



HOUSE OF COMMONS
CHAMBRE DES COMMUNES
CANADA

Standing Committee on Justice and Human Rights

JUST • NUMBER 155 • 1st SESSION • 42nd PARLIAMENT

EVIDENCE

Tuesday, June 4, 2019

—
Chair

Mr. Anthony Housefather

Standing Committee on Justice and Human Rights

Tuesday, June 4, 2019

• (1610)

[Translation]

The Chair (Mr. Anthony Housefather (Mount Royal, Lib.)):
We apologize for being late.

[English]

We had a vote in the chamber. I'm sorry for being late, especially to our witness.

Good afternoon, everyone, and welcome to the Standing Committee on Justice and Human Rights, as we resume our study of online hate.

Today it is an enormous pleasure to be joined by Colin McKay, head of government affairs and public policy at Google Canada. We really appreciate Google's participation and yours to enable us to have a better study. Thank you so much.

Mr. McKay, the floor is yours.

Mr. Colin McKay (Head, Government Affairs and Public Policy, Google Canada): Thank you, Chair.

Thank you to all members of the committee for the opportunity to speak with you today.

I don't mind the delay. It's the business of Parliament, and I'm just happy to be a part of it today.

As the chair just mentioned, my name is Colin McKay, and I'm the head of government affairs and public policy for Google in Canada.

We, like you, are deeply troubled by the increase in hate and violence in the world. We are alarmed by acts of terrorism and violent extremism like those in New Zealand and Sri Lanka. We are disturbed by attempts to incite hatred and violence against individuals and groups here in Canada and elsewhere. We take these issues seriously, and we want to be part of the solution.

At Google, we build products for users from all backgrounds who live in nearly 200 countries and territories around the world. It is essential that we earn and maintain their trust, especially in moments of crisis. For many issues, such as privacy, defamation or hate speech, local legislation and legal obligations may vary from country to country. Different jurisdictions have come to different conclusions about how to deal with these complex issues. Striking this balance is never easy.

To stop hate and violent extremist content online, tech companies, governments and broader society need to work together. Terrorism

and violent extremism are complex societal problems that require a response, with participation from across society. We need to share knowledge and to learn from each other.

At Google we haven't waited for government intervention or regulation to take action. We've already taken concrete steps to respond to how technology is being used as a tool to spread this content. I want to state clearly that every Google product that hosts user content prohibits incitement to violence and hate speech against individuals or groups, based on particular attributes, including race, ethnicity, gender and religion.

When addressing violent extremist content online, our position is clear: We are agreed that action must be taken. Let me take some time to speak to how we've been working to identify and take down this content.

Our first step is vigorously enforcing our policies. On YouTube, we use a combination of machine learning and human review to act when terrorist and violent extremist content is uploaded. This combination makes effective use of the knowledge and experience of our expert teams, coupled with the scale and speed offered by technology.

In the first quarter of this year, for example, YouTube manually reviewed over one million videos that our systems had flagged for suspected terrorist content. Even though fewer than 90,000 of them turned out to violate our terrorism policy, we reviewed every one out of an abundance of caution.

We complement this by working with governments and NGOs on programs that promote counter-speech on our platforms—in the process elevating credible voices to speak out against hate, violence and terrorism.

Any attempt to address these challenges requires international coordination. We were actively involved in the drafting of the recently announced Christchurch Call to Action. We were also one of the founding companies of the Global Internet Forum to Counter Terrorism. This is an industry coalition to identify digital fingerprints of terrorist content across our services and platforms, as well as sharing information and sponsoring research on how to best curb the spread of terrorism online.

I've spoken to how we address violent extremist content. We follow similar steps when addressing hateful content on YouTube. We have tough community guidelines that prohibit content that promotes or condones violence against individuals or groups, based on race, ethnic origin, religion, disability, gender, age, nationality, veteran status, sexual orientation or gender identity. This extends to content whose primary purpose is inciting hatred on the basis of these core characteristics. We enforce these guidelines rigorously to keep hateful content off our platforms.

We also ban abusive videos and comments that cross the line into a malicious attack on a user, and we ban violent or graphic content that is primarily intended to be shocking, sensational or disrespectful.

Our actions to address violent and hateful content, as is noted in the Christchurch call I just mentioned, must be consistent with the principles of a free, open and secure Internet, without compromising human rights and fundamental freedoms, including the freedom of expression. We want to encourage the growth of vibrant communities, while identifying and addressing threats to our users and their broader society.

We believe that our guidelines are consistent with these principles, even as they continue to evolve. Recently, we extended our policy dealing with harassment, making content that promotes hoaxes much harder to find.

What does this mean in practice?

From January to March 2019, we removed over 8.2 million videos for violating YouTube's community guidelines. For context, over 500 hours of video are uploaded to YouTube every minute. While 8.2 million is a very big number, it's a smaller part of a very large corpus. Now, 76% of these videos were first flagged by machines rather than humans. Of those detected by machines, 75% had not received a single view.

We have also cracked down on hateful and abusive comments, again by using smart detection technology and human reviewers to flag, review and remove hate speech and other abuse in comments. In the first quarter of 2019, machine learning alone allowed us to remove 228 million comments that broke our guidelines, and over 99% were first detected by our systems.

We also recognize that content can sit in a grey area, where it may be offensive but does not directly violate YouTube's policies against incitement to violence and hate speech. When this occurs, we have built a policy to drastically reduce a video's visibility by making it ineligible for ads, removing its comments and excluding it from our recommendation system.

Some have questioned the role of YouTube's recommendation system in propagating questionable content. Several months ago we introduced an update to our recommendation systems to begin reducing the visibility of even more borderline content than can misinform users in harmful ways, and we'll be working to roll out this change around the world.

It's vitally important that users of our platforms and services understand both the breadth and the impact of the steps we have taken in this regard.

We have long led the industry in being transparent with our users. YouTube put out the industry's first community guidelines report, and we update it quarterly. Google has long released a transparency report with details on content removals across our products, including content removed upon request from governments or by order from law enforcement.

While our users value our services, they also trust them to work well and provide the most relevant and useful information. Hate speech and violent extremism have no place on Google or on YouTube. We believe that we have developed a responsible approach to address the evolving and complex issues that have seized our collective attention and that are the subject of your committee's ongoing work.

Thank you for this time, and I welcome any questions.

• (1615)

The Chair: Thank you very much for your opening statement.

We will go to Mr. MacKenzie.

Mr. Dave MacKenzie (Oxford, CPC): If I don't use all my time, Mr. Chair, Mr. Barrett will take it.

The Chair: Absolutely.

Mr. Dave MacKenzie: Thank you for being here today, Mr. McKay.

You're in an enviable position of trying to harness whatever is going on in the world through your medium. I wonder what you would define as "hateful messages".

Mr. Colin McKay: If you'll permit me to look at my notes, I have a very specific definition.

For us, hate speech refers to content that promotes violence against, or has the primary purpose of inciting hatred against, individuals or groups based on the attributes I mentioned in my opening remarks.

Mr. Dave MacKenzie: When we have that definition and somebody puts something on YouTube that may come from a movie or a television show, it seems to me as though, at times, those topics would be part of the broadcast. If those show up, what happens?

Mr. Colin McKay: If those show up and they are flagged for review—a user flags them or they're spotted by our systems—we have a team of 10,000 who review videos that have been flagged to see if they violate our policies.

If the context is that something is obviously a clip from a movie or a piece of fiction, or it's a presentation of an issue in a particular way, we have to carefully weigh whether or not this will be recognized by our users as a reflection of cultural or news content, as opposed to something that's explicitly designed to promote and incite hatred.

Mr. Dave MacKenzie: A couple of weeks ago a group of youngsters attacked and beat a woman in a park. I believe only one was 13; I think the rest of them were young. It showed up on the news. Would that end up in a YouTube video?

Mr. Colin McKay: Speaking generally and not to that specific instance, if that video were uploaded to YouTube, it would violate our policies and would be taken down. If they tried to upload it again, we would have created a digital fingerprint to allow us to automatically pull it down.

The context of how a video like that is shown in news is a very difficult one. It's especially relevant not just to personal attacks, but also to terrorist attacks. In some ways, we end up having to evaluate what a news organization has determined is acceptable content. In reviewing it, we have to be very careful that it's clear to the viewer that this is part of a commentary either by a news organization or another organization that places that information in context.

Depending on the length and type of the video, it may still come down.

Mr. Dave MacKenzie: Okay. I appreciate that, because one of the things that I think did occur was that it showed up over and over on news. If you say that Google doesn't accept that video in YouTube, I think that's very appropriate. I'm not sure how we deal with it in everyday newscasts, so in some respects, I think you're ahead of where we are with the news.

Is that equally true of other mediums? We're talking about videos. Can somebody google—as opposed to YouTube—some hateful speech that takes place? I don't know if “censor” is the right word, but do you have a means to take it down or locate it?

Mr. Colin McKay: I was being very specific about YouTube, because that's somewhere that people consciously upload to our platform. In “search” we apply similar processes, but they have to be broader. We are just providing answers to questions that are posed to us by users about specific instances and events. If a news organization is presenting information in this way, it will surface in a normal way within our systems.

With specific elements, if there are specific speeches or pieces of content that are illegal within a country, then those will be taken down. We follow the law and legislation of the countries within which operate.

• (1620)

Mr. Dave MacKenzie: Okay. Thank you.

I'll pass to Mr. Barrett.

The Chair: You have two minutes, Mr. Barrett.

Mr. Michael Barrett (Leeds—Grenville—Thousand Islands and Rideau Lakes, CPC): Thanks.

Mr. McKay, Mr. MacKenzie satisfied my curiosity as it relates to this, but I do have a question. If Google has said this year that it is not planning to allow political ads, is it the intent of Google to allow political ads in the next election? Or is the current plan just not workable ever?

Mr. Colin McKay: We were faced with a difficult decision. The legislation was passed in December, and we had to have a system in place for the end of June. We went through the evaluation internally

as to whether or not we could take political ads in Canada within that time frame, and it just wasn't workable.

The reality around transparency in political ads is that we already have products in the United States—and we're rolling them out elsewhere—that provide transparency around political advertising. Those products are evolving as we go through election after election. Europe just had one. India just had one. Brazil just had one. Our goal is to continue developing those products to a point where we hope it will reach parity with what's identified in the Elections Act.

Mr. Michael Barrett: Do you have an expectation for a timetable as to when, based on the current legislation, you think Google would be able to comply?

Mr. Colin McKay: No.

Mr. Michael Barrett: Okay.

My next question has to do with public safety. Rural Canadians have expressed concerns that mapping software that often relies on Google Maps doesn't identify rural streets. That can pose problems for emergency services. Is there a mechanism or a plan for a mechanism to be made available to rural Canadians to be able to identify to Google either missing streets or missing mapping data for instances like the one I mentioned?

Mr. Colin McKay: They can right now; on Google Maps, you can use the feedback mechanism to identify a particular element of the map. You can identify whether that road is closed indefinitely and it just isn't marked on the map, or whether there has been development since we last mapped that area and there is now a municipal building or some other facility that needs further recognition. They can send those signals to us through the mapping product. That is actually the quickest way to do it. Those feedback comments go directly to the mapping team and then are evaluated for inclusion.

Mr. Michael Barrett: Thanks very much for your answers to my questions.

The Chair: Thank you very much.

Mr. McKinnon.

Mr. Ron McKinnon (Coquitlam—Port Coquitlam, Lib.): Thank you for being here. I'm very interested to hear about the AI work you're doing to track down malicious content and so forth. I'm interested more particularly in tracking the provenance of such content. I submit that anonymity can be a big problem in encouraging bad behaviour online. I understand that Google has a very broad universe in which it operates. It has many different products.

I'm most particularly interested in commentary. I'm wondering whether Google has considered not necessarily requiring users to be authenticated, whether by an authentication authority such as Verisign or by more homegrown approaches such as webs of trust like PGP...and identifying people with an icon of some kind to indicate whether or not these people are authenticated. The next part of that would be to allow them to filter out content that came from unauthenticated sources. Do you have any comments on that?

Mr. Colin McKay: I have a two-part response if you'll be patient with me. I think the first is that if we're speaking specifically about YouTube and a platform where you're able to upload information, there isn't a process of verification/authentication, but you do need to provide some reference points for yourself as an uploader. This can be limited to an email address and some other data points, but it does create a bit of a marker, especially for law enforcement who may want to track back the behaviour of a particular video uploader.

One area we focus on, though, is that we're very conscious that many users rely on anonymity or pseudonymity to be able to take positions, especially in politically sensitive or socially heightened environments, particularly if they're advocates of a particular position using our platforms. The process of verification/authentication in those circumstances is actually detrimental to them.

What I will speak to is that in responding to incidents of hate and online violent extremist content, we have made conscious efforts both in Google Search and our Google News product, as well as YouTube, in the moments after a crisis especially, when there isn't a reliable, factual content available about the immediate crisis, to focus, as our responsibility, on the authenticity and authority of those sources that are reporting and commenting on the crisis.

Within our systems, particularly in YouTube, you will see that if you're looking at a particular incident, the other material that is recommended to you comes from reliable sources that you likely have had contact with before. We try to send those signals. In addition to making information that's relevant to your query available, we're trying to make it clear that we're also trying to provide that level of reassurance, if not certainty.

• (1625)

Mr. Ron McKinnon: I understand your point about anonymity being sometimes desirable, and many people might need it in certain circumstances.

What I'm looking for as an end-user is to be able to, say, exclude from my feed, by my choice, content that was not authenticated, perhaps. Right?

Mr. Colin McKay: To further refine my point, it's important to realize that YouTube is a platform for content creators, so essentially the people who are on YouTube are people who are trying to build an audience and a community around shared interests. The vast majority of them have given you the information to be able to verify and have some level of certainty about their authority if not their actual identity, whether you're talking about repairing small engines, doing model trains or political commentary, and whether you're talking about traditional news organizations or purely online news organizations.

For the vast material or content, you are operating in an environment where you are able to identify and then qualify what

you are looking at. A lot of my opening statement, as well as a lot of our online and automated processes, are based on that immediate response to a crisis or that attempt to incentivize violent extremism or hatred online, where we're trying to do exactly what you're trying to describe. That is to say, wait a second, what's the outlier, what's the uploader, who is the user of our service trying to pursue a negative outcome, and can we identify them and qualify them and then apply our policies against them, whether it means limiting the availability of those videos or taking them down from the system?

We are trying to pursue that goal, not just within the context of authentication.

Mr. Ron McKinnon: I've got a few seconds left.

Perhaps you can talk to us about bots. I understand that sometimes they're beneficial and other times they're malicious. I understand you can identify when a bot is at play. Perhaps you could tell us about standards and what action could be taken to control them.

Mr. Colin McKay: I think from our point of view, I'd back away from bots to a wider perspective, which is that, across our system, we've long had experience with automated attacks on our infrastructure as well as our services. What we have focused on over time is providing the signals to our users that they are being subjected to an automated attack that is trying to either compromise their Google account, their Gmail account, or to present misinformation or disinformation to them. That goes all the way back to providing notices to users that they could be subject to a state-sponsored attack on their Gmail account.

Through this sort of deep-level analysis that I described, which analyzes videos writ more broadly across our infrastructure, we are trying both to identify when we see systemic attempts to breach the security of our systems and also to raise the profile and popularity of content, whether it's on search or whether it's on YouTube, to battle that. From our point of view, it's a very different context from the other services, but it's something in which we've historically invested a lot of money and time in both combatting and then also providing flags to our users so they're aware that they're being subject to these attacks or that there's an attempt to try to influence them in this way.

Mr. Ron McKinnon: Thank you.

[Translation]

The Chair: Thank you.

Ms. Moore, you have the floor.

Ms. Christine Moore (Abitibi—Témiscamingue, NDP): Thank you.

I'd like to invite my colleagues to pick up their phones, if they wish.

[English]

If you do a quick Google search of the phrase “how to pimp,” there are countless videos available that inform people of how to take part in human trafficking. Why haven't these sites and videos been removed?

Mr. Colin McKay: I'm sorry. What was the phrase?

• (1630)

Ms. Christine Moore: How to pimp.

Mr. Colin McKay: How to pimp, okay.

The reality is that we're constantly fighting against issues and content like this. It's a constant effort to identify the context within which there are comments in a video, and then to create what we call a “classifier”, which is an automated system to identify them on a broad scale.

If you see content like that, there is the opportunity to flag it right there on the YouTube video on your mobile device or the page. We use that as a signal to recognize that, wait a second, there is behaviour here and content that needs to be removed. Obviously, that's something we don't want in our system.

Ms. Christine Moore: Okay. So when you remove those videos and websites explaining how to pimp, is there any information that is transmitted, for example, to police forces or local authorities? If someone has a complete guide explaining how to engage in human trafficking, I think it involves criminal activity. Do you flag the police of that country, for example, and say, “Maybe that guy is involved in something criminal, and you could take a look,” or do you just remove the video because there are just too many of them all the time and you don't have the time to follow up?

Mr. Colin McKay: I'm not sure in this case. I can follow up with you.

[Translation]

Ms. Christine Moore: Okay.

In addition, the English keywords are accurate, because most people use that language. However, what about algorithms in other languages that may be used less often, such as French? Is it easier to spread hate if you use a language other than English, because the algorithms for keywords are not as well developed?

[English]

Mr. Colin McKay: I think we have addressed that broadly, because we've long focused on providing our services in many languages. We have actually developed the artificial intelligence translation systems to be able to translate upwards of 200 languages. Our systems aren't focused solely on English terms and English challenges. It's broader. It's international, and the review teams that I described are also international.

Our team that reviews content is made up of of 10,000 people distributed around the world, specifically so that they can have the linguistic, cultural and societal background to understand the context within which they are seeing comments and material, and making decisions about whether or not the particular content or account needs to be taken down.

We recognize that challenge and we're still using a combination of automated processes that use some of the best individual language specialists and language translation software in the world to filter into the process.

Ms. Christine Moore: How quickly are you able to remove a video that has already been removed and been modified, for example, using sound that goes faster. They do that. I've often seen that from my daughter. There are people producing *Paw Patrol* a little faster so that it is not recognized by the system and they are able to publish their video.

In terms of hate videos, are you able to quickly remove a video that has already been removed once and has been modified just to avoid those controls?

Mr. Colin McKay: Yes, we are.

I recognize the example you described. I've seen that as well. That is one of the challenges, especially immediately after a crisis. We're seeing content being uploaded and they are playing with it a little bit to try to confuse our systems.

What we do, particularly in the case of hate content and violent content online, is to tighten the standards within which we identify videos so that we're taking them down even more quickly.

Even in the context of *Paw Patrol*, I think your daughter will likely find that if she goes back to the same channel two weeks later, they may not have the *Paw Patrol* content because it will have been recognized and taken down.

[Translation]

The Chair: You have one minute left.

[English]

Ms. Christine Moore: Okay.

I would like to know a little bit more about the process of reviewing flagged videos, and who reviews them when it's not done by a computer.

Also, are the workers reviewing these videos provided with any services, because having to listen to these kinds of things all the time causes a lot of distress to people? What services are you providing to these workers to make sure they do not go crazy from listening to all of these things all the time?

Mr. Colin McKay: To begin with the process itself, as I mentioned, especially in the context of hate content, we are dealing with such a quantity that we rely on our machine learning and image classifiers to recognize content. If the content has been recognized before and we have a digital hash of it, we automatically take it down. If it needs to be reviewed, it is sent to this team of reviewers.

They are intensely trained. They are provided with local support, as well as support from our global teams, to make sure they are able to deal with the content they're looking at and also the needed supports. That is so that as they look at what can be horrific content day after day, they are in a work environment and a social environment where they don't face the same sorts of pressures that you're describing. We are very conscious that they have a very difficult job, not just because they're trying to balance rights versus freedom of expression versus what society expects to find when online, but also because they have the difficult job of reviewing material that others do not want to review.

For us, whether they're based in one office or another around the world, we are focused on giving them training and support so they can do their job effectively and have work-life balance.

•(1635)

The Chair: Thank you very much.

Ms. Khalid.

Ms. Iqra Khalid (Mississauga—Erin Mills, Lib.): Thank you, Mr. McKay, for coming in today.

I'm going to follow up Madame Moore's line of questioning.

How many reviewers do you have to review specifically Canadian content within Google Canada?

Mr. Colin McKay: We have a global team that doesn't treat the content by jurisdiction or region. Depending on what the pressure point may be or where the flow of content may be coming from, they will deal with that as a flow.

Where they get their insight and their expertise on Canada is in part from guidance from my team and my colleagues who work for Google in Canada. Also, we have a sophisticated mechanism for ensuring that the cultural, social and political context within which content is being reviewed is recognized within that review process.

Ms. Iqra Khalid: How long does it take you to remove something once it's reported or flagged to you? What's the specific timeline?

Mr. Colin McKay: It varies, depending on the context and the severity of the material.

We've already had examples in our conversation today about whether or not it's commentary or it's news reporting, or it's actual video of a violent attack. In the context of the Christchurch attack, we found that there were so many people uploading the videos so quickly that we had to accelerate our artificial intelligence review of the videos and make on-the-fly decisions about taking down video, based on its being substantially similar to previous uploads.

In that process, the manual review was shortened extremely because we were facing a quantity.... In a case where there's broader context to be considered, there's still a commitment to review it quickly, but we do need a process of deliberation.

Ms. Iqra Khalid: In your opening remarks, you spoke about different countries having different legislation. This is something we've heard before this committee, that our government needs to set requirements for providers such as yourself to remove all posts that

would constitute hate speech, and failure to do so in a timely manner would result in accountable action or significant fines, etc.

Can you talk a little about what other jurisdictions are doing? How do you keep your global team updated with all the varying legislation within the different countries?

Mr. Colin McKay: Sure.

First, speaking specifically to Europe, which has in place a code of conduct around hate speech and very clear reporting obligations, we've arrived at a point where 83.8% of the content that has been flagged for review is assessed in less than 24 hours, and 7.9% in less than 48 hours.

That gives you a bit of an idea of the window within which content that deals with hate can be reviewed appropriately. Obviously, from our point of view, we're trying to improve on that.

The way we work within this broad organization of 10,000 is that we have very clear-cut internal review, and established guidelines for those review teams, around what the expectations and obligations are within each jurisdiction and what is explicitly illegal, and then what we would consider borderline illegal that requires some level of intervention on YouTube to restrict access to that content.

Internal to the company, like any multinational, we have a team that's dedicated to identifying both the differences and the similarities, and ensuring that we are in compliance.

•(1640)

Ms. Iqra Khalid: Have you seen, through the work that you do in the very different countries you operate in, that some countries are more successful in curbing hate speech online through their legislation than others?

Mr. Colin McKay: This is an observation that's just off the top of my head, and it's personal. I would say that we are seeing a variety of efforts to deal with this challenge. They're based within, as I said, the social and political context of each country, and the level of immediacy and severity being applied to the issue reflect local pressures. The difficulty for us still remains understanding those social, economic and political pressures and the context within which we can interpret them, using our systems to deliver a result that's acceptable to those jurisdictions, governments and societies. From country to country, one thing we've seen is that, if there's a more coordinated and collaborative effort to arrive at complementary and similar approaches, if not shared principles and legislation, that effort can have a broader and more recognizable impact, especially for users.

I'll point to an example. You have a juxtaposition between New Zealand and Australia in reaction to the Christchurch attack, where the Prime Minister of New Zealand took on this approach to develop a call that brought in all the stakeholders to develop an aggressive approach to dealing with this, but not an immediate approach. Australia went the other way and implemented legislation, which, it was quickly realized, needs to be reconsidered in Parliament. That's not to say the intent and execution of that legislative process was wrong; it's just that it still needs further deliberation. I think that's the challenge we face. We're in the space now where, as I said, we all share concern, we all want to act on it, and we want to act on it in a way that has impact.

Ms. Iqra Khalid: Thank you.

You spoke about how you have one global team and that you train them. You also addressed the challenge of understanding the social and cultural factors within each country. Do you think that you would be better helped if you had teams in the countries you operate in who understand specifically the social and cultural factors that impact hate speech in that specific country?

Mr. Colin McKay: I think the first step is to have a clear idea of what the boundaries are for terms like “hate speech” and “violent extremist content”, because as a company, we're still interpreting and trying to define our perception of what society finds acceptable and what you as legislators in government find acceptable. The first step for us would be what a clear definition is, so that we can act upon it, because that's often where we have points of contention as to what exactly is the expectation around takedown and restriction or limiting access on content, especially if it's related to hate more than violent extremism.

Ms. Iqra Khalid: Do you think there's one common definition of hate speech across the world or common threads of a definition that you think we could work with internationally?

Mr. Colin McKay: I mentioned the definition that we act upon, and it's very broad. We find that a reliable reference point for our activities. Often it's in commentary and political discourse where it's challenging to interpret whether or not that line has been crossed. There are baseline documents that already exist on human rights and legal obligations that we certainly reference, and we speak regularly to both government and legislators as well as to NGOs to make sure

that we're aware of how that conversation has evolved and that we're filtering in the right way.

The Chair: Thank you so much.

I just have one question to follow up on Ms. Khalid's and Ms. Moore's question. How many of the 10,000 people who do the vetting are based in Canada?

Mr. Colin McKay: Very few.

• (1645)

The Chair: Would it be zero?

Mr. Colin McKay: It's not zero, no.

The Chair: Just to follow up the other question that Ms. Moore asked about the French language, I understand that you have translation software. Everybody's seen Google Translate. It's a great help to meet people in a baseline sense, but obviously that's not an effective way to understand the terminology used online. Do you have people with native language skills in all of these multiple languages who put in the search terms?

Mr. Colin McKay: I can't confidently say that for every language.

The Chair: Let's say French.

Mr. Colin McKay: In French, yes.

The Chair: Okay.

Your assistance here today and the fact that Google is willing to work with us in this way is incredibly appreciated, Mr. McKay. I really want to thank you.

We have another meeting, but it's in camera. What I would ask is that we take a five-minute suspension, and then we will ask everybody who shouldn't be here to clear the room.

[*Translation*]

I want to thank everyone.

We'll suspend the meeting for five minutes.

[*English*]

Mr. Colin McKay: Thank you very much.

The Chair: The meeting is adjourned.

Published under the authority of the Speaker of
the House of Commons

SPEAKER'S PERMISSION

The proceedings of the House of Commons and its Committees are hereby made available to provide greater public access. The parliamentary privilege of the House of Commons to control the publication and broadcast of the proceedings of the House of Commons and its Committees is nonetheless reserved. All copyrights therein are also reserved.

Reproduction of the proceedings of the House of Commons and its Committees, in whole or in part and in any medium, is hereby permitted provided that the reproduction is accurate and is not presented as official. This permission does not extend to reproduction, distribution or use for commercial purpose of financial gain. Reproduction or use outside this permission or without authorization may be treated as copyright infringement in accordance with the *Copyright Act*. Authorization may be obtained on written application to the Office of the Speaker of the House of Commons.

Reproduction in accordance with this permission does not constitute publication under the authority of the House of Commons. The absolute privilege that applies to the proceedings of the House of Commons does not extend to these permitted reproductions. Where a reproduction includes briefs to a Committee of the House of Commons, authorization for reproduction may be required from the authors in accordance with the *Copyright Act*.

Nothing in this permission abrogates or derogates from the privileges, powers, immunities and rights of the House of Commons and its Committees. For greater certainty, this permission does not affect the prohibition against impeaching or questioning the proceedings of the House of Commons in courts or otherwise. The House of Commons retains the right and privilege to find users in contempt of Parliament if a reproduction or use is not in accordance with this permission.

Also available on the House of Commons website at the following address: <http://www.ourcommons.ca>

Publié en conformité de l'autorité
du Président de la Chambre des communes

PERMISSION DU PRÉSIDENT

Les délibérations de la Chambre des communes et de ses comités sont mises à la disposition du public pour mieux le renseigner. La Chambre conserve néanmoins son privilège parlementaire de contrôler la publication et la diffusion des délibérations et elle possède tous les droits d'auteur sur celles-ci.

Il est permis de reproduire les délibérations de la Chambre et de ses comités, en tout ou en partie, sur n'importe quel support, pourvu que la reproduction soit exacte et qu'elle ne soit pas présentée comme version officielle. Il n'est toutefois pas permis de reproduire, de distribuer ou d'utiliser les délibérations à des fins commerciales visant la réalisation d'un profit financier. Toute reproduction ou utilisation non permise ou non formellement autorisée peut être considérée comme une violation du droit d'auteur aux termes de la *Loi sur le droit d'auteur*. Une autorisation formelle peut être obtenue sur présentation d'une demande écrite au Bureau du Président de la Chambre.

La reproduction conforme à la présente permission ne constitue pas une publication sous l'autorité de la Chambre. Le privilège absolu qui s'applique aux délibérations de la Chambre ne s'étend pas aux reproductions permises. Lorsqu'une reproduction comprend des mémoires présentés à un comité de la Chambre, il peut être nécessaire d'obtenir de leurs auteurs l'autorisation de les reproduire, conformément à la *Loi sur le droit d'auteur*.

La présente permission ne porte pas atteinte aux privilèges, pouvoirs, immunités et droits de la Chambre et de ses comités. Il est entendu que cette permission ne touche pas l'interdiction de contester ou de mettre en cause les délibérations de la Chambre devant les tribunaux ou autrement. La Chambre conserve le droit et le privilège de déclarer l'utilisateur coupable d'outrage au Parlement lorsque la reproduction ou l'utilisation n'est pas conforme à la présente permission.

Aussi disponible sur le site Web de la Chambre des communes à l'adresse suivante : <http://www.noscommunes.ca>