Project at the Belfer Center for Science and International Affairs at Harvard Kennedy School.

Finally, the Company has already disclosed its assessments of the risks that election interference, as well as actions taken by the Company to prevent election interference, may pose to the Company's finances, operations, and reputation. In the Risk Factors section of its Annual Report for the fiscal year ended December 31, 2018, the Company highlighted how cybersecurity risks or changes to the Company's policies and practices could impact these areas. For example, the Company stated that "changes to our . . . advertising policies or the practices or policies of third parties may affect the type of ads and/or manner of advertising that we are able to provide which could also have an adverse effect on our business." The Company further described that "[It] experience[s] cyber attacks of varying degrees and other attempts to gain unauthorized access to [its] systems on a regular basis" and that "failure to abide by [its] privacy

policies, inadvertent disclosure that results in the release of [its] users' data, or in [its] or [its] users' inability to access such data, could result in government investigations and other liability, legislation or regulation, seriously harm [the Company's] reputation and brand and, therefore, [its] business, and impair [its] ability to attract and retain users. [The Company] expect[s] to continue to expend significant resources to maintain state-of-the-art security protections that shield against bugs, theft and, misuse or security vulnerabilities or breaches." Finally, the Company noted that "While [it has] dedicated significant resources to privacy and security incident response, including dedicated worldwide incident response teams, [the Company's] response process may not be adequate, may fail to accurately assess the severity of an incident, may not respond quickly enough, or may fail to sufficiently remediate an incident, among other issues. As a result, [the Company] may suffer significant legal, reputational, or financial exposure, which could adversely affect [its] business and results of operations."

Alphabet considers any attempted subversion of its services to be an attack upon its fundamental mission and values. It is deeply committed to preventing its services from being exploited in such a way, and has invested a great deal of effort into detecting and removing abusers from its platforms. It invests considerable resources to constantly improve the identification of deceptive content, while promoting content that is authoritative, relevant and current. Alphabet also believes that its users, advertisers, and creators must be able to trust in their security and safety. In the face of constantly evolving and unprecedented challenges, such as the Russian Interference, the Company has executed, and continues to execute, a meticulous investigation and has communicated its findings and solutions in great detail to the public while simultaneously creating channels for ongoing disclosure. Through these actions, the Company has already addressed the Proposal's underlying concern and essential objective and rendered the purpose of the Proposal moot.

B. *Substantial implementation of the Proposal does not require word-for-word adoption.*

Rule 14a-8(i)(10) permits exclusion of a proposal when a company has already substantially implemented the underlying concerns and essential objectives of the proposal, even when the manner by which a company implements the proposal does not correspond precisely to the actions sought by the proponent. In 1983, the Commission adopted the current interpretation of the exclusion, noting that, for a proposal to be omitted under this rule, it need not be implemented in full or precisely as presented. "In the past, the Staff has permitted the exclusion of proposals under Rule 14a-8(c)(10) [the predecessor provision to Rule 14a-8(i)(10)] only in those cases where the action requested by the proposal has been fully effected. The Commission proposed an interpretative change to permit the omission of proposals that have been 'substantially implemented by the issuer. While the new interpretative position will add more subjectivity to the application of the provision, the Commission has determined that the previous formalistic application of this provision defeated its purpose.'" SEC Release No. 34-20091 (Aug. 16, 1983). The 1998 amendments to the proxy rules reaffirmed this position. See SEC Release No. 34-40018 (May 21, 1998) at n.30 and accompanying text.

The Staff has consistently taken the position that a company need not comply with every detail of a proposal or implement every aspect of a proposal in order to make a determination that the proposal has been substantially implemented and, therefore, can be

excluded under Rule 14a-8(i)(10). See, e.g., Ford Motor Company (avail. Feb. 22, 2016); Symantec Corporation (avail. June 3, 2010); Bank of America Corp. (avail. Jan. 4, 2008); AutoNation Inc. (avail. Feb. 10, 2004); and AMR Corporation (avail. Apr. 17, 2000). In each of these letters, the Staff concurred that a company may omit a shareholder proposal from its proxy materials under Rule 14a-8(i)(10) where the proposal was not implemented exactly as proposed.

The Proposal requests that the Company issue a report “reviewing the efficacy of its enforcement of Google's terms of service related to content policies and assessing the risks posed by content management related to election interference, to the company's finances, operations, and reputation.” However, the nine paragraphs preceding the resolved clause of the Proposal make clear that the Proponent's underlying concern and essential objective is to have a report on election interference. As noted above, a broader interpretation that the Proposal is making a request that the Company review the general efficacy of its enforcement of Google's terms of service related to content policies would be impermissibly vague and indefinite under 14a-8(i)(3). As discussed in Section A, the Company has already conducted, and continues to conduct, several investigations into abusers of Google's terms of service, and how such abuses can affect election interference. Although the Company may not have complied with every detail of the Proposal, it has nevertheless substantially implemented it, which is all that is required under Rule 14a-8(i)(10).

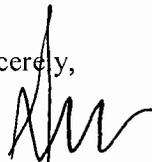
For the reasons noted above, the Company has substantially implemented the requirements of the Proposal. These actions have therefore addressed the underlying concern and essential objective of the Proposal. In order to “avoid the possibility of shareholders having to consider matters which already have been favorably acted upon by . . . management,” SEC Release No. 34-12598 (July 7, 1976), the Company respectfully requests the Staff's concurrence in the omission of the Proposal as having been substantially implemented pursuant to Rule 14a-8(i)(10).

* * * * *

Conclusion

By copy of this letter, the Proponents are being notified that for the reasons set forth herein the Company intends to omit the Proposal from its Proxy Statement. We respectfully request that the Staff confirm that it will not recommend any enforcement action if the Company omits the Proposal from its Proxy Statement. If we can be of assistance in this matter, please do not hesitate to call me.

Sincerely,



Pamela L. Marcogliese

Enclosures

Securities and Exchange Commission, p. 9

cc: Patrick Doherty, Office of the New York State Comptroller, on behalf of the New York State Common Retirement Fund; and

Natasha Lamb, Managing Partner, Arjuna Capital, on behalf of Lisa Stephanie Myrkalo and Andrea Louise Dixon

EXHIBIT A

THOMAS P. DINAPOLI
STATE COMPTROLLER



STATE OF NEW YORK
OFFICE OF THE STATE COMPTROLLER

DIVISION OF CORPORATE GOVERNANCE
59 Maiden Lane-30th Floor
New York, NY 10038
Tel: (212) 383-3931
Fax: (212) 681-4468

December 6, 2018

Mr. David C. Drummond
Corporate Secretary
Alphabet Inc.
1600 Amphitheatre Parkway
Mountain View, California 94043

Dear Mr. Drummond:

The Comptroller of the State of New York, Thomas P. DiNapoli, is the trustee of the New York State Common Retirement Fund (the "Fund") and the administrative head of the New York State and Local Retirement System. The Comptroller has authorized me to inform of his intention to offer the enclosed shareholder proposal for consideration of stockholders at the next annual meeting.

I submit the enclosed proposal to you in accordance with rule 14a-8 of the Securities Exchange Act of 1934 and ask that it be included in your proxy statement.

A letter from J.P. Morgan Chase, the Fund's custodial bank verifying the Fund's ownership of Alphabet Inc. shares, continually for over one year, is enclosed. The Fund intends to continue to hold at least \$2,000 worth of these securities through the date of the annual meeting.

We would be happy to discuss this initiative with you. Should Alphabet decide to endorse its provisions as company policy, the Comptroller will ask that the proposal be withdrawn from consideration at the annual meeting. Please feel free to contact me at (212) 383-1428 and/or email at pdoherty@osc.ny.gov should you have any further questions on this matter.

Very truly yours,

A handwritten signature in black ink, appearing to read "Patrick Doherty", written over a large, stylized, light-colored scribble or watermark.

Patrick Doherty
Director of Corporate Governance

Enclosures

Report on Content Governance

WHEREAS: With an estimated 1.2 trillion searches per year worldwide – and billions of users – Alphabet Inc.'s Google sits at the center of global controversy regarding its role in Russia's reported election interference during the 2016 United States presidential election and what experts say is an ongoing threat to the democratic process.

Shareholders are concerned that Google's failure to have proactively addressed this issue poses substantial regulatory, legal, and reputational risk to shareholder value.

In October 2017, *Bloomberg* reported Google found evidence Russian agents bought Google ads to interfere with the 2016 presidential campaign, using YouTube and Google's main search advertising systems.

Richard Clark, cybersecurity adviser to President George W. Bush, and Robert Knake, cybersecurity adviser to President Barack Obama, wrote: "Russia could well interfere in the 2020 presidential vote, or the 2018 midterm elections...They will be back. And when they are, we better be ready with a plan that's suited to our current moment."

We believe Google has an obligation to demonstrate how it manages content to prevent violations of its terms of service. Yet, disclosures have been inadequate. Content policies appear reactive, not proactive.

Congressional committees have launched multiple investigations into Russian interference. Lawmakers plan to introduce legislation to require internet companies to disclose more information about political ad purchases. A United States Senator stated, "If Vladimir Putin is using Facebook or Google or Twitter to, in effect, destroy our democracy, the American people should know about it."

The *New York Times* reported, "Despite Google's insistence that its search algorithm undergoes a rigorous testing process to ensure that its results do not reflect political, gender, racial or ethnic bias, there is growing political support for regulating Google and other tech giants like public utilities and forcing it to disclose how exactly it arrives at search results."

Foreign ministers of the Group of Seven countries, including the United States, said, "we are increasingly concerned about cyber-enabled interference in democratic political processes." Germany enacted a law with fines of up to 50 million Euros if social media platforms don't promptly remove posts containing unlawful content. The U.K. government is considering regulating Google as a news organization.

Advertisers have raised alarms about fake user accounts. Some companies have reduced expenditures on digital advertising. Nomura Securities estimated YouTube has lost up to 750 million dollars in revenue due to advertiser fear of being associated with objectionable content.

RESOLVED: Shareholders request Alphabet Inc. issue a report to shareholders at reasonable cost, omitting proprietary or legally privileged information, reviewing the efficacy of its enforcement of Google's terms of service related to content policies and assessing the risks posed by content management related to election interference, to the company's finances, operations, and reputation.

SUPPORTING STATEMENT: Proponents recommend the report include assessment of the scope of platform abuses and address related ethical concerns.

J.P.Morgan

Miriam G. Awad

Vice President
CIB Client Service Americas

December 6, 2018

Mr. David C. Drummond
Corporate Secretary
Alphabet Inc.
1600 Amphitheatre Parkway
Mountain View, California 94043

Dear Mr. Drummond,

This letter is in response to a request by The Honorable Thomas P. DiNapoli, New York State Comptroller, regarding confirmation from JP Morgan Chase that the New York State Common Retirement Fund has been a beneficial owner of Alphabet Inc. continuously for at least one year as of and including December 06, 2018.

Please note that J.P. Morgan Chase, as custodian for the New York State Common Retirement Fund, held a total of 830,661 shares of common stock as of December 06, 2018 and continues to hold shares in the company. The value of the ownership stake continuously held by the New York State Common Retirement Fund had a market value of at least \$2,000.00 for at least twelve months prior to, and including, said date.

If there are any questions, please contact me at (212) 623-8481.

Regards,



Miriam Awad

cc: Patrick Doherty – NYSCRF
Tana Goldsmith – NYSCRF
Kyle Seeley - NYSCRF

Arrived in mail slot
12/26/2018

ARJUNA  CAPITAL
ENLIGHTENED INVESTING

December 21, 2018

VIA OVERNIGHT MAIL

Alphabet Inc.
Attn: Corporate Secretary
1600 Amphitheatre Parkway
Mountain View, Ca 94043

To whom it may concern:

Arjuna Capital is an investment firm focused on sustainable and impact investing.

I am hereby authorized to notify you of our intention to co-lead file the enclosed shareholder resolution with Alphabet Incorporated on behalf of our clients Lisa Stephanie Myrkalo and Andrea Louise Dixon. Please expect an additional co-lead filing from the New York State Common Retirement Fund. Arjuna Capital submits this shareholder proposal for inclusion in the 2019 proxy statement, in accordance with Rule 14a-8 of the General Rules and Regulations of the Securities and Exchange Act of 1934 (17 C.F.R. § 240.14a-8). Per Rule 14a-8, Lisa Stephanie Myrkalo and Andrea Louise Dixon holds more than \$2,000 of GOOGL common stock, acquired more than one year prior to today's date and held continuously for that time. Our clients will remain invested in this position continuously through the date of the 2019 annual meeting.

Enclosed please find verification of this position and letter from Lisa Stephanie Myrkalo and Andrea Louise Dixon authorizing Arjuna Capital to undertake this filing on their behalf. We will send a representative to the stockholders' meeting to move the shareholder proposal as required by the SEC rules.

We would welcome discussion with Alphabet Incorporated about the contents of our proposal.

Please direct any written communications to me at the address below or to natasha@arjuna-capital.com. Please also confirm receipt of this letter via email.

Sincerely,



Natasha Lamb
Managing Partner
Arjuna Capital
1 Elm Street
Manchester, MA 01944

Enclosures

Report on Content Governance

WHEREAS, shareholders are concerned that Alphabet's Google is failing to effectively address content governance concerns, posing risks to shareholder value. These concerns extend globally across multiple Google platforms and products.

Google's attempts thus far to address content governance have misfired. For example, in October 2017, Google acknowledged its automated system incorrectly flagged Google Docs content as violating its terms of service, blocking user-generated content and inconveniencing users. In June 2018, YouTube apologized for repeatedly filtering or demonetizing content created by LGBTQ users.

Google's YouTube continues to provide a home for extremist content. In June 2018, U.S. news outlets reported that white supremacist and white nationalist content "found a home" on social network Google Plus, violating the user policy. Data & Society research institute says: "The [YouTube] platform, and its parent company, have allowed racist, misogynist, and harassing content to remain online — and in many cases, to generate advertising revenue — as long as it does not explicitly include slurs." The Network Contagion Research Institute, a group tracking the spread of hate speech, found that the man charged in a mass shooting at a Pittsburgh synagogue linked to racist and antisemitic YouTube videos 71 times.

In December 2018, the *Washington Post* reported: "A year after YouTube's chief executive promised to curb 'problematic' videos, it continues to harbor and even recommend hateful, conspiratorial videos, allowing racists, anti-Semites and proponents of other extremist views to use the platform as an online library for spreading their ideas. . . . The struggle to control the spread of such content poses ethical and political challenges to YouTube and its embattled parent company, Google."

These controversies have drawn regulatory scrutiny. The European Union, for example, announced measures intended to pressure Google and other companies to combat disinformation ahead of EU parliament elections in May 2019. Security Commissioner Julian King said: "No excuses, no more foot-dragging, because the risks are real.... They've got to get serious about this stuff." CEO Sundar Pichai was summoned to Congress to respond to questions about the spread of conspiracy theories on YouTube, during which he said: "We are constantly undertaking efforts to deal with misinformation ... we are looking to do more ... it's an area we acknowledge there's more work to be done." Shareholders are concerned that Google's inability to address these issues proactively poses substantial regulatory, legal, and reputational risks to long-term value.

RESOLVED: Shareholders request Alphabet Inc. issue a report to shareholders at reasonable cost, omitting proprietary or legally privileged information, reviewing the efficacy of its enforcement of Google's terms of service related to content policies and assessing the risks posed by content management controversies related to election interference, freedom of expression, and the spread of hate speech, to the company's finances, operations, and reputation.

SUPPORTING STATEMENT: Proponents recommend the report include assessment of the scope of platform abuses and address related ethical concerns.

December 5, 2018

Natasha Lamb
Managing Partner
Arjuna Capital
353 W. Main Street
Durham, NC 27701

Dear Ms. Lamb,

We hereby authorize Arjuna Capital to file a shareholder proposal on our behalf Alphabet Inc Class A (GOOGL) regarding Content Governance, asking for a report evaluating the Company's current strategies and policies on content governance, including the impact on human rights, democracy, and freedom of expression, as well as the reputational, regulatory, and financial risks posed by content governance controversies, for the company's annual meeting in 2019.

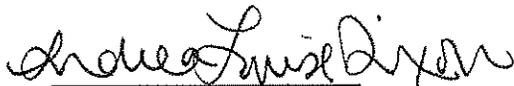
We are the beneficial owners of more than \$2,000 worth of common stock in Alphabet Inc Class A (GOOGL) that we have held continuously for more than one year. We intend to hold the aforementioned shares of stock through the date of the company's annual meeting in 2019.

We specifically give Arjuna Capital full authority to deal, on our behalf, with any and all aspects of the aforementioned shareholder proposal. We understand that our names may appear on the corporation's proxy statement as the filer of the aforementioned proposal.

Sincerely,



Lisa Stephanie Mrykalo



Andrea Louise Dixon

c/o Arjuna Capital
353 W. Main Street
Durham, NC 27701

charles SCHWAB

December 21, 2018

To WHOM IT MAY CONCERN:

Re: Nusbaum Family Irrevocable Trust;
Andrea Louise Dixon & Lisa Stephanie Mrykalo /; ***

This letter is to confirm that Charles Schwab & Co. is the record holder for the beneficial owner of the account above (# ***), which Arjuna Capital manages and which holds 14 shares of common stock in Alphabet Inc Class A (GOOGL).*

As of December 21, 2018, Nusbaum Family Irrevocable Trust; Andrea Louise Dixon & Lisa Stephanie Mrykalo held, and have held continuously for at least one year, 14 shares of Alphabet Inc Class A (GOOGL) stock.

This letter serves as confirmation that the account holder listed above is the beneficial owner of the above referenced stock.

Sincerely,



5/9/16


Jonnalee Owens
Relationship Specialist/Advisor Services

Case ID # AM-2563754

Independent investment advisors are not owned by, affiliated with, or supervised by Charles Schwab & Co., Inc. ("Schwab"). ©2016 Charles Schwab & Co., Inc. All rights reserved.

Member SIPC. CRS 00038 (0609-9534) 09/16 SGC48613-00



EXHIBIT B



Written Testimony of Kent Walker
Senior Vice President and General Counsel, Google
Senate Select Committee on Intelligence
Hearing on “Social Media Influence in the 2016 US Elections”
Written Congressional Testimony
November 1, 2017

Chairman Burr, Vice-Chair Warner, and members of the Committee, thank you for the opportunity to appear before you this morning.

My name is Kent Walker. I am Senior Vice President and General Counsel at Google and I lead our Legal, Policy, Trust and Safety, and Philanthropy teams. I’ve worked at the intersection of technology, security, and the law for over 25 years, including a tour early in my career as an Assistant US Attorney at the Department of Justice focusing on technology crimes.

We believe that we have a responsibility to prevent the misuse of our platforms and we take that very seriously. Google was founded with a mission to organize the world’s information and make it universally accessible and useful. The abuse of the tools and platforms we build is antithetical to that mission.

Google is deeply concerned about attempts to undermine democratic elections. We are committed to working with Congress, law enforcement, others in our industry, and the NGO community to strengthen protections around elections, ensure the security of users, and help combat disinformation.

We are dealing with difficult questions that balance free expression issues, unprecedented access to information, and the need to provide high quality content to our users. There are no easy answers here, but we are deeply committed to getting this right. We recognize the importance of this Committee’s mandate, and we welcome the opportunity to share information and talk about solutions.

Of course disinformation and propaganda campaigns aren't new, and have involved many different types of media and publications. When it comes to online platforms, for many years we have seen nation states and criminals attempt to breach our firewalls, game our search results, and interfere with our platforms. These attempts range from large-scale threats, such as distributed denial of service attacks, which we are able to identify and thwart quickly, all the way down to small-scale, extremely targeted attacks, such as attempts to gain access to email accounts of high-profile individuals.

We take these threats very seriously. We serve billions of users every day, so our solutions need to work at scale. We've built industry-leading security systems and we've put these tools into our consumer products. Back in 2007, we launched the first version of our Safe Browsing tool, which helps protect users from phishing, malware, and other attack vectors. Today, Safe Browsing is used on more than three billion devices worldwide. If we suspect that users are subject to government-sponsored attacks we warn them. And last month, we launched our Advanced Protection Program, which integrates physical security keys to protect those at greatest risk of attack, like journalists, business leaders, and politicians. We face motivated and resourceful attackers, and we are continually evolving our tools to stay ahead of ever-changing threats.

Our tools don't just protect our physical and network security, but also detect and prevent artificially boosting content, spam, and other attempts to manipulate our systems. On Google News, for example, we label links so users can see if the content is locally sourced, an OpEd, or an in-depth piece. For Google Search, we have updated our quality guidelines and evaluations to help identify misleading information, helping surface more authoritative content from the web. We have updated our advertising guidelines to prohibit ads on sites that misrepresent themselves. And on YouTube we employ a sophisticated spam and security-breach detection system to detect anomalous behavior and catch people trying to inflate view counts of videos or numbers of subscribers.

We have deployed our most advanced technologies to increase security and fight manipulation, but we realize that no system is going to be 100% perfect. It is hard to rapidly identify all untruthful content at massive scale, and harder yet to understand the motives and potential connections of the people posting that content. But we have made substantial progress

in preventing and detecting abuse, and we are seeing continued success in stopping bad actors attempting to game them. And as threats evolve, we will continue to adapt in order to understand and prevent new attempts to misuse our platforms.

With respect to the Committee's work on the 2016 election, we have looked across our products to understand whether individuals apparently connected to government-backed entities were using those products to disseminate information with the purpose of interfering with the US election. We based this review on research into misinformation campaigns from Alphabet's Jigsaw group, our information security team's own methods, and leads provided by other companies.

While we did find activity associated with suspected government-backed accounts, that activity appears to have been limited on our platforms. Of course, any activity like this is more than we would like to see. We have provided the relevant information to the Committee, have issued a public summary of the results of our review, and will continue to cooperate with the Committee's investigation.

Starting with our ads products, we found two accounts that appear to be associated with this effort. These accounts spent approximately \$4700 dollars in connection with the 2016 presidential election, representing less than 0.0002 percent of the total amount spent on that race. We believe that the activity we found was limited because of various safeguards that we had in place in advance of the 2016 election, and the fact that Google's products didn't lend themselves to the kind of micro-targeting or viral dissemination that these actors seemed to prefer.

As part of our investigation, we also looked at our other services. Let me share a few key points. On YouTube, we did find 18 channels on YouTube with roughly 1,100 videos, a total of 43 hours of content, uploaded by individuals who we suspect are associated with this effort and which contained political content. That compares with the 400 hours of YouTube content uploaded every minute, and the over one billion hours of YouTube content watched every day. These videos generally had very low view counts; only around 3% percent had more than 5,000 views. The videos were not targeted to any particular sector of the US population as that's not

feasible on YouTube. Additionally, we found a limited number of Gmail accounts that appear to have been primarily used to set up accounts on social media platforms.

We continue to expand our use of cutting-edge technology to protect our users and will continue working with governments to ensure that our platforms aren't used for nefarious purposes.

We will also be making political advertising more transparent, easier for users to understand, and even more secure.

- In 2018 we'll release a transparency report for election ads, sharing data about who is buying election ads on our platforms and how much money is being spent. We'll pair our transparency report with a database of election ad creatives from across our ads products. And we will make the database available for public research.
- We're also going to make it easier for users to understand who bought the election ads they see on our networks. Going forward, users will be able to find the name of any advertiser running an election ad on Search, YouTube, and the Google Display Network with one click on an icon above the ad.
- We will continue enhancing our existing safeguards to ensure that we only permit US nationals to buy US election ads. We already tightly restrict which advertisers can serve ads to audiences based on their political leanings. Moving forward, we'll go further by verifying the identity of anyone who wants to run an election ad or use our political-interest-based tools and confirming that person is permitted to run that ad.

We certainly can't do this alone. Combating disinformation campaigns requires efforts from across the industry. We'll continue to work with other companies to better protect the collective digital ecosystem, and, even as we take our own steps, we are open to working with governments on legislation that promotes electoral transparency.

Our commitment to addressing these issues extends beyond our services. Google has supported significant outreach to increase security for candidates and campaigns across the

United States, France, Germany, and other countries. We've offered in-person briefings and introduced a suite of digital tools designed to help election websites and political campaigns protect themselves from phishing, unauthorized account access, and other digital attacks. We've partnered with the National Cyber Security Alliance to fund and advise on security training programs that focus specifically on elected officials, campaigns, and staff members. We are also increasing our long-standing support for the bipartisan Defending Digital Democracy Project at the Belfer Center for Science and International Affairs at Harvard Kennedy School.

Let me conclude by recognizing the importance of the work of this Committee. Our users, advertisers, and creators must be able to trust in their security and safety. We share the goal of identifying bad actors who have attempted to interfere with our systems and abuse the electoral process. We look forward to continued cooperation, both with the members of this Committee and with our fellow companies, to both provide access to tools that help citizens express themselves while avoiding abuses that undercut the integrity of elections.

Thank you for the opportunity to tell you about our ongoing efforts in this space. We look forward to continuing to work with Congress on these important issues, and I'm happy to answer any questions you might have.

EXHIBIT C

SAFETY AND SECURITY

An update on state-sponsored activity

Kent Walker

SVP of Global Affairs

Published Aug 23, 2018



We've invested in robust systems to detect [phishing and hacking attempts](#), identify influence operations launched by foreign governments, and protect political campaigns from digital attacks through our [Protect Your Election](#) program.

Our Threat Analysis Group, working with our partners at [Jigsaw](#) and Google's Trust & Safety team, identifies bad actors, disables their accounts, warns our users about them, and shares intelligence with other companies and law enforcement officials.

This week, there has been a lot of news about attempted state-sponsored hacking and influence campaigns. We wanted to provide an update on some of our ongoing work in this area:

State-sponsored phishing attacks

Technical attribution of a recently-reported influence campaign from Iran

Detection and termination of activity on Google properties

State-sponsored phishing attacks

Phishing—attempts to trick users into providing a password that an attacker can use to sign into an account—remains a threat to all email users. Our improving technology has enabled us to [significantly decrease the volume](#) of phishing emails that get through to our users. Automated protections, account security (like security keys), and specialized warnings give Gmail users industry-leading security. As part of our security efforts, for the past

eight years, we've displayed prominent warnings to Gmail users who are at risk of phishing by potentially state-sponsored actors (even though in most

cases the specific phishing attempt never reaches the user's inbox).

In recent months, we've detected and blocked attempts by state-sponsored actors in various countries to target political campaigns, journalists, activists, and academics located around the world. When we've seen these types of attacks, we've notified users as well as law enforcement.

On Monday morning, we issued our most recent series of notifications to Gmail users who were subject to suspicious emails from a wide range of countries. We posted about these sorts of warnings [here](#)—if you received this type of warning, please read the blog post and take action immediately.

Iran and FireEye

To complement the work of our internal teams, we engage FireEye, a leading cybersecurity group, and other top security consultants, to provide us with intelligence. For the last two months, Google and Jigsaw have worked closely with FireEye on the influence operation linked to Iran that [FireEye identified this week](#). We're grateful to FireEye for identifying some suspicious Google accounts (three email accounts, three YouTube channels, and three Google+ accounts), which we swiftly disabled. [FireEye's full report](#) has just been published today. It's worth reading.

In addition to the intelligence we received from FireEye, our teams have investigated a broader range of suspicious actors linked to Iran who have engaged in this effort. We've updated U.S. lawmakers and law enforcement about the results of our investigation, including its relation to political content in the United States. We wanted to provide a summary of what we told them.

Connections to IRIB: forensic evidence

Our technical research has identified evidence that these actors are associated with the [IRIB](#), the Islamic Republic of Iran Broadcasting.

We can't go into all the technical details without giving away information that would be helpful to others seeking to abuse our platforms, but we have observed the following:

Technical data associated with these actors is strongly linked to the official IRIB IP address space.

Domain ownership information about these actors is strongly linked to IRIB account information.

Account metadata and subscriber information associated with these actors is strongly linked to the corresponding information associated with the IRIB, indicating common ownership and control.

These facts, taken together with other technical signals and analysis, indicate that this effort was carried out as part of the overall operations of the IRIB organization, since at least January 2017. This finding is consistent with internet activity we've warned about in recent years from Iran.

We've regularly sent warnings to Gmail users about phishing attempts coming from Iran (including on Monday).

We detected and warned users about a covert man-in-the-middle [security attack in Iran in 2011](#).

We uncovered and warned users [about politically-motivated phishing in Iran in 2013](#).

Detecting and terminating activity on Google properties

Actors engaged in this type of influence operation violate our policies, and we swiftly remove such content from our services and terminate these actors' accounts. Additionally, we use a number of robust methods, including IP blocking, to prevent individuals or entities in Iran from opening advertising accounts.

We identified and terminated a number of accounts linked to the IRIB organization that disguised their connection to this effort, including while sharing English-language political content in the U.S.:

39 YouTube channels that had 13,466 total US views on relevant videos;

6 blogs on Blogger

13 Google+ accounts

Our investigations on these topics are ongoing and we will continue to share our findings with law enforcement and other relevant government entities in the U.S. and elsewhere, as well as with others in the industry.

The state-sponsored phishing attacks, and the actors associated with the IRIB that we've described above, are clearly not the only state-sponsored actors at work on the Internet. For example, last year we disclosed information about actors linked to the [Internet Research Agency \(IRA\)](#). Since then, we have continued to monitor our systems, and broadened the range of IRA-related actors against whom we've taken action. Specifically, we've detected and removed 42 YouTube channels, which had 58 English-language political videos (these videos had a total of fewer than 1,800 U.S. views). We've also identified and terminated the account associated with one blog on Blogger.

* Since our last update in August, we have identified and terminated a limited number of accounts linked to coordinated influence operations, including while sharing English-language political content in the U.S.:

With respect to Iran-linked operations: 6 G+ pages; 1 Blogger blog; and 34 YouTube channels that had 20,794 total US views on English-language political videos.

With respect to Russia-linked operations: 1 YouTube channel linked to the Internet Research Agency, which had only a single English-language political video with 94 US views.

We will continue to deal with attempts to abuse our systems by identifying bad actors, terminating their accounts, and sharing relevant information with Congress and law enforcement.

**As of November 20, 2018, we've revised this post with updates since August 23, 2018, the post's original publication date.*

EXHIBIT D



Written Testimony of Kent Walker
Senior Vice President, Global Affairs & Chief Legal Officer
Senate Select Committee on Intelligence
Hearing on, “Foreign Influence Operations’ Use of Social Media Platforms”
Written Congressional Testimony
September 5th, 2018

Chairman Burr, Vice Chairman Warner, and members of the Committee, thank you for the opportunity to provide an update on the efforts we’re making to secure our platforms ahead of the 2018 midterm elections in the US and for future elections around the world.

My name is Kent Walker. I am Senior Vice President, Global Affairs and Chief Legal Officer at Google, and I lead our Legal, Policy, Trust and Safety, and Google.org teams. I’ve worked at the intersection of technology, security, and the law for over 25 years, including time spent early in my career as an Assistant US Attorney at the Department of Justice focusing on technology crimes.

We believe that we have a responsibility to prevent the misuse of our platforms and we take that very seriously. Google was founded with a mission to organize the world’s information and make it universally accessible and useful. The abuse of the tools and platforms we build is antithetical to that mission.

In my testimony to the Committee last fall, I described the investigation we had conducted to understand whether individuals apparently connected to government-backed entities were using our products to disseminate information with the purpose of interfering with the 2016 US election. We based that review on research into misinformation campaigns from our Jigsaw group, our information security team’s own methods, and leads provided by other companies. We identified limited activity and we took swift action, disabling any accounts we found.

We have continued our efforts and work diligently to identify and remove actors from our products who mislead others regarding their identity. We've continued to investigate activity by the Internet Research Agency and other Russia-affiliated entities since we testified before the Committee last year. When we have found activity, we've removed the relevant accounts. We've also investigated activity linked to Iranian influence efforts, similarly removing the accounts we have linked to that activity. We recently issued an updated summary of the results of our review, and are happy to continue to provide our findings to Congress, law enforcement, and our peers. While the activity on our products remains limited, any activity like this is more than we would like to see. We will continue to invest resources to address this issue and to work with law enforcement, Congress, and other companies.

Google remains deeply concerned about attempts to undermine democratic elections. As we promised the Committee last year, we have now fulfilled all four of our commitments to provide increased transparency in election advertising:

First, we've rolled out a **Verification Program** for anyone who wants to purchase a federal election ad on Google in the U.S. We now require that advertisers provide government-issued identification information and other key information to confirm they are a U.S. citizen or lawful permanent resident or a U.S.-based organization, as required by law.

Second, to help people better understand who is paying for an election ad we've incorporated **In-ad Disclosures**, which means we now identify by name advertisers running election-related campaigns on Search, YouTube, Display and Video 360, and the Google Display Network.

Third, we have launched a "Political advertising on Google" **Transparency Report** for election ads, which provides data about the entities buying election-related ads on our platforms, how much money is spent across states and congressional districts on such ads, and who the top advertisers are overall. The report also shows the keywords advertisers have spent the most money on ads of political importance during the current U.S. election cycle from May 31st, 2018 onwards.

Finally, we now offer a **searchable election Ad Library** within our public Transparency Report which shows things like which ads had the highest views, what the latest election ads running on our platform are, and deep dives into specific advertisers' campaigns. The data shows the overall amount spent and number of ads run by each election advertiser, and whether the advertiser targeted its ad campaigns geographically or by age or gender, as well as the approximate amount spent on each individual ad, the approximate impressions generated by each ad, and the dates each ad ran on our platform. In addition, the data from the report and Ad Library is publicly available through Google Cloud's BigQuery. Using BigQuery's API, anyone can write code and run their own unique queries on this data set. Researchers, political watchdog groups and private citizens can use our data set to develop charts, graphs, tables or other visualizations of political advertising on Google Ads services. Together with the Transparency Report, these tools provide unprecedented, data-driven insights into election ads on our platform. We're updating the Transparency Report and Ad Library every week, so as we head into election season, anyone can see new ads that get uploaded or new advertisers that decide to run Google ads.

In addition to our transparency efforts, we have implemented a number of initiatives to improve the cybersecurity posture of candidates, campaigns, and the election infrastructure writ large.

First, in October 2017, we unveiled the **Advanced Protection Program**, which provides the strongest account protection that Google offers. The Advanced Protection Program is designed to thwart targeted spearphishing attacks; it is designed for users that may need an extra layer of protection against such attacks. This includes elected officials and candidates for public office. We have made continuous improvements to the Advanced Protection Program since I last testified before this Committee, including adding support for iOS and making it easier for interested users to purchase security keys. We have conducted extensive outreach to campaigns, candidates, journalists, and human rights activists to help drive use of security keys that are at the heart of the Advanced Protection Program.

Second, in May 2018, Jigsaw, our incubator dedicated to building technology to address significant security challenges, announced the availability of **Project Shield** to U.S. political organizations (e.g., candidates, campaigns, political action committees) registered with the

appropriate electoral authorities. Project Shield is a free service that uses Google technology to prevent distributed denial of service (DDoS) attacks that block access to content. DDoS attacks can threaten the integrity of elections by depriving citizens of important information about candidates and elections, and we have seen an increase in such attacks over time. Project Shield is now protecting a diverse array of political organizations. Project Shield also protects state, county, and local electoral board and voter information sites. Google's and Jigsaw's [Protect Your Election](#) site provides further information on the tools and resources we have made available to help campaigns, election officials, and news organizations defend against digital attacks.

Third, we **continue to issue warnings to users** when we believe they are at risk of state-sponsored efforts to hijack their accounts. We have issued these warnings for many years, and we will continue to do so in light of account hijacking campaigns we see from state-sponsored actors during this election cycle.

We have deployed our most advanced technologies to increase security and fight manipulation, but we realize that no system is going to be 100% perfect. Our algorithms are designed to identify content that many people find relevant and useful. We are constantly looking to find signals that help us identify deceptive content, while promoting content that is authoritative, relevant, and current. We have made substantial progress in preventing and detecting abuse, and are seeing continued success in stopping bad actors attempting to game our systems. And as threats evolve, we will continue to adapt in order to understand and prevent new attempts to misuse our platforms.

We certainly can't do this important work alone. Combating disinformation campaigns requires efforts from across the industry. We'll continue to work with other companies to better protect the collective digital ecosystem, and, even as we take our own steps, we are open to working with governments on legislation that promotes electoral transparency.

Our commitment to addressing these issues extends beyond our services. Google has supported significant outreach to increase security for candidates and campaigns across the United States, France, Germany, and other countries. We've [partnered with the National Cyber Security Alliance](#) to help promote better account security, which includes security training programs that focus specifically on elected officials and staff members. We also continue to

support the bipartisan Defending Digital Democracy Project at the Belfer Center for Science and International Affairs at Harvard Kennedy School. And in 2007, we launched the first version of our Safe Browsing tool, which helps protect users from phishing, malware, and other attack vectors by examining billions of URLs, software, and website content. We have made this tool free and publicly available to webmasters and developers so that they can protect their websites and applications from malicious actors. Today, Safe Browsing is used on more than three billion devices worldwide.

Let me conclude by recognizing the important work of this Committee. Our users, advertisers, and creators must be able to trust in their security and safety. We share the goal of identifying bad actors who have attempted to interfere with our systems and abuse the electoral process. We look forward to continued cooperation, both with the members of this Committee and with others in the security community and the broader industry, to provide tools that help citizens express themselves and to address abuses that undercut the integrity of elections.

Thank you for the opportunity to outline our ongoing efforts in this space. We look forward to continuing to work with Congress on these important issues, and to answering any further questions you might have.