Responses to Written Questions Submitted by Honorable John Thune to Maggie Stanphill

Question 1. Dr. Stephen Wolfram, a witness at this hearing who has spent his life working on the science and technology of artificial intelligence, described Google and other internet platforms as "automated content selection businesses," which he defined as entities that "work by getting large amounts of content they didn't themselves generate, then using what amounts to [artificial intelligence] to automatically select what content to deliver or to suggest to any particular user at any given time—based on data they've captured about the user." Does Google agree with these characterizations of its business by Dr. Wolfram? If not, please explain why not.

Response. Our mission is to organize the world's information and make it universally accessible and useful.

We have many different products that are designed differently and serve this mission in different ways, including:

● Google Search organizes information about webpages in our Search index.

● YouTube provides a platform for people to upload videos to the open web with ease, and makes it easy for people to access those videos.

● Our advertising products allow businesses large and small to reach customers around the world and grow their businesses.

In many cases, we use automated processes to organize the vast array of information available on our platforms and the Web in order to provide relevant, useful information to users in a timely and accessible manner.

We believe in ensuring our users have choice, transparency, and control over how they engage with all of our products; for instance, Google Search and YouTube have options that allow users to operate them without any input from their personal data or browsing data, as well as the ability to turn off autoplay of videos suggested by YouTube's recommendation system.

Question 2. In Dr. Wolfram's prepared testimony, he formulates possible market-based suggestions for large internet platforms to consider that would "leverage the exceptional engineering and commercial achievements of the [automated content selection] businesses, while diffusing current trust issues about content selection, providing greater freedom for users, and inserting new opportunities for market growth." Specifically, Dr. Wolfram asked "Why does every aspect of automated content selection have to be done by a single business? Why not open up the pipeline, and create a market in which users can make choices for themselves?"

a. In what he labels "Suggestion A: Allow Users to Choose among Final Ranking

Providers" Dr. Wolfram suggests that the final ranking of content a user sees doesn't have to be done by the same entity. Instead, there could be a single content platform but a variety of "final ranking providers", who use their own programs to actually deliver a final ranking to the user. Different final ranking providers might use different methods, and emphasize different kinds of content. But the point is to let users be free to choose among different providers.

Some users might prefer (or trust more) some particular provider—that might or might not be associated with some existing brand. Other users might prefer another provider, or choose to see results from multiple providers. Has Google considered Dr. Wolfram's suggestion to allow users to choose among final ranking providers? If so, please provide Google's reaction to Dr. Wolfram's proposal. If not, will Google commit to considering Dr. Wolfram's suggestion and providing a briefing to the Committee on its efforts to consider this suggestion?

Response. Today, users have myriad choices when it comes to finding and accessing all types of content online. There are a variety of providers that organize information in different ways.

For general-purpose search engines, consumers can choose among a range of options: Bing, Yahoo, and many more. DuckDuckGo, for instance, a relatively new search engine provider, hit a record 1 billion monthly searches in January 2019, demonstrating that a new entrant can compete in this space.

There are many ways consumers find and access news content on the Internet. They navigate directly to sites and use dedicated mobile apps. They access news articles via social media services like Twitter and Facebook. And they use aggregators like News 360 and Drudge Report.

It has never been easier for a new entrant to build and become a new 'final ranking provider' for end users. Developers today can build on free repositories of web index data, like Common Crawl, to build new search engines. This is the kind of underlying, common content "platform" Dr. Wolfram seems to describe.

b.   In what he labels "Suggestion B: Allow Users to Choose among Constraint Providers" Dr. Wolfram suggests putting constraints on results that automated content businesses generate, for example forcing certain kinds of balance. Much like final ranking providers in Suggestion A, there would be constraint providers who define sets of constraints. For example, a constraint provider could require that there be on average an equal number of items delivered to a user that are classified (say, by a particular machine learning system) as politically left-leaning or politically right-leaning. Constraint providers would effectively define computational contracts about properties they want results delivered to users to have. Different constraint providers would define different computational contracts. Some might want balance; others might want to promote particular types of content, and so on. But the idea is that users could decide what constraint provider they wish to use. Has Google considered Dr. Wolfram's suggestion to allow users to choose among constraint providers? If so, please provide Google's reaction to Dr. Wolfram's proposal. If not, will Google commit to considering Dr. Wolfram's suggestion and providing a briefing to the Committee on its efforts to consider this suggestion?

Response. Google cares deeply about giving users transparency, choice and control in our products and services. We offer a number of resources to help users better understand the products and services we provide. For example, users can control what Google account activity is used to customize their experiences, including adjusting what data is saved to their Google account, at myaccount.google.com. If users wish to consume content in a different way, there are many other platforms and websites where they can do so, as discussed above.

Question 3. Does Google believe that algorithmic transparency is a policy option Congress should be considering? If not, please explain why not.

Response. Transparency has long been a priority at Google to help our users understand how our products work. We must balance this transparency with the need to ensure that bad actors do not game our systems through manipulation, spam, fraud and other forms of abuse. Since Google launched our first Transparency Report in 2010, we've been sharing data that sheds light on how government actions and policies affect privacy, security, and access to information online. For Search, our How Search Works site provides extensive information to anyone interested in learning more about how Google Search, our algorithms, and Search features operate. The site includes information on our approach to algorithmic ranking. We offer extensive resources to all webmasters to help them succeed in having their content discovered online. We also publish our 160 page Search Quality Evaluator Guidelines, which explain in great detail what our search engine is aiming to achieve, and which form a crucial part of the process by which we assess proposed changes to our algorithms.

It's important to note, however, that there are tradeoffs with different levels of transparency, and we aim to balance various sensitivities. For example, disclosing the full code powering our product algorithms would make it easier for malicious actors to manipulate or game our systems, and create vulnerabilities that would represent a risk to our users–while failing to provide meaningful, actionable information to well-meaning users or researchers, notably due to the scale and the pace of evolution of our systems. Extreme model openness can also risk exposing user or proprietary information, causing privacy breaches or threatening the security of our platforms.

Regarding transparency in AI algorithms more broadly, in our own consumer research, we've seen that access to underlying source code is not useful to users. Rather, we have found that algorithmic explanation is more useful. We've identified a few hallmarks of good explanations: it accurately conveys information regarding the system prediction or recommendation; is clear, specific, relatable, and/or actionable; boosts understanding of the overall system; and takes appropriate account of context. In our research we have been demonstrating progress in designing interpretable AI models, model understanding, and data and model cards for more transparent model reporting (see our Responsible AI Practices for a full list of technical recommendations and work). And we've outlined more details where government, in collaboration with civil society and AI practitioners, has a crucial role to play in AI explainability standards, among other areas, in our paper Perspectives on Issues in AI Governance.

Question 4. Does Google believe that algorithmic explanation is a policy option that Congress should be considering? If not, please explain why not.

Response. Transparency has long been a priority at Google to help our users understand how our products work. We must balance this transparency with the need to ensure that bad actors do not game our systems through manipulation, spam, fraud and other forms of abuse. Since Google launched our first Transparency Report in 2010, we've been sharing data that sheds light on how government actions and policies affect privacy, security, and access to information online. For

Search, our How Search Works site provides extensive information to anyone interested in learning more about how Google Search, our algorithms, and Search features operate. The site includes information on our approach to algorithmic ranking. We offer extensive resources to all webmasters to help them succeed in having their content discovered online. We also publish our 160 page Search Quality Evaluator Guidelines, which explain in great detail what our search engine is aiming to achieve, and which form a crucial part of the process by which we assess proposed changes to our algorithms.

It's important to note, however, that there are tradeoffs with different levels of transparency, and we aim to balance various sensitivities. For example, disclosing the full code powering our product algorithms would make it easier for malicious actors to manipulate or game our systems, and create vulnerabilities that would represent a risk to our users–while failing to provide meaningful, actionable information to well-meaning users or researchers, notably due to the scale and the pace of evolution of our systems. Extreme model openness can also risk exposing user or proprietary information, causing privacy breaches or threatening the security of our platforms.

Regarding transparency in AI algorithms more broadly, in our own consumer research, we've seen that access to underlying source code is not useful to users. Rather, we have found that algorithmic explanation is more useful. We've identified a few hallmarks of good explanations: it accurately conveys information regarding the system prediction or recommendation; is clear, specific, relatable, and/or actionable; boosts understanding of the overall system; and takes appropriate account of context. In our research we have been demonstrating progress in designing interpretable AI models, model understanding, and data and model cards for more transparent model reporting (see our Responsible AI Practices for a full list of technical recommendations and work). And we' outlined more details where government, in collaboration with civil society and AI practitioners, has a crucial role to play in AI explainability standards, among other areas, in our paper Perspectives on Issues in AI Governance.

Question 5. At the hearing, I noted that the artificial intelligence behind internet platforms meant to enhance user engagement also has the ability, or at least the potential, to influence the thoughts and behaviors of literally billions of people. Does Google agree with that statement? If not, please explain why not.

Response. We strongly believe that AI can improve lives in a number of ways, though we also recognize that AI is a rapidly evolving technology that must be applied responsibly. For these reasons, we assess all of our AI applications in accordance with our Google AI Principles: be socially beneficial, avoid creating or reinforcing unfair bias, be built and tested for safety, be accountable to people, incorporate privacy design principles, uphold high standards of scientific excellence, and be made available for uses in accordance with these principles. Following these principles, we do not build AI products for the purpose of manipulating users. Furthermore, it would not be in our business interest to engage in activities that risk losing user trust.

Rather, like our other technologies, we are using AI to provide a better experience for our users, and our efforts are already proving invaluable in different ways. For example, nearly 1 billion unique users use Google Translate to communicate across language barriers, and more than 1

billion users use Google Maps to navigate roads, explore new places, and visualize places from the mountains to Mars. We also recently introduced an AI-powered app called Bolo to help improve childrens' reading skills; early results in India demonstrate that 64% of children showed an improvement in reading proficiency in just 3 months.

These opportunities to use AI for social good come with significant responsibility, and we have publicly outlined our commitment to responsible AI development—including algorithmic accountability and explainability—in the Google AI Principles (also see our Responsible AI Practices for a full list of technical recommendations and work).

Question 6. YouTube has offered an autoplay feature since 2015. The company also offers users the option of disabling autoplay. To date, what percentage of YouTube users have disabled autoplay?

Response. Autoplay is an optional feature we added based on user feedback, as users wanted an option for a smoother YouTube experience, like listening to the radio or having a TV channel on in the background. We added an easy on/off toggle for the feature so that users can make a choice about whether they want to keep autoplay enabled, depending on how they are using the platform in a given session. Many users have chosen to disable autoplay in some situations and enable it in others. Our priority is to provide users with clear ways to use the product according to their specific needs.

Question 7. How many minutes per day do users in the United States spend, on average, watching content from YouTube? How has this number changed since YouTube added the autoplay feature in 2015?

Response. YouTube is a global platform with over 2 billion monthly logged-in users. Every day people watch over a billion hours of video and generate billions of views. More than 500 hours of content are uploaded to YouTube every minute. We are constantly making improvements to YouTube's features and systems to improve the user experience, and would not attribute changes in user behavior over the course of four years to a single product change.

Question 8. What percentage of YouTube video views in the United States and worldwide are the result of clicks and embedded views from social media?

Response. YouTube provides a number of ways for users to discover content, including through social media. The ways users choose to engage with the platform vary depending on their individual preferences, the type of content, and many other factors.

Question 9. What percentage of YouTube video views in the United States and worldwide are the result of YouTube automatically suggesting or playing another video after the user finishes watching a video?

Response. A significant percentage of video views on YouTube come from recommendations. Overall, a majority of video views on YouTube come from recommendations. This includes what people see on their home feed, in search results and in Watch Next panels. Recommendations are a popular and useful tool that helps users discover new artists and creators

and surface content to users that they might find interesting or relevant to watch next. The ways users choose to engage with recommendations and YouTube's autoplay feature vary depending on user preferences, the type of content they are watching, and many other factors. Many users like to browse the Watch Next panel and choose the next video they want to play. In some cases, users want to continue to watch videos without having to choose the next video, for example if they are using YouTube to listen to music or to follow a set playlist of content. To provide users with choices, YouTube has an easy toggle switch to turn autoplay off if users do not want to have videos automatically play.

Question 10. What percentage of YouTube video views in the United States and worldwide are the result of users searching YouTube.com?

Response. When users first start using YouTube, they often begin by searching for a video. YouTube search works similar to Google search – users type a search query into the search box, and we present a list of videos or YouTube channels that are relevant to that search query. Videos are ranked based on a number of factors including how well the title and description match the query, what is in the video content, and how satisfied previous users were when they viewed these videos. The ways users choose to find content, including through the YouTube home page, searches, and recommendations vary depending on user preferences, the type of content, and many other factors.

Question 11. In 2018, YouTube started labeling videos from state-funded broadcasters. What impact, if any, have these labels had on the rate that videos from these channels are viewed, clicked on, and shared by users?

Response. If a channel is owned by a news publisher that is funded by a government, or publicly funded, an information panel providing publisher context may be displayed on the watch page of the videos on its channel. YouTube also has other information panels, including to provide topical context for well-established historical and scientific topics that have often been subject to misinformation online, like the moon landing. We have delivered more than 2.5 billion impressions across all of our information panels since July 2018.

Question 12. During the hearing, I discussed my efforts to develop legislation that will require internet platforms to give its users the option to engage with the platform without having the experience shaped by algorithms driven by user-specific data. In essence, the bill would require internet platforms like Google to provide users with the option of a "filter bubble-free" view of services such as Google search results, and enabling users to toggle between the opaque artificial intelligence driven personalized search results and the "filter bubble-free" search results. Does Google support, at least in principle, providing its users with the option of a "filter bubble-free" experience of its search results?

Response. There is very little personalization in organic Search results based on users' inferred interests or Search history before their current session. It doesn't take place often and generally doesn't significantly change organic Search results from one person to another. Most differences that users see between their organic 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. One of the most common reasons results may differ between people involves localized organic search results, when listings are customized to be relevant for anyone in a particular area. Localization isn't personalization because everyone in the same location gets similar results. Localization makes our search results more relevant. For example, people in the US searching for "football" do not generally want UK football results, and vice versa. People searching for "zoos" in one area often want locally-relevant listings.

Search does include some features that personalize results based on the activity in their Google account. For example, if a user searches for "events near me" Google may tailor some recommendations to event categories we think they may be interested in. These systems are designed to match a user's interests, but they are not designed to infer sensitive characteristics like race or religion. Overall, Google strives to make sure that our users continue to have access to a diversity of websites and perspectives.

Anyone who doesn't want personalization using account-based activity can disable it using the Web & App Activity setting. Users can also choose to keep their search history stored but exclude Chrome and app activity.

Question 13. In 2013, former Google Executive Chairman Eric Schmidt wrote that modern technology platforms like Google "are even more powerful than most people realize." Does Google agree that it is even more powerful than most people realize? If not, please explain why not.

Response. We are committed to providing users with powerful tools, and our users look to us to provide relevant, authoritative information. We work hard to ensure the integrity of our products, and we've put a number of checks and balances in place to ensure they continue to live up to our standards. We also recognize the important role of governments in setting rules for the development and use of technology. To that end, we support federal privacy legislation and proposed a legislative framework for privacy last year.

Question 14. Does Google believe it is important for the public to better understand how it uses artificial intelligence to make inferences from data about its users?

Response. Automated predictions and decision making can improve lives in a number of ways, from recommending music to monitoring a patient's vital signs, and we believe public explainability is crucial to being able to question, understand, and trust machine learning systems. We've identified a few hallmarks of good explanations: they accurately convey information regarding the system prediction or recommendation; are clear, specific, relatable, and/or actionable; boost understanding of the overall system; and take appropriate account of context.

We've also been taken numerous steps in our technical research to make our algorithms more understandable and transparent (see our Responsible AI Practices for a full list of technical recommendations and work), including:

● We've developed a lot of research and tools to help people better understand their data and design more interpretable models.

● We're also working on visualizing what's going on inside deep neural nets.

● And explainability is built into some projects such as predicting cardiovascular risk from images of the retina – our model shows what parts of the image most contributed to the prediction.

Explainability when it comes to machine learning is something we take very seriously, and we'll continue to work with researchers, academics, and public policy groups to make sure we're getting this right. It's important to note that government, in collaboration with civil society and AI practitioners, also has a crucial role to play in AI explainability standards, and we've outlined more details in our paper Perspectives on Issues in AI Governance.

Question 15. Does Google believe that its users should have the option to engage with their platform without being manipulated by algorithms powered by its users' own personal data? If not, please explain why not.

Response. Google cares deeply about giving users transparency, choice and control in our products and services. We offer a number of resources to help users better understand the products and services we provide. These resources include plain-English and easy-to-understand instructions about how users can make meaningful privacy and security choices on Google products and more generally, online. For example, Google's Privacy Policy (available at https://policies.google.com/privacy) includes short, educational videos about the type of data Google collects.

Question 16. Does Google design its algorithms to make predictions about each of its users?

Response. There are indeed some places in our products where we endeavor to make predictions about users in order to be more helpful, for example in our Maps products we might suggest that a user plan to leave early for a trip to the airport depending on the user's settings and the data we have. Specifically, this might happen when the user has received an email confirmation from an airline suggesting the user may be flying that day; combining this with traffic data that shows an accident has stalled traffic on a nearby road may trigger us to prompt the user to leave early to allow for additional traffic.

As described in response to other answers, we offer a number of resources to help users better understand the products and services we provide including our uses of data. These resources include plain-English and easy-to-understand instructions about how users can make meaningful privacy and security choices on Google products and more generally, online. For example, Google's Privacy Policy (available at https://policies.google.com/privacy) includes short, educational videos about the type of data Google collects.

Question 17. Does Google design its algorithms to select and display content on its Search service in a manner that seeks to optimize user engagement?

Response. The purpose of Google Search is to help users find the information they are looking for on the web. Keeping them on the Google Search results page is not our objective.

Question 18. Does Google design its algorithms to select and display content on its YouTube service in a manner that seeks to optimize user engagement?

Response. We built our YouTube recommendation system to help users find new content, discover their next favorite creator, or learn more about the world. We want to provide more value to our users, and we work hard to ensure that we only recommend videos that will create a satisfying and positive user experience.

We update our systems continuously, and have been focusing on information quality and authoritativeness, particularly in cases like breaking news, or around sensitive or controversial topics. In January of this year, we announced the latest of our improvements to our recommendation system is to greatly reduce recommendations of borderline content and content that could misinform users in harmful ways. In June, we launched new features that give users more control over what recommendations appear on the homepage and in their 'Up Next' suggestions. These features make it easier for users to block channels from recommendations, give users the option to filter recommendations on Home and on Up Next, and give users more information about why we are suggesting a video.

Question 19. Does Google design its algorithms to select and display content on its News service in a manner that seeks to optimize user engagement?

Response. The algorithms used for our news experiences are designed to analyze hundreds of different factors to identify and organize the stories journalists are covering, in order to elevate diverse, trustworthy information.

Question 20. Tristan Harris, a witness at this hearing who was formerly an employee of Google, stated that what we're experiencing with technology is an increasing asymmetry of power between internet platforms and users, and that internet platforms like Google essentially have a supercomputer pointed at each user's brain that can predict things about the user that the user does not even know about themselves.

a.   Does Google agree that there is an asymmetry of power between it and its users?

Response. Users have transparency, choice, and control when it comes to how they use our platforms, and what information they choose to provide to us in order for us to customize their user experience. Users are in control of how they use our products, and if we do not earn their trust, they will go elsewhere.

b.   What predictions does Google seek to make about each user?

There are indeed some places in our products where we endeavor to make predictions about users in order to be more helpful, for example in our Maps products we might suggest that a user plan to leave early for a trip to the airport depending on the user's settings and the data we have. Specifically, this might happen when the user has received an email confirmation from an airline suggesting the user may be flying that day; combining this with traffic data that shows an

accident has stalled traffic on a nearby road may trigger us to prompt the user to leave early to allow for additional traffic.

We offer a number of resources to help users better understand the products and services we provide including our uses of data. These resources include plain-English and easy-to-understand instructions about how users can make meaningful privacy and security choices on Google products and more generally, online. For example, Google's Privacy Policy (available at https://policies.google.com/privacy) includes short, educational videos about the type of data Google collects.

c.   Does Google agree with Tristan Harris's characterization that internet platforms like Google essentially have a supercomputer pointed at each user's brain?

No, we do not agree with that characterization. We work hard to provide search results that are relevant to the words in a user's search, and with some products, like YouTube, we are clear when we are offering recommendations based on a user's preferences, but users retain control through their settings and controls to optimize their own experience.

Question 21. Does Google seek to optimize user engagement?

Response. We seek to optimize user experience. We have a multitude of tools and options to help our users interact with our products and platforms in ways that work best for them. We are committed to keeping our users safe online, and providing them with positive experiences. We do this through technological innovation, strong community guidelines, extensive education and outreach, and providing our users with choice, transparency and control over their experience. Our Digital Wellbeing Initiative focuses on these issues. More information about how we help our users find the balance with technology that feels right to them can be found on our Digital Wellbeing site.

Question 22. How does Google optimize for user engagement?

Response. As mentioned above in question 21, we optimize for user experience rather than user engagement, and give our users a number of tools to control their use of our platforms through our Digital Wellbeing product features. We continue to invest in these efforts to help users find the balance with technology that is right for them.

Question 23. How does Google personalize search results for each of its users?

Response. Search does not require personalization in order to provide useful organic search results to users' queries. In fact, there is very little personalization in organic 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 organic Search results from one person to another. Most differences that users see between their organic Search results and those of another user typing the same Search query are better explained by other factors such as a user's location.

For instance, if a user in Chicago searches for "football", Google will most likely show results about American football first. Whereas if the user searches "football" in London, Google will

rank results about soccer higher. Overall, Google strives to make sure that our users have access to a diversity of websites and perspectives.

Anyone who doesn't want personalization using account-based activity can disable it using the Web & App Activity setting. Users can also choose to keep their search history stored but exclude Chrome and app activity. "Incognito" search mode or a similar private browsing window can also allow users to conduct searches without having account-based activity inform their search results.

Search ads are ranked in a similar manner to organic Search results. The match between a user's search terms and the advertisers' selected keywords is the key factor underlying the selection of ads users see.

In relation to Google Ads, users can turn off personalized ads at myaccount.google.com. Once they've turned off personalization, Google will no longer use Account information to personalize the user's ads. Ads can still be targeted with info like the user's general location or the content of the website they are visiting.

Question 24. How does Google personalize what content it recommends for its users to see on YouTube?

Response. A user's activity on YouTube, Google and Chrome may influence their YouTube search results, recommendations on the Home page, in-app notifications and suggested videos among other places.

There are several ways that users can influence these recommendations and search results. They can remove specific videos from their watch history and queries from their search history, pause their watch and search history, or start afresh by clearing their watch and search history.

Question 25. How does Google personalize what content its users see on its News service?

Response. Whether our users are checking in to see the top news of the day or looking to dive deeper on an issue, we aim to connect them with the information they're seeking, in the places and formats that are right for them. To this end, Google provides three distinct but interconnected ways to find and experience the news across our products and devices: top news stories for everyone, personalized news, and additional context and perspectives.

1. Top News for everyone: For users who want to keep up with the news, they need to know what the important stories are at any point in time. With features such as Headlines in Google News and Breaking News on YouTube, we identify the major stories news sources are covering. This content is not personalized to individuals, but does vary depending on region and location settings. Google's technology analyzes news across the web to determine the top stories for users with the same language settings in a given country, based primarily on what publishers are writing about. Once these stories are identified, algorithms then select which specific articles or videos to surface and link to for each story, based on factors such as the prominence and freshness of the article or video, and authoritativeness of the source.

2.   Personalized news: Several Google news experiences show results that are personalized for our users. These include Discover, For you in Google News, and the Latest tab of the YouTube app on TVs. Our aim is to help our users stay informed about the subjects that matter to them, including their interests and local community. Google relies on two main ways to determine what news may be interesting to our users. In the experiences mentioned above, users can specify the topics, locations, and publications they're interested in, and they will be shown news results that relate to these selections. Additionally, depending on their account settings, our algorithms may suggest content based on a user's past activity on Google products. Algorithms rank articles based on factors like relevance to their interests, prominence and freshness of the article, and authoritativeness of the source. Google's news algorithms do not attempt to personalize results based on the political beliefs or demographics of news sources or readers. Users can control what account activity is used to customize their news experiences, including adjusting what data is saved to their Google account, at myaccount.google.com. In some Google products, such as Google News and Discover, users can also follow topics of interest, follow or hide specific publishers, or tell us when they want to see similar articles more or less frequently.

3.   Additional contexts and perspectives: A central goal of Google's news experiences is to provide access to context and diverse perspectives for stories in the news. By featuring unpersonalized news from a broad range of sources, Google empowers people to deepen their understanding of current events and offers an alternative to exclusively personalized news feeds and individual sources that might only represent a single perspective.

a.   Search experiences: When users search for something on Google, they have access to information and perspectives from a broad range of publishers from across the web. If they search for a topic that's in the news, their results may include some news articles labeled "Top stories" at the top of the results, featuring articles related to the search and a link to more related articles on the News tab. Users can also search for news stories and see context and multiple perspectives in the results on news.google.com, news on the Assistant, and within the "Top News" section of search results on YouTube. These results are not personalized. Our algorithms surface and organize specific stories and articles based on factors like relevance to the query, prominence and freshness of the article, and authoritativeness of the publisher. Users can always refine the search terms to find additional information.

b.   In-product experiences: In some news experiences, such as "Full coverage" in Google News, we show related articles from a variety of publishers alongside a given article. These results are not personalized. In providing additional context on a story, we sometimes surface videos, timelines, fact check articles, and other types of content. Algorithms determine which articles to show, and in which order, based on a variety of signals such as authoritativeness, relevance, and freshness.

Question 26. Does Google engage in any effort to change its user's attitudes? [response below]

Question 27. Does Google engage in any effort to change its user's behaviors? [response below]

Question 28. Does Google engage in any effort to influence its users in any way? [response below]

Question 29. Does Google engage in any effort to manipulate its users in any way? [response below]

Question 30. Do rankings of search results provided by Google have any impact on consumer attitudes, preferences, or behavior?

Response. We answer questions 26, 27, 28, 29, and 31 together. When users come to Google Search, our goal is to connect them with useful information as quickly as possible. That information can take many forms, and over the years the search results page has evolved to include not only a list of blue links to pages across the web, but also useful features to help users find what they're looking for even faster. For our Knowledge Graph allows us to respond to queries like "Bessie Coleman" with a [Knowledge Panel] with facts about the famous aviator. Alternatively, in response to queries like "how to commit suicide", Google has worked with the National Suicide Prevention Hotline to surface a results box at the top of the search results page with the organization's phone number and website that can provide help and support. The goal of this type of result is to connect vulnerable people in unsafe situations to reliable and free support as quickly as possible.

For other questions, Search is a tool to explore many angles. We aim to make it easy to discover information from a wide variety of viewpoints so users can form their own understanding of a topic. We feel a deep sense of responsibility to help all people, of every background and belief, find the high-quality information they need to better understand the topics they care about and we try to make sure that our users have access to a diversity of websites and perspectives.

When it comes to the ranking of our search results – the familiar "blue links" of web page results – the results are determined algorithmically. We do not use human curation to collect or arrange the results on a page. Rather, we have automated systems that are able to quickly find content in our index – from the hundreds of billions of pages we have indexed by crawling the web – that are relevant to the words in the user's search. To rank these, our systems take into account a number of factors to determine what pages are likely to be the most helpful for what a user is looking for. We describe this in greater detail in our [How Search Works] site.

Question 31. The website moz.com tracks every confirmed and unconfirmed update Google makes to its search algorithm. In 2018, Google reported 3,234 updates. However, moz.com reported that there were also at least six unconfirmed algorithm updates in 2018. Does Google publicly report every change it makes to its search algorithm? If not, why not?

Response. We report the number of changes we make to Google Search each year on our [How Search Works website]. To prevent bad actors from gaming our systems, we do not publicly report on the nature of each change.

Question 32. Does an item's position in a list of search results have a persuasive impact on a user's recollection and evaluation of that item?

Response. We aim to make it easy to discover information from a wide variety of viewpoints so users can form their own understanding of a topic. We feel a deep sense of responsibility to help all people, of every background and belief, find the high-quality information they need to better

understand the topics they care about and we try to make sure that our users have access to a diversity of websites and perspectives.

When it comes to the ranking of our search results – the familiar "blue links" of web page results – the results are determined algorithmically. We do not use human curation to collect or arrange the results on a page. Rather, we have automated systems that are able to quickly find content in our index – from the hundreds of billions of pages we have indexed by crawling the web – that are relevant to the words in the user's search. To rank these, our systems take into account a number of factors to determine what pages are likely to be the most helpful for what a user is looking for. We describe this in greater detail in our How Search Works site.

Question 33. A study published in 2015 in the Proceedings of the National Academy of Sciences entitled "The Search Engine Manipulation Effect (SEME) and its Possible Impact on the Outcomes of Elections" discussed an experiment where the study's authors (one of whom is a former editor in chief of Psychology Today) sought to manipulate the voting preferences of undecided eligible voters throughout India shortly before the country's 2014 national elections. The study concluded that the result of this and other experiments demonstrated that (i) biased search rankings can shift the voting preferences of undecided voters by 20% or more, (ii) the shift can be much higher in some demographic groups, and (iii) search ranking bias can be masked so that people show no awareness of the manipulation. This is a rigorously peer-reviewed study in the Proceedings of the National Academy of Sciences of the United States of America, one of the world's most-cited scientific journals, which strives to publish only the highest quality scientific research. Has Google carefully reviewed this study and taken steps to address the conclusions and concerns highlighted in this study? If so, please describe the steps taken to address this study. If Google has not taken steps to address this study, please explain why not?

Response. Google takes these allegations very seriously. Elections are a critical part of the democratic process and Google is committed to helping voters find relevant, helpful, and accurate information. Our job — which we take very seriously — is to deliver to users the most relevant and authoritative information out there. And studies have shown that we do just that. It would undermine people's trust in our results, and our company, if we were to change course. There is absolutely no truth to Mr. Epstein's hypothesis. Google is not politically biased and Google has never re-ranked search results on any topic (including elections) to manipulate user sentiment. Indeed, we go to extraordinary lengths to build our products and enforce our policies in an analytically objective, apolitical way. We do so because we want to create tools that are useful to all Americans. Our search engine and our platforms reflect the online world that is out there.

We work with external Search Quality Evaluators from diverse backgrounds and locations to assess and measure the quality of search results. Any change made to our Search algorithm undergoes rigorous user testing and evaluation. The ratings provided by these Evaluators help us benchmark the quality of our results so that we can continue to meet a high bar for users of Google Search all around the world. We publish our Search Quality Evaluator Guidelines and make them publicly available on our How Search Works website.

On Google Search, we aim to make civic information more easily accessible and useful to people globally as they engage in the political process. We have been building products for over a decade that provide timely and authoritative information about elections around the world and help voters make decisions that affect their communities, their cities, their states, and their countries. In 2018, for example, we helped people in the US access authoritative information about registering to vote, locations of polling places, and the mechanics of voting. We also provided information about all US congressional candidates on the Search page in Knowledge Panels, and provided the opportunity for those candidates to make their own statements in those panels. On election day, we surfaced election results for US congressional races directly in Search in over 30 languages. We have also partnered with organizations like the Voting Information Project, with whom we've worked since 2008 to help millions of voters get access to details on where to vote, when to vote, and who will be on their ballots. This project has been a collaboration with the offices of 46 Secretaries of State to ensure that we are surfacing fresh and authoritative information to our users.

In addition to Search results about election information, we have made voting information freely available through the Google Civic Information API, which has allowed developers to create useful applications with a civic purpose. Over 400 sites have embedded tools built on the Civic Information API; these include sites of candidates, campaigns, government agencies, nonprofits, and others who encourage and make it easier for people to get to the polls.