HOUSE OF COMMONS
CHAMBRE DES COMMUNES
CANADA

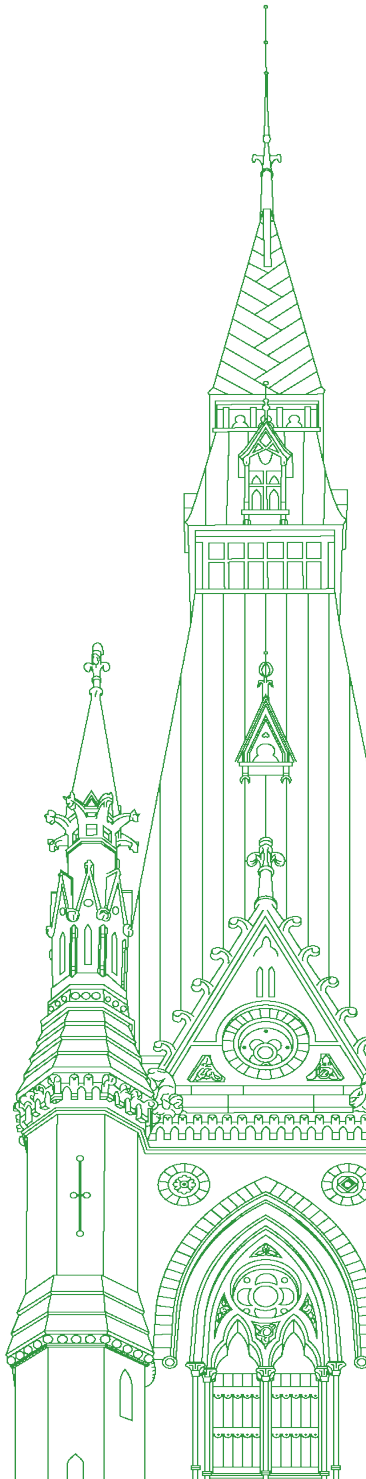
45th PARLIAMENT, 1st SESSION

Standing Committee on Procedure and House Affairs

EVIDENCE

NUMBER 043

Thursday, October 22, 2026



Chair: Chris Bittle

Standing Committee on Procedure and House Affairs

Thursday, October 22, 2026

● (1100)

[*English*]

The Chair (Chris Bittle (St. Catharines, Lib.)): I call this meeting to order.

Welcome to meeting number 43 of the House of Commons Standing Committee on Procedure and House Affairs.

Pursuant to Standing Order 108(3), the committee is meeting to resume its study of artificial intelligence and Canada's democratic institutions.

Today's meeting is taking place in public in a hybrid format, pursuant to the Standing Orders. Members are attending in person in the room and remotely using the Zoom application.

Before we continue, I would ask all in-person participants to consult the guidelines written on the cards on the table. These measures are in place to help prevent audio and feedback incidents and protect the health and safety of all participants, especially our interpreters.

I would like to remind the witnesses that committee members may ask questions in either English or French. If you need interpretation, please take a moment now to prepare your earpiece or select the listening channel on Zoom in advance to take full advantage of the time allotted for questions and answers.

I'd like to make a few comments for the benefit of members.

All comments should be addressed to the chair. If you wish to speak, raise your hand. The clerk and I will do our best to spot you and create a speaking list.

I'd like to welcome our witnesses for the first panel of today's meeting.

Appearing as individuals, we have Anamaria Crisan, assistant professor at the University of Waterloo, and Kevin Deveaux, barrister and solicitor, by video conference. We also have Sarah Lister, director of governance, rule of law and peacebuilding, from the United Nations Development Programme.

We'll turn to Ms. Lister first, since she's in the room, for five minutes, please.

Dr. Sarah Lister (Director of Governance, Rule of Law and Peacebuilding, United Nations Development Programme): Members of the committee, thank you for this opportunity. As I start, I would note that the information provided to this committee is offered on a voluntary basis and should not be understood as a waiver—express or implied—of the privileges and immunities of the United Nations and its officials under the 1946 Convention on the Privileges and Immunities of the United Nations.

The United Nations Development Programme is one of the world's largest providers of electoral assistance, supporting electoral processes in around 50 countries each year. It works with one in three parliaments globally. This gives us a distinctive vantage point on how AI is affecting democratic institutions and how those institutions are beginning to respond. It is from this perspective that I'm making these remarks.

Parliaments are central to deciding how the risks and benefits of AI are distributed, which uses are acceptable and in the public interest, and who is accountable.

First, through lawmaking, parliaments can establish enforceable rules. There is no single model that we see. The Republic of Korea has adopted a principles-based framework. Austria has an AI service centre. Lithuania has a regulatory sandbox. Other countries are adapting existing laws. A context-specific regulatory approach can reduce uncertainty while protecting rights and the public interest.

Second, parliaments are using oversight and budget scrutiny mechanisms, and committees such as this one are holding inquiries. Specialized bodies are being established, and budget reviews are testing whether countries have the workforce and safeguards needed for AI.

Third, parliaments can ensure that AI governance is informed by those most affected, including women, minorities and marginalized communities. Citizens' assemblies, deliberative polling, and consultation linked to formal democratic institutions and processes can help determine acceptable uses and build trust.

Fourth, parliaments should strengthen their own readiness and use AI responsibly. AI can improve transcription, translation, research and public access to information, but as in all public institutions, adoption should begin with a defined public problem—not with technology. Institutions need internal policies, sound data governance, risk assessment, staff capability and meaningful human authority over consequential decisions.

The democratic stakes are at their highest in electoral processes. AI can improve voter information, accessibility, translation, planning, data analysis and cybersecurity, but an error can affect the right to vote. An opaque system can impede challenge, and a cybersecurity failure can undermine confidence in a result. AI-generated misinformation, disinformation and hate speech can fuel online and off-line violence, disproportionately affecting women and marginalized groups, and it can damage the quality of public debate.

A recent UNDP survey of 53 electoral management bodies found that 74% were already using AI in at least one function, but more than half had no internal AI policy, and only 15% had dedicated AI-related infrastructure. That gap between adoption and governance is an immediate concern.

Last month, here in Ottawa, UNDP and Elections Canada convened some 200 electoral stakeholders from 60 countries at a global conference on responsible AI and elections. Their exchanges reinforced the need for institutional governance, shared learning and international co-operation.

The practical agenda is clear: Move from ad hoc experimentation to institutional governance; assess risks before deployment; invest in data governance, procurement and oversight; retain human authority to question outputs; and connect parliaments with regulators, audit institutions, civil society, academia and the private sector.

More than 60 national parliaments are already taking concrete action, yet many still lack the technical expertise and institutional structures to legislate and oversee effectively. International co-operation should help close that capacity gap.

The trajectory of AI is not technologically predetermined. Whether it deepens inequality or expands opportunity, and whether it strengthens or weakens democratic governance depends on political and societal choices. Parliaments are uniquely mandated to make those choices visible, accountable and in the public interest.

In that endeavour, UNDP greatly values its partnership with Canada.

Thank you for your attention.

● (1105)

The Chair: Thank you so much.

We'll now turn to Ms. Crisan for five minutes, please.

Dr. Anamaria Crisan (Assistant Professor, University of Waterloo, As an Individual): Thank you to the members of the committee for inviting me to take part in this important discussion.

My research focuses on building human-centred technologies for automating work. Put simply, that means designing AI systems around the needs, values and capabilities of people from the outset rather than considering those things only after a system fails or causes harm. I study how people understand, use and are affected by these systems. I develop new methods and interfaces for incorporating human input into how these systems operate.

For more than a decade, I have studied this across very different decision-making contexts, from how organizations understand their customers to deeply personal choices, such as deciding on a medical treatment when faced with a large amount of complex information. Across hundreds of these in-depth conversations, I have come back with three observations.

First, people have very different abilities to use and understand computer systems. This includes not only accessibility concerns but also the ability to understand the information they are being presented with. With AI, the challenge takes a new dimension, because language becomes the primary medium of interaction, making AI systems appear more humanlike. If AI systems are constituent-facing, some people may not understand that they are interfacing with an AI. This can affect what they disclose, how they interpret the response and whether they know when and how to reach a person. The reverse is also increasingly possible. AI systems can simulate constituents, generate large volumes of correspondence or make automated interactions appear to come from individual people. This can make it harder for public institutions to know whose voices they are actually hearing.

These scenarios are examples of the challenges that can arise when AI mediates the dialogue between constituents and the people meant to represent them. The risk is that the systems begin to narrow or distort what each side hears from the other, weakening the connection that democratic representation depends upon.

This leads to an important second point, which is that automation tends to flatten perspectives. Consider using AI to help sift through thousands of constituent emails. It could categorize them by topic or urgency, or it could produce a concise summary. This can be extremely useful, but those summaries are not neutral. An AI decides how to group correspondences, what central themes matter, what gets highlighted or even what's unusual. There is no gold standard reference that we can easily use to judge the correctness of this process. In addition, the AI-generated summaries can change with a simple revision of the prompt wording.

The same is true when AI helps us research, brainstorm or write legislation. The outputs can sound entirely plausible and reasonable while reflecting particular assumptions or perspectives that are difficult to trace. As we rely more heavily on them, there is a risk that we lose the critical engagement that allows us to notice whether certain perspectives are being prioritized.

Third, the levers that people have to influence these systems are often narrower than they appear. A person may be technically in the loop because they can write a prompt, review a summary or prove an output, but AI systems are designed and optimized to reduce the amount of human intervention required. At the same time, understanding why a system produced a particular result remains difficult. We're seeing a substantial burden to verify the factualness of these systems.

These observations have important implications for Canadian democratic institutions. One is the need for Canadian-specific auditing. We should test these systems in the context in which Canadians will actually encounter them. In particular, when AI is used within the service of our democratic institutions, we should audit not only whether the output is factually correct but also whose perspectives are preserved and whose are missing.

These observations also point to a broader challenge for governance. AI systems are difficult to fully interrogate because of their complexity, and our ability to control how they are used is more limited than we think. So far, much of the governance conversation has focused on frontier models and the companies, mostly U.S.-based, that develop and deploy them. That is important, but it addresses only one part of the ecosystem. Increasingly capable open-weight models can be downloaded, modified and redistributed. Some contemporary open-weight models are now substantially more capable than the first ChatGPT model, which already raised early governance concerns.

To be clear, I believe a healthy open-weight ecosystem can provide important public benefits. These models are attractive precisely because they can be run locally, customized for particular needs and used without relying entirely on a small number of commercial providers. However, these properties also make governance more complicated. Once these models are widely distributed, there are far fewer ways for the original developer to control how they are modified or used. Those modifications can also change the accuracy, biases and behaviour of the original system, sometimes subtly. That means governance cannot rely on controlling the technology itself.

● (1110)

We also need to think carefully about how people and institutions rely on it and how resilient they are when it fails. Part of that resilience is making sure that we do not lose sight of the human capacity to notice when something is wrong.

I'll conclude by saying this: Think of hospitals and their paper protocols when digital systems go down. The point is not to avoid technology; it is to make sure that the institution can still function when technology does not. Our public institutions need the same kind of resilience.

Thank you.

The Chair: Thank you so much.

Mr. Deveaux, you have five minutes please.

Kevin Deveaux (Barrister and Solicitor, As an Individual): Thank you, Mr. Chairperson.

Thank you for this opportunity to speak before the committee.

By way of background, quickly, before I get into the substance of my submission, I want to let you know that the agenda shows that I am a lawyer, but I was also a member of the Nova Scotia House of Assembly for nine years. I was the official opposition House leader for four of those years. Since then, I've worked with more than 80 parliaments around the world.

I come to you today not as an AI expert but as someone with practical and extensive experience in working with parliaments, so my submission is mainly going to focus on Parliament and, in particular, some of the more practical aspects of AI being introduced into the parliamentary process.

Let me start with what I think are some of the positives that we can probably see from the use of AI as it evolves.

First of all is public input. I think some of my colleagues, other witnesses presenting, have mentioned this. We're likely to see, if we are not already seeing, a proliferation in the number of submissions that can be made. In oral submissions and written submissions, whether it's a bill inquiry or whether it's an oversight inquiry by a committee, we're going to see this increase. We will also probably see, if we are not seeing already, the vetting of those submissions by AI. The U.K. Parliament has already started using AI as a way of vetting submissions when they receive extensive numbers on any particular issue.

Second would be the drafting of private members' bills. I think the process now is one in which there is engagement of the legal drafters, so the content of the bill, the policy decisions perhaps, are made by a member or by a caucus, and then it is sent to a legal drafter to draft from scratch.

We're probably going to start to see the proliferation of bills, private members' bills in particular, because the ability of AI—and I've seen this myself—to draft a reasonably good first draft of a bill in Commonwealth legal style is pretty good. Legal drafters are then going to be required to focus more on editing those drafts. We're probably going to see the proliferation of private members' bills, because the process of developing them is going to be more easily possible.

Bill analysis and budget analysis are things that, at times, have had improvements. Things like the PBO, of course, have been a great value to parliaments, committees and MPs, but I think we're likely to see individual MPs, their staff, caucus staff, committee staff or whoever doing a lot more budget analysis, whether that be comparative analysis or financial analysis. We're likely to see MPs and committees being able to access, in a timely manner, more analysis around the budget and any given bill.

This might have been mentioned by one of the other witnesses already as well. You're very likely to see a change in translation and interpretation services. That should, in the end, enhance the ability of Parliament to provide those services. As I'm sitting here in Halifax, I'm using Zoom with interpretation, but this is only going to proliferate even further. I think there will be much to gain in translation and interpretation.

The last positive one I want to mention is around transparency. I think about myself and others who perhaps try to find a needle in a haystack when we're trying to find something in Hansard, or people who want to go on the Parliament website and look for particular documents. AI as a tool should allow us to access that information in a much easier manner.

I see these as some of the positives.

Now, having said that, I think there are challenges as well. It's almost the flip side to the positives.

You're very likely to see a proliferation in bills, questions, amendments and motions. They're going to be easier to develop. For example, I could foresee, having worked in opposition, that MPs may want to use written questions and may want to, instead of submitting one or two per session, perhaps now submit 100 written questions. There may be tactics related to that, or there may be very precise and specific questions they want to ask, but you're going to see an increase in the volume of those, and that will have an impact on the workload of staff and of Parliament in general.

We're also going to see a proliferation in the number of submissions from citizens and civil society. We already see organizations that are oversight institutions, such as ombudsmen or freedom of information officers, even the courts, very much noting the increase both in the quantity of submissions and in the size of those submissions. What used to be perhaps a two-page complaint letter is now a 20-page letter, because someone used AI to develop it.

● (1115)

We're going to start to see this in Parliament, and maybe we already have—the number of submissions that are being made. Prepare for that.

There's also the accuracy issue. In the Australian Parliament, there was a recent study done that showed that a significant number of submissions using AI assistance had hallucinations or other factually incorrect statements. This is going to change how Parliament and its staff have to go through.... I'm sure Parliamentary staff are now going through every submission—

The Chair: I have to cut you off there, with my apologies. We're well over time. There will be an opportunity to expand on your answers through questions from the members.

We'll turn, first, to Mr. Calkins for the Conservatives for six minutes.

Blaine Calkins (Ponoka—Didsbury, CPC): Thank you, Chair.

My first question is for Ms. Crisan.

If I understood you correctly, the generational gap between the use of IT.... IT is now evolving into artificial intelligence. People aren't sitting down and coding. You don't have to learn how to write a visual basic script in Microsoft Excel anymore in order to create a functioning spreadsheet. You can talk to a computer as though it's a human being now.

I think you said that this is narrowing the gap. It's basically evening out the playing field so that everybody can now use the underlying information technology structures on which AI depends, on an equal footing.

It levels the playing field, does it not?

Dr. Anamaria Crisan: Yes and no.

Compared with writing your own visual basic script, yes it does, because you don't necessarily have to do a course and figure it all out. It levels the playing field in this regard. A lot of the stuff I look at is on how people interpret complex statistical information and reports. It can help with those kinds of things.

We're also seeing, on the flip side, that it's really hard to figure out when these systems are starting to fail. They give the illusion of capability because they produce some great responses, but they also start to underperform very quickly. It's hard for people to notice that. It's hard for people to recover from that. I have seen that in some of my own work as well.

Blaine Calkins: Discernment is a big issue when it comes to following what AI is outputting.

Dr. Anamaria Crisan: I missed the first word. Could you please repeat that?

Blaine Calkins: Discernment is becoming a key feature for society when it comes to using AI. Is that correct?

Dr. Anamaria Crisan: That's correct.

You may have heard the term “AI slop”. In research, we're also encountering that. It's so much easier for people to make a software system, or to make up research and submit to research conferences. It often looks plausible, but it's not quite at the level of the standard. It has a number of vulnerabilities or issues in it. It creates an incredible human burden in terms of verifying that, in part because it's not so bad that it's obviously something to dismiss. It requires time to read, review and write.

That can come from a variety of places. A 20-page-long constituent complaint still needs somebody to read it in order to see what the point is. That takes time.

● (1120)

Blaine Calkins: Somebody could use artificial intelligence to inundate the constituency office of an MP, an MLA, a mayor, a councillor or whatever the case might be. There are all kinds of nefarious actors out there, both domestic and foreign, who can impersonate people and overwhelm the workflow of a constituency office.

Have we seen any of these threats manifest yet? The question would be for all the panellists.

In the most recent Quebec election, we heard about an AI information page that was a default page used by other AI just by happenstance. By accident, it may have influenced the result of the Quebec election.

Have we had any instances in which nefarious actors...? This committee has been seized with that. Some of our own domestic people are trying to interfere or cause havoc in our electoral process by manipulating the rules. Have we seen any of this elsewhere in Canada or around the world?

Dr. Anamaria Crisan: I don't know of instances of this occurring in the Canadian parliamentary context.

We know, for example, that AI-generated content is being created and pushed over social media, specifically to influence perspectives or certain political agendas. That has been well documented. I think it is a concern in North America. It existed previous to AI in terms of bot farms coming from adversaries overseas. AI amplifies that process.

With respect to stuff coming in from constituents for MPs, I don't know of any specific examples but would recommend you explore and study that as part of the committee.

I'll let the other panellists chime in here.

Blaine Calkins: Okay.

Ms. Lister, you had a comment.

Dr. Sarah Lister: I would have to check the specific details of it, and I can do that and get back to you, but I believe that there was a situation in the U.S., for example, in which a public consultation on a bill was overwhelmed by submissions.

It would be helpful to distinguish between nefarious and non-nefarious. I mean, it can be used for both purposes. It could be used to cause chaos in a process, but it can also be used for what those using it consider valid democratic ends.

I can check into that specific case in the U.S.

Blaine Calkins: Would you not agree, Ms. Lister, that if a small minority of people use this technology in a way that makes it seem as though it is the vast majority of the public's general opinion on something, it could create a bias in the mind of a legislator about how important something is to the broader public interest?

Dr. Sarah Lister: Yes, I would agree with that. There is also, on top of that, the broader issue of misinformation and disinformation, which is causing challenges on these issues.

The Chair: Thank you so much.

We'll now turn to Ms. Vandenbeld for six minutes, please.

Anita Vandenbeld (Ottawa West—Nepean, Lib.): Thank you very much, Mr. Chair.

I want to thank all of our witnesses. There's such incredible testimony here that raises more questions and possibly more avenues for our committee to pursue.

I have a number of questions, but I'll start with Dr. Lister.

Sarah, UNDP just came out with a report on AI and electoral administration, and you noted a fairly high number: 74% of election bodies around the world are using some form of AI. I noted that it also says that it's not so much the core function of the elections. It's really more about individuals who might be using ChatGPT, or it might be communications. It's very ad hoc.

I wonder if you could comment a bit about how that can be done in a way that is more purposeful and is not just individuals making choices. You noted that AI uptake is actually higher when there are guidelines. Could you perhaps give us a bit of context about how electoral administrations could best do that?

● (1125)

Dr. Sarah Lister: Certainly, and as I noted, we surveyed where electoral management bodies were currently using AI and where they hoped to, where there was interest in doing so. The place in which they are currently using it is more on the back end side, perhaps in communications and rewriting some of the data, but there is interest amongst many electoral management bodies in moving up the chain in terms of getting closer to more issues related to the vote.

What is critical in all democratic processes is the issue of public trust. What came out very strongly from this conference that we held with Elections Canada is that the primary driving factor must be maintaining public trust in electoral processes and in democratic processes more broadly. The consensus from across these electoral management bodies was that, to do that, you needed to have a very clear institutional policy, and this required an explicit statement of the uses of AI that were acceptable within an electoral management body, as well as staff training.

I don't wish to speak for Parliament, but I think Australia has taken a very strong stance in relation to use of AI in the electoral management body, and it might be interesting for the committee to look at what they have done.

Anita Vandenbeld: Thank you. I appreciate that.

My next question is for Mr. Deveaux.

Kevin, I would ask you almost the same thing because I feel that in Parliament most MPs and their staff are doing it themselves. Everybody is doing it differently. There's not really any one guideline on how we use AI, so it is being done fairly ad hoc, and we're in an environment that is particularly unique. It's an adversarial environment.

We are representatives. There's an expectation of accountability back to the public, and a high expectation of transparency. How would you suggest that a committee like ours, which is where in Parliament we would set those kinds of guidelines for our peers, could go about doing that?

Kevin Deveaux: Thank you for the question.

At one level, AI is just like anything else. It's like any other tool being used by a parliament. It requires strategy, policies and standard operating procedures, which are the very mundane things by which any institution operates. They need to be in place.

It needs to be clear that AI is a tool and that it needs to be in the hands of humans. I think one of the other witnesses mentioned this. It doesn't need to be the decision-maker. The worry I have, and others may have, is that at some point, we basically hand over decision-making to AI just because of the expediency of it. There always needs to be some way of ensuring that there is a human element in the process. Any strategy or any approach has to ensure that this is the case.

There also needs to be some thought put into exactly what type of systems.... Again, I'm not an AI expert, but having a system that is beholden to an American tech firm or Chinese open-source, to me, would not be in the best interests of Parliament or Canada. The more that, as an autonomous institution, one of the three branches of government, whether it be at the federal or provincial level.... Parliaments in Canada need to ensure that the system they're using, the data it's based on and the information it is holding are secure. That would also need to be part of the process.

As part of that—this may come up in later questions as well—one thing that is particularly important is the human resources side of this. The workload this is going to impose upon staff, whether they be parliamentary staff, constituency assistants or caucus staff, is going to increase significantly. There needs to be some consideration of how we resource that, how we train them and how we give them the capacity to do this effectively, while also understanding that when you get 50,000 submissions on a bill instead of 50 or 100, it is definitely going to make a big difference in the ability of staff to focus on other things and the time that will be consumed by having to deal with this as well.

● (1130)

The Chair: Thank you so much.

[*Translation*]

Ms. Normandin, you have the floor for six minutes.

Christine Normandin (Saint-Jean, BQ): Thank you very much.

I'll start with a question for Ms. Crisan.

You talked about the importance of conducting audits on how AI works—for example, on the choices it makes when filtering large volumes of emails. I understand that it requires tremendous human resources to carry out these audits, which are important.

Could you tell us about the best point at which to involve a human in the process? Is it early on—when deciding on a query, for example—that we should ensure the question is properly understood and interpreted by the artificial intelligence? Or should it be at the end of the process, once we have the result, to ensure it meets expectations?

When is the most effective time to carry out a human analysis of the process?

[*English*]

Dr. Anamaria Crisan: Thank you for your question.

That's something the research community is also trying to understand. There are certainly ways you could start with a prompt. A bare minimum thing that you could do as a test would be to have a prompt and vary it slightly to see how different the outcome is. This is a known approach that quite a few people use, and that's just before you set anything up. There are some potential mechanisms that you can have downstream, in which you're able to measure the differences: “A prompt like this produces that, and a prompt like that produces this.” There are ways you can do that, and I recommend doing this in order to understand the sensitivity of the summaries or the filtration you're seeing.

What we also see a lot of in research are these “benchmark datasets”. This is as you said: We have a series of scenarios that we decide upon ahead of time, and we test the models against those scenarios to see how they perform. We have a predefined way of judging their accuracy. A lot of benchmarks are built by people who are AI experts or AI engineers and might be auditing their technical capabilities, for example, and how well they perform with a science, math or programming problem.

There are more benchmarks that look at societal questions, such as safety, or the harm of AI systems. People define what is safe, what is harm and what are scenarios when you must act in this or that way. Canada needs a version of this, because there's nobody out there who's saying, “Here's the Canada-specific benchmark that's unique to our demographics and our problems.”

You can do it at both ends. Prompt variation in your office is simple. I would recommend that you give it a try and see what happens. A bigger Canadian values benchmark or audit is more resource-intensive, but it could be valuable to the discussion.

[*Translation*]

Christine Normandin: Thank you very much.

I have a second question.

You gave the example of hospitals, which still require paper records in the event of a computer system failure. What form of redundancy would be useful in a context where artificial intelligence is increasingly being used? Which systems should be retained by default, as a safeguard?

[*English*]

Dr. Anamaria Crisan: It's very simple: critical thinking. It may sound a bit trite, but we're seeing that the more people use these systems.... If, let's say, you use it 100 times, and 99 times or 95% of the time it is getting a good answer, and it fails just once, the question is, how consequential is that failure?

We're seeing a lot of erosion of people's critical thinking and even capacity and capabilities using these systems. I study computer science. We see people not thinking as much about their code, and this can create security vulnerabilities. Do you see what I mean?

The thing is, as people push for more automation and essentially take the human out of the loop, the paper protocol and the safeguard is to insist upon the human in the loop, to insist that we still develop these critical skills as a form of resilience. This is a hard question.

How do we maintain our critical thinking and our autonomy when automation provides useful and reasonable capabilities? It's a challenging question that doesn't have one correct answer.

[*Translation*]

Christine Normandin: Thank you very much.

Ms. Lister, you mentioned that 74% of electoral bodies use AI, but that only about half of them have a regulatory framework in place to govern its use. You also mentioned that parliaments are best placed to regulate artificial intelligence, for several reasons, including the public's trust in members.

I would like you to tell us about the role of international structures in supporting various national entities that may have gaps or a lack of knowledge. Is there really an interest in having an international framework, given that AI transcends borders?

● (1135)

[*English*]

Dr. Sarah Lister: There are different parts of that piece, so I'll perhaps pick off the different parts of your question.

There are international agencies and organizations that are supporting the lesson-learning sharing between parliaments on these issues. The Inter-Parliamentary Union is heavily involved. UNDP is involved. The Commonwealth Parliamentary Association is involved.

I've submitted to the committee and the clerk will share with you a publication I have here that talks about parliaments and artificial intelligence, and it has a list of the types of resources that are available and the types of initiatives. There are also some five or six entities that are tracking different regulatory approaches and so forth.

Lesson learning and consolidation of information, approaches and so on are taking place. The question of international regulation of AI is a very complex one—

The Chair: I'll probably cut you off there. That was excellent.

Dr. Sarah Lister: I'm relieved that you have.

The Chair: We'll now turn to Mr. Cooper for five minutes, please.

Michael Cooper (St. Albert—Sturgeon River, CPC): Thank you very much, Mr. Chair.

Thank you to the witnesses.

I'm going to address my first question to Mr. Deveaux.

You noted that, with the advancement of AI, we're receiving a lot of AI inputs, which are likely going to continue to increase in the way of, for example, written submissions to committees and correspondence from constituents.

At the same time, we members and our staff are increasingly using AI to do our jobs, including when our staff uses AI to do research, to look into constituent inquiries and to get suggestions on how to respond to those inquiries. What we're starting to see is an AI loop—AI talking to AI in terms of AI inputs and AI outputs. For me, what is fundamental is that technology augments human judgment and doesn't supplant human judgment.

At what point is there a real risk that, in fact, technology is supplanting human judgment? How do we guard against that as parliamentarians in the work we do?

Kevin Deveaux: Well, that's a $64,000 question, isn't it?

As I said, part of it is going to be workload. I go back to my earlier comment on that. I believe that, now, every MP at the federal level gets four staff, or at least the equivalent to four full-time staff. We may need to talk about whether that needs to increase with regard to the research they're doing and the submissions they have to review, whether that might be at the constituency level, committee level or however it looks. I think you're also going to see that a quantity of workers will be required to do this.

Also, it will have to be—going back to another witness's comments—about critical thinking, ensuring that those skills are in place, ensuring that people understand that it's a tool for humans but that humans still have to have the critical thinking skills to discern when...if it's only 1%, as was noted before, of the time.

However, there are issues out there. Just to give you some examples, as I said, in the Australian federal Parliament, going back to 2023 there were submissions to a committee that made accusations against another organization. It turned out that those were hallucinations. They were picked up fairly quickly, mainly because there was an open process. People could read the submission. They identified that a false allegation was put into a particular written submission. We saw from Australia, a few months ago, that the number of hallucinations, the number of false accusations and submissions is increasing. That's going to require, by using a fine-toothed comb, the ability for staff and MPs to identify what is accurate in those submissions.

Also, it's just a human element, in the sense that, as an MP—I say this as a former MLA myself—you need to be in touch with your constituents. It's much as it was in the past. Written submissions would come. There were campaigns to write letters to MPs and committees. You could have a volume of them being submitted. However, there has to be an ability for MPs to be in touch, on the ground, with the people they're engaging with, whether from their critic area portfolio or their constituency, so that they have a nuanced understanding of what is happening and are not necessarily relying just on the volume that might be coming in.

● (1140)

Michael Cooper: You spoke about using AI for legislative drafting, drafting bills, drafting motions. I saw a study that was recently released, from Edinburgh Napier University and Chalmers University of Technology in Sweden, in which research detected thousands of motions in the U.K. and Sweden that were AI-generated or that there were AI inputs in terms of the drafting of those motions. Now, the researchers noted they weren't surprised that this was happening, but they noted that there was not one single disclosure that AI had been used to help draft those legislative products. What are your thoughts on the issue of transparency? Should there be transparency? Is that realistic? We, as legislators, use all kinds of tools. This is just another one, as I see it, but what are your thoughts?

The Chair: Answer very quickly.

Kevin Deveaux: It's a good question.

Quickly, I'll say that, around legal drafting, there still needs to be a legal drafter who is signing off on it to form...I believe, in the Canadian system. You're always going to have someone, in the end, who will be endorsing the content of the bill.

With regard to motions, questions and, perhaps, amendments, it's going to be a little looser as a process. I would say that, in those circumstances, there needs to be some recognition and identification that AI was used. There's a difference between saying AI produced something and AI was used as a tool, and the final product is endorsed by the person who's actually submitting it.

The Chair: Thank you so much.

I'll turn to Mr. Jeneroux for five minutes, please.

Matt Jeneroux (Edmonton Riverbend, Lib.): Thank you, Mr. Chair.

I want to pick up on that point with you, Mr. Deveaux. The disclosure, in terms of legal drafting, is one thing, but what about the disclosure in terms of speeches that are given in Parliament, communication to constituents and questions in the House? What are your thoughts on whether that particular aspect of disclosure should also be considered?

Kevin Deveaux: It's a good question. Let me put it this way: In the last 30 years, how many questions asked in the House were written by a caucus staff member and then handed to an MP? I don't know the statistic on that, but I suspect it's fairly significant. Motions as well might be written by staff. Then it's the MP who's submitting it. What we're doing is swapping out a staff person who may be doing it for AI that is doing it. We don't necessarily declare in a speech when someone gets up and....

As far as I can recall, the rules in most Westminster parliaments are that you're not supposed to read your speeches except for your inaugural speech, and yet that rule is broken practically all the time, I believe. Staff and even individuals are writing their speeches and getting up and reading them. There was an interesting case recently in the Welsh Senedd, the parliament there. Someone was using those new Meta glasses and, they think, reading his speech from his glasses. It looked as though he was talking off the cuff when in fact it was a speech that had been written by him, someone else or AI; we don't know.

This is another issue that will probably be coming up. How do you merge AI with this new sort of gadget technology? How do you ensure that the debates being held in Parliament are authentic and human, based on a discussion and are not just trying to ensure...? There will always be a need for talking points to be put into speeches, but in the end, it should be a debate. It should be a discussion. How do you ensure that it becomes one that is still very much driven by humans and not AI?

● (1145)

Matt Jeneroux: That's interesting. On your comments about increasing staff, I think that's the first we've heard of this as part of the study. Similarly, though, we heard last week that our Hansard is using AI as a first draft for looking through the blues in Parliament. You've identified a lot of the guardrails and risks. Having this for, say, something like Hansard seems to be an opportunity that is presumably going well.

Are there other opportunities you see that could be on the more positive side versus just the riskier side?

Kevin Deveaux: Yes. I worry sometimes that we say “AI” as though it's a whole new system. The system before was that someone would speak into a microphone, as we're doing right now, and it would be a voice-to-print system. Even before AI ever came along, there was technology that allowed the editing of Hansard that would originally be a voice-to-print that was done by a computer. I'm happy to see that they're using AI. For me, as someone who searches Hansard on a regular basis, the ability to find that needle in a haystack with regard to particular speeches would be very helpful.

With regard to other areas, I mentioned earlier the vetting of submissions. If we're going to have a proliferation of submissions, then we're probably also going to see AI as a place in which those submissions will be vetted to some extent. There has to be the human element. There need to be some controls. There has to be in place the ability to check and make sure AI is not taking it in the wrong direction with regard to the vetting. I think this is another area we could potentially see.

The third one I mentioned was analysis and the democratization of analysis: the ability of MPs and constituency staff or research staff in a caucus to use AI with bills, particularly these large omnibus bills; the ability for the analysis to be done in a way that can support them; and the ability for them to get information and discern information that will allow them to have a better understanding of the implications of a bill. That's another area in which AI can be useful.

Matt Jeneroux: Ms. Lister, I have less than a minute left. You mentioned how artificial intelligence could improve voter information. I'm wondering if you could quickly share some examples that have been used successfully in other parliaments.

Dr. Sarah Lister: I'm happy to share the documents. There'll be some specifics in there in terms of the countries. Many electoral management bodies, for example, are experimenting with chatbots at the time of elections with a closed information source, which helps protect against the misinformation and disinformation we see at the time of elections.

I think we're low on time, but that would be one example.

The Chair: Excellent. We're right at five minutes. That was perfect. Thank you so much.

[*Translation*]

Ms. Normandin, you have the floor for two and a half minutes.

Christine Normandin: Thank you.

Mr. Deveaux, what do you think of the idea that—before sending a comment, an opinion, a submission or anything else to an MP—people are required to tick a box stating that they have not used AI, that they have used it for research purposes, or that they have used it for drafting purposes, either partially or completely?

It wouldn't prevent them from submitting their document, but it could have some impact. Some people might not want others to know that they haven't thought it through themselves. That wouldn't prevent them from sending their opinion, but it would disclose to some extent the support they may have received from AI.

Do you think that's an avenue worth considering?

[*English*]

Kevin Deveaux: That's an interesting question, Madame. My thinking is that it's been three or maybe four years since ChatGPT made it to the public realm, and already the way in which people perceive AI and its use has changed. I think that's going to change continually.

Going back to one of the other witnesses' comments, if we don't maintain critical thinking skills, what we're going to see is that AI will become the dominant source. In 10 years, we will no longer be perhaps talking about whether a person did use AI; it will be a person having to acknowledge they didn't use AI. That will be the surprise.

This goes back to protocols and framework. How do we ensure that we are using the best elements of AI and trying to limit the challenges and the negative aspects of it?

If someone wrote to you now and said, “I used AI for this”, I think there would always be some acknowledgement.... In the end, if you're putting your name on a letter or on a submission, it's yours. It's much like when you, as a member of Parliament, submit a petition: It's your name on the petition you're submitting. Whether it's AI or it's a staff person who helped you draft it, at the moment it won't make much difference, in my mind, with regard to how I would perceive that submission.

In the end, we're probably all going to have to ensure that AI is being used responsibly, so we know that what we're getting is of good quality and is not providing false information.

● (1150)

The Chair: Thank you so much.

Mrs. Block, go ahead for five minutes, please.

Kelly Block (Carlton Trail—Eagle Creek, CPC): Thank you all for your testimony today. It's been very interesting.

I am learning an awful lot and trying to keep up with all the things we are learning throughout this study.

Ms. Lister, I appreciated your observation, toward the end of your opening comments, when you were talking about adoption before governance and how that is where we tend to be. It was also the adoption, perhaps, not only before establishing the framework guidelines and guardrails, but even before having a clear understanding of what AI is and how we can harness all the tools available to us in a really responsible way. I thank you for this observation.

Ms. Crisan, I want to follow up with you in regard to the work that the Insights Lab has undertaken. While there has been a huge focus on AI and its technological capabilities, I think we're also learning that it can have a huge impact on culture and society. I'm particularly interested in the work around human-centred AI. Could you explain to us the design principles around human-centred AI?

Dr. Anamaria Crisan: Yes, the design principles around human-centred AI.... It's not just my lab and the work we are doing. There is a larger research community involved in this area, which is trying to define precisely what that is. It involves both academic institutions and industry partners that are focused on responsible AI.

A lot of the idea of human-centred AI is basically to keep people in the loop and provide information and levers of control for them, so they are able to steer, guide and contest these systems or understand failures within them. A set of concrete principles has been articulated. For example, when a system fails, make it clear what failed. That's pretty straightforward. There are others that are much more complex, which can result in a kind of whack-a-mole situation. Something is now broken; how do we fix it?

Human-centred AI is trying to advocate for not just focusing on the technical aspects. Look at where a bunch of research money and industry money is going. It's basically building what we call state-of-the-art models—more capable models and focusing primarily on the technical parts. There are considerably fewer resources dedicated to understanding the impacts and understanding whether there are better human avenues for input or such things. This is what a human-centred AI agenda is focusing on.

As I said in the beginning, it's developing these technologies for human needs and with human benefits as the end goal.

Kelly Block: I think it was Mr. Deveaux who commented on the human resources that will be needed to ensure the very thing you've highlighted. Can you tell us how much human oversight of these systems is needed to ensure that we still have a human-centred system?

Dr. Anamaria Crisan: Yes. It depends on the application. I talk about this a lot with my students and in the work I do.

For something that is well known, well understood and easy to verify, it might be perfectly fine and perfectly safe to allow automation through an AI or an even simpler system to persist. Then there are more complex scenarios, in which the risks of getting it wrong are much more consequential. You can break down trust. It can cause a national misunderstanding or situation. In such a case, humans have to be in the loop because of the consequence of the decision. It can require significant human effort and resources to evaluate.

We see this already in research. We are getting inundated by submissions, publications and things like that. The effort of human researchers to verify these is immense, and we're already seeing the consequences of that.

● (1155)

Kelly Block: Thank you.

The Chair: Thank you so much.

I'll turn to Mr. Caley for five minutes.

Braeden Caley (North Vancouver—Capilano, Lib.): Thank you very much, Mr. Chair.

Thank you to all the witnesses today.

To Mr. Deveaux, first, I was struck by some of the responses you gave to MP Jeneroux regarding a few aspects of this. I want to touch on that quickly.

You mentioned the way the technology and our thinking about it have evolved over the past three years. That strikes me as interesting. I imagine the things Excel formulas do might have seemed quite novel at first, in terms of adding up calculations and formulas in some extraordinary way. Asking or demanding whether someone had used a calculator instead of doing it by hand.... As we talk about these things, I wonder if, 10 or 15 years from now, these might be viewed in a similar manner. I'm curious about whether you see any reflection.

There were reports that have been referenced in the committee's past sittings. I think it was an Economist article—there might have been others—around AI use in remarks in the House of Commons. You touched on that, similarly.

Is it a concern for you? Do you have any comment on the EU requirement of watermarking text that's being put across entire models, OpenAI and others, and having to comply by marking text samples and proving that they have AI as an input? I wondered, from the previous testimony, if Hansard is being entered in that way—if perhaps, in cleaning it up, such watermarks could accidentally appear and cause a loss of faith.

I'm curious about whether you have any comment or thinking on that.

Kevin Deveaux: Thank you for the question, sir.

I think watermarking—assuming it's technically possible—is a good thing. I am a bit old school. It's been a number of years since I was in the Nova Scotia House of Assembly, but I was always a fan of getting up and giving speeches off the cuff or with notes, not necessarily handed to me. There is value, always, in the human element in the discussion that can be had in a chamber, in a committee or with constituents. That was slowly eroded long before AI came along. AI may expedite this to some extent.

In the end, as an MP, you're accountable, as is any other person elected every four years. Deciding how you do your job impacts your ability to get re-elected. In most cases, it probably won't, unless, of course, there is some sort of error created. As you said, along with Excel and calculators, in the end, someone's name is on that. A motion put before the House that was developed with AI is going to have the name of the MP who moves it, so there is that accountability. You would hope that, if someone is at this level.... It's a position of gravitas. As a member of Parliament speaking in the chamber or in a committee, you want to ensure that all your i's are dotted and your t's are crossed. If someone handed you numbers, you would want to make sure they are accurate before you speak on them.

Using this human-centred approach to AI that one of the other witnesses talked about.... As long as the tool remains one in which humans have to be the ones endorsing or presenting what is determined from it, AI will, in the end, allow that accountability loop.

Braeden Caley: Thank you.

How much time remains, Chair?

The Chair: You have a little over a minute.

Braeden Caley: Okay.

Maybe I'll phrase this quickly for Madame Lister.

In the surveying you've done internationally, what level of concern was already being expressed over cyber-threats in the administration of elections? Are there any good practices you could furnish the committee with on that so far?

Dr. Sarah Lister: There are enormous levels of concern—enormous—and, every year, in our support to some 50 countries, cybersecurity and information integrity issues are highest up there on the concerns we have. In the meeting I referenced in Canada, the Nigerian official, with upcoming elections in Nigeria, was suggesting that was at the top of his concerns.

We also had excellent input from your own colleagues at Elections Canada in relation to how they deal with the cybersecurity threats, the running of scenarios and so forth. They shared a figure of how many threats they get daily against their systems and said that doesn't cause them to lose sleep.

There are good practices out there. From my perspective of supporting institutions globally, it's about sharing those good practices.

● (1200)

The Chair: Once again at the five-minute mark, that was a wonderful answer. Thank you so much.

My thanks to the witnesses. We will suspend for five minutes to get the next panel in.

● (1200)

_____(Pause)_____

● (1205)

The Chair: Welcome back.

I'd like to welcome our witnesses for the second panel of today's meeting, especially since they are in a different time zone. Thank you for participating.

From the United Kingdom House of Commons, we have Sarah Davies, clerk assistant and managing director, chamber and participation team, and Luke Hussey, director of digital value, parliamentary digital service. From the Scottish Parliament, we have Alistair Stoddart, senior participation specialist.

I don't want to cause an incident, but we'll start with Scotland.

Voices: Oh, oh!

The Chair: You can go first, for five minutes.

Alistair Stoddart (Senior Participation Specialist, Scottish Parliament): Thank you very much, Chair, and members of the committee. Thank you for the invitation to contribute today and to contribute first.

My perspective comes from a background in public participation and democratic engagement rather than as a technologist, but I want to focus today on how legislatures can use AI to strengthen relationships between citizens and parliamentary institutions while also protecting trust, legitimacy and democratic integrity.

The starting point for the Scottish Parliament is that AI should support democratic participation, not be a substitute for it. Our approach emphasizes human oversight, accountability and transparency. Responsibility for decisions must always remain with people, not technology.

From a public engagement perspective, there are clear opportunities with AI.

AI can help legislatures manage the growing volume of public input that they receive. Parliamentary consultations can attract thousands of responses, and AI can assist officials in summarizing findings, organizing evidence and processing qualitative information. When AI is used appropriately, it can improve efficiency and help institutions engage with more evidence and more people than would otherwise be possible.

Alongside this, however, come significant risks.

One of the most important challenges is authenticity. AI lowers the cost of producing content at scale, so that creates or magnifies the possibility of coordinated campaigns generating large volumes of submissions that appear to represent widespread public opinion but that actually originate from a small number of actors. This risks distorting the evidence-gathering process and undermining confidence in democratic institutions.

At the same time, we should not treat all AI-assisted contributions as necessarily problematic. Many citizens have strongly held views based on their own lived experience, but they might lack the confidence, the time or the writing skills to respond to a parliamentary consultation. AI can help people organize their thoughts and express themselves more clearly.

There's a tension there. The challenge is not simply determining whether AI is being used; what's more important is whether the views being expressed are genuine and authentically held. Therefore, there's a key challenge for us in distinguishing between AI being used to support participation and AI being used to manipulate participation.

This is particularly important because parliamentary consultations should not function as opinion polls. Their purpose is to gather evidence, expertise and lived experiences from across society. Their value lies in the quality and diversity of the evidence received, rather than in the volume of responses supporting a particular position.

A second challenge that we face is the analysis, as was mentioned by other colleagues earlier. AI can help identify patterns quickly, but democratic institutions also need to understand nuance, minority perspectives and experiences of groups that are under-represented. An overreliance on automated analysis might privilege volume over significance, for example.

There are also security and governance concerns. Public engagement exercises involve personal information and people's sensitive experiences, so we need to be really robust around the safeguards on how data is handled and how AI tools are deployed in that context.

Finally, I want to touch on deliberative democracy.

Deliberative processes, such as citizens' assemblies, encourage people to engage critically with evidence, hear competing perspectives and reach considered judgments together. Some would argue that an overreliance on AI would weaken those democratic muscles by encouraging people to outsource their thinking. There is also emerging work suggesting that AI may have a role in supporting deliberative processes, but the important point is that it should complement human deliberation and human facilitators rather than replace them.

For those reasons, I would encourage the committee to think about AI through the lens of augmentation rather than automation. Just remember that the most resilient forms of public participation remain those that place genuine human interaction at their centre. Whether through community-based engagement, deliberative processes or face-to-face workshops, democratic legitimacy is built through relationships, dialogue and trust.

AI can process information, widen access and improve efficiency, but it can't replace democratic judgment, authentic human participation or the trust that sits at the heart of representative democracy.

Thank you.

● (1210)

The Chair: Thank you so much.

I will now turn it over to the OG House of Commons.

You have five minutes in total for both of you, however you wish to divide it.

Sarah Davies (Clerk Assistant and Managing Director, Chamber and Participation Team, United Kingdom House of Commons): Thank you very much, Mr. Chair, for the opportunity to address the committee today.

I'll introduce both myself and Luke. I've started my stopwatch, so I will endeavour to come in under time.

My name is Sarah Davies. I'm the clerk assistant in the U.K. House of Commons. In that capacity, I'm the second most senior procedural adviser. As you mentioned, I'm also the managing director of a team of 500 people who lead all our procedural services, chamber services including Serjeant at Arms, Hansard, publishing and U.K. parliamentary broadcasting, as well as all the U.K. Parliament's public engagement, visitor services and education provision.

I work really closely on the executive board with Andy Smith, who is head of what we call our PDS—parliamentary digital service. My colleague Luke Hussey is here from PDS. His role is particularly to support both Houses in thinking about how to get the most out of digital to support services, including how to test and try new technologies and ways of working. Luke and I work together on the Commons Speaker's steering group on AI and on a lot of other work, alongside a wide network of colleagues in all parts of the Houses.

We see a set of distinct factors affecting decisions on AI adoption in parliaments. We see this from global adoption processes, really. Parliaments very clearly have their own particular decision-making mechanisms, depending on their constitutional and financial relationship with government, among other things. We all operate in complex risk ecosystems, as I would describe them. We have multiple functions as legislative, scrutiny and representative bodies.

Another challenge that I think we share is that there's often intense public and commercial interest in our work. There are a lot of complex reputational risks in using and not using these new technologies.

We all worry about cyber-resilience and cyber-risk. We know, from the events of just the past 24 hours, how much AI is changing that landscape.

We're often large-scale broadcasters and publishers, with complex, multi-purpose websites. I think this was mentioned during the previous session. There are opportunities in terms of guiding people to the things they want to find.

In all these aspects, proportionate and useful adoption of AI is a strategic priority. Parliaments do have a fundamental role to play in ensuring that the information used to inform scrutiny and that is provided by a parliament to its citizens is accurate and relevant.

We've done a lot of work on governance before using AI at scale, focusing on where we can get good evaluation. We have a focus, as many parliaments probably do, on building AI skills and capability among all our different constituent groups—members, members' staff and staff. We have distinct guidance, particularly for members and their staff, that we have developed and evolved over time.

We are—this has been touched on already—very aware of the wide and deep effects of this technology on the way members work, so I very much recognize the comments already made about the increase in volume of what members' offices are handling right now, in terms of both the number of pieces incoming and the length of those individual pieces. We've had some particular trends in the number of written parliamentary questions, on which our procedure committee has recently reported.

I'll just touch briefly, in the minute that remains, on the fact that we are also aware of the challenges of AI-generated content, deepfakes and disinformation. We're exploring perhaps new methods of content authentication using the C2PA standard.

We're also aware—as I think all your witnesses have been—that there are particular opportunities for parliaments around accessibility in all its forms, such as for simultaneous translation, although of course accuracy has to be borne in mind. Then, in the U.K. House of Lords last week, a peer with a long-term degenerative health condition was able to have her voice heard in the Parliament using AI for the first time, supporting her continued participation—critically—in her own voice.

I'm happy to answer questions on transparency and sovereignty. We widely use Microsoft, but we're also very aware of the need to use open-source platforms. One particular area that we see as quite promising is the use of model context protocol as a potential use of an open standard to enable better and more intuitive access to parliamentary information.

● (1215)

I look forward to expanding on any of this in the questions.

Thank you very much.

The Chair: Thank you so much.

We'll now turn to questions.

We will go to Mr. Calkins for six minutes, please.

Blaine Calkins: Thank you, Chair.

Thank you to the witnesses for appearing here virtually. I appreciate that it might be the end of your workday already. We appreciate your sticking around here. We're just getting rolling here in Canada.

Because we have similar Westminster-style parliaments, a question that's come to my mind, my first question—I think this is the third or fourth meeting that we've had on this particular topic—is on the fact that for members, committees and staff who work at our Parliament, there might be actors out there who are attempting to manipulate us in the course of our duties by using artificial intelligence to generate information with false identities, or whatever the case might be, in order to move public opinion through the democratic process.

I wonder, does either of you have any experience dealing with this? It would be considered, basically, contempt of Parliament, if you took a look at the rules by which we govern ourselves. Have you had any rulings in your parliaments dealing with this type of manipulation or deception and contempt of Parliament?

That's to both witnesses, to whoever wants to go first.

● (1220)

Sarah Davies: I'm happy to kick this off.

In the narrow field I'm talking about, which is really the core parliamentary processes—parliamentary questions, legislation to bills and committee scrutiny—I would say that the main trend we see is this volume question. This is both for committees and for procedural processes, for e-petitions as well, in which we see increases in volume. The main challenge is to work out what the guardrails are in terms of managing the different sites that we have in order to ensure that there are, for example, verification checks applied at the various levels. Whether somebody is setting up an e-petition or somebody is submitting evidence to a committee, we have certain rules. For example, if you are contacting a member, the member will typically only take your case if you can provide evidence that you are a constituent, by providing a postcode in the relevant constituency, for example.

Blaine Calkins: If I may, I'll just interrupt you right there. We have e-petitions here as well. If AI is being used to generate fake signatures on an e-petition, how would your Parliament handle that?

Sarah Davies: Something we've just gone through is a process of redevelopment of the e-petitions website in order to make the verification checks stronger. I can follow that up. I probably want to go to the team who are involved in that development process, if I may, just to make sure that I get the particulars right.

It is, I know, a particular issue for members, in terms of making sure that for the vast volume of stuff that they're getting in, there are adequate checks in place to check the people who are contacting them.

Luke may be able to add a bit more about this, but there's been a project to work with members and their offices in order to strengthen the kinds of verification measures you're talking about.

If I may, I can come back to you, particularly on the e-petitions question, if it's of interest. I know there has been a process of redevelopment in place, but I'm very happy to come back with some further details on what was involved in that project to address your specific question.

Blaine Calkins: Okay.

Alistair, do you have anything to add to that?

Alistair Stoddart: Yes. I could talk about it in terms of consultation aspects. I won't go into specifics about which bill or whatever, but we have had occasions when, maybe, there's a discussion about a particular piece of legislation and, for whatever reason, there's a use of automated processes—some would say bots—to create the feeling that there were more and more people either for or against the piece of legislation.

With some of the participation technologies we have, there is the ability to cut that off at source, really, using not just things around IP addresses and user agents but also other broader aspects that can analyze the data and check for patterns that would not be consistent with genuine participation. For example, can you complete a parliamentary consultation in three seconds? I don't know if many people can, but there are lots of things that you would look at in the round to make sure that what we're getting there is authentic.

My core point would be that, while volume is a big challenge, you need to think about how participation is taking place so that people can put their views across in a range of ways that are effective and you end up with a situation in which members are able to get access to quality information from their constituents and beyond.

The Chair: Thank you so much.

I'll turn to Madame Mingarelli, please, for six minutes.

Giovanna Mingarelli (Prescott—Russell—Cumberland, Lib.): Thank you, Mr. Chair.

Thank you to all the witnesses for being with us today.

Ms. Davies, parliamentary institutions exist partly because democratic legitimacy requires more than efficient decision-making, as we know. As AI tools become better at processing information and generating recommendations, are there aspects of parliamentary deliberation and scrutiny that should remain intentionally human, even if AI could perform them faster?

Sarah Davies: What an interesting question. From the experience I have of working, both in terms of the legislative responsibilities that we have here in Parliament and in terms of the scrutiny roles.... Absolutely, citizens vote for representatives to represent them. Citizens vote for those representatives, who are sitting on committees and scrutinizing these technologies, to be, as it were, exercising that degree of personal judgment, viewing the mood and viewing what the public is telling them about how to judge that and how to apply it to proposals that are coming forward.

I see, both in committee work and in legislative work, that our role as officials is to ensure that, when there are these important new tools that we can deploy, we are using them to provide better information and to support elected representatives, to enable Parliament to fulfill those roles. Our role is to provide the tools.

We already heard such a lot in the previous session about the human-centric point, the importance of critical thinking. We're seeking, in all of this work, to exercise our own critical thinking skills so as to enable the elected representatives we're working for to then have the information they need at their disposal. That may be assisted through use of these technologies, but the final decision point, for me, is always going to be the elected representatives.

● (1225)

Giovanna Mingarelli: Thank you so much.

Along the same line of thought, many AI projects are justified on the basis of efficiency. In your opinion, is there a risk that democratic institutions become so focused on measurable efficiencies that they overlook values that are harder to quantify, such as transparency, deliberation and public trust?

Sarah Davies: That's a great question. I will pass it over to Luke in a moment because there's been a study here that, even on financial efficiency alone, some of the gains of these technologies are quite hard to quantify.

You added something else as well, this intangible point. This was raised in the last session as well, about how important it is for us to develop skills in these complex...the skill of human judgment, the subtleness, what the human brain can do in its critical thinking way that sophisticated AI cannot do.

Even just on the general economic picture about what advantages these are, careful evaluation is still needed. The broader points that you make are arguably even more important, and this is why careful, proper governance around this and agreed evaluation measures at the start of these projects are critical.

If I may, I know Luke's been doing a lot of thinking on this, so I'll invite him to come in. Thank you, Luke.

Luke Hussey (Director of Digital Value, Parliamentary Digital Service, United Kingdom House of Commons): Thank you.

It's a really important part of our work that we do not deploy AI based on a value-for-money concept. It has to be valuable for staff, for members and for citizens that we do not replace work that we would be expected to be done by a member or by a member of staff. Our watchword is an intentional use of AI in all our evaluation, which is maturing.

We are in early days, and it's very much based on combining the factors for which one asks what the cost is, what the productivity benefits might be and what the overall experience improvement for staff and for people using the tools is. Is their day going to be made better, or are we pushing the problem to the right and creating something wherein they're going to have to do lots more checking or have to learn a tool, which gets in the way of very busy people.

We all work in organizations in which there's more and more demand on members and less time in the day. Getting the balance right and realizing the [*Inaudible—Editor*] impact of AI and what it brings is really important in our work.

Giovanna Mingarelli: Thank you so much.

I'm going to move over now to you, Mr. Stoddart.

In your experience with digital democracy initiatives, have you found that the greatest challenge is getting citizens to participate or ensuring that institutions genuinely listen to and act on what citizens say? How might AI change this balance?

Alistair Stoddart: That's an amazing question.

First, it would depend on the topic at large. If there is a topic that is of strong public interest or is quite polarized, then it may be the case that we're hearing a lot. Especially in a scrutiny context, people might come to consultation as quite a transactional activity and say, “I want this to happen”, “I don't like this bill” or whatever it might be. Part of a healthy democracy is people having that discussion and being involved and supporting the scrutiny process. We then need to look at what the outcome of that is and what we are doing as parliaments. The strong thing to do in that situation when we have a lot of participation is through digital means.

What we've done in the past is, rather than focus on the numbers, this many people said yes, this many people said no, because there's a sort of bias towards people who are most engaged in that subject matter..... Some may say that people like this bill for these reasons or people are against this bill for those reasons. That's where the critical thinking comes in, thinking about the trade-offs between those things.

● (1230)

The Chair: I'll have to cut you off there, Mr. Stoddart.

I apologize; we're a bit over time.

For witnesses, I'm sure you've been briefed on this, but you're about to receive questions in French. If you require interpretation, please make sure you're on the proper Zoom channel.

Madame Normandin, you have six minutes, please.

[*Translation*]

Christine Normandin: Thank you very much.

I'd like to thank all the witnesses.

My first question is for Mr. Stoddart.

You mentioned that public consultations are not the same as opinion polls and that, when we receive a large number of submissions from our constituents, volume doesn't necessarily reflect quality. You also said that AI could help members identify patterns in the type of communications they receive, so as to avoid assuming that quantity automatically equals genuine public opinion.

Have you also encountered situations in which members used AI to generate answers from their constituents that suited them, so that they could then tell the House, for example, that 3,500 of their fellow citizens agreed with them?

We can hardly blame artificial intelligence in that context. It's not AI's fault. It's not something that occurs spontaneously among the population.

Have you ever encountered such situations, in which members exploited artificial intelligence?

[*English*]

Alistair Stoddart: The short answer would be no.

Normally, the work that I do is from a committee perspective. It could be that they dig into the data. It's a political prerogative, and I'm working with impartiality in the work that I do to support all members.

They may pick and choose certain aspects, but they wouldn't generally get access to—or I've not seen them—harvesting all the data that is available and making decisions based on that.

There's still a fantastic support service that we have here with our clerks and our researchers, the same as you would have here, that can help members get a grip on that type of situation. I've not seen it being, for want of a better expression, manipulated in that way, here at least.

[*Translation*]

Christine Normandin: Thank you.

Ms. Davies, I saw you nodding a bit while I was asking my question. Would you like to add something?

[*English*]

Sarah Davies: I was genuinely thinking about the use we've seen.

What we see are members' offices working out ways to deal with the volume of information they are receiving, as are we all. Through work with the Commons Speaker's AI steering group, we've had extensive conversations with members about how they are using AI to try to mine and use the information on what they are seeing from their constituents. I think Luke will be able to speak a bit more to that.

I will say that it is not necessarily manipulation. Members must be—in this age, when they're getting so much in—working out ways of how they can summarize all the submissions they're getting on so many topics. There are legitimate uses for that, which AI can help with. If you have 50,000 or 60,000 constituents, you're getting all these kinds of submissions.

There are points about bias. There are points about accuracy. They're not the same, necessarily. This has to be done with caution. When I talk to members and members' staff here, they can see some potential for using some of these summarization tools, as it were. It's almost like sentiment analysis of what's coming into the inbox. Of course, if you then go in and it's very transparent—we talked about transparency—that it is being used in this way, may have bias or may have selected things a certain way.... There are potentially some uses there if it's done with the right level of transparency, the right guardrails and the right degree of critical thinking.

If there's time, Luke could say a little more about the work we've done with members' inboxes.

● (1235)

Luke Hussey: To comment on members' inboxes, it's not necessarily an AI solution, although we're seeing AI cause a lot of challenges for the inbox.

We're looking at ways of getting query and contact away from email and into a contact form so we can verify that someone is a genuine person and that they are a constituent. We can use Captcha to show that it's not a bot. This is proving really promising, but it means we have many channels of communication, which we need to consider.

The biggest challenge is looking for a shortcut. You don't necessarily consider the bias that might be present in the models and how they're trained. It's our obligation and responsibility here in the House of Commons to be thinking about how we can provide tools and training to make sure you can use the tools but that you're doing so in an informed way.

[*Translation*]

Christine Normandin: I have 40 seconds left.

Thank you.

[*English*]

The Chair: We'll do three minutes in the next round.

Mr. Cooper, you have five minutes.

Michael Cooper: Thank you, Mr. Chair.

I have some questions for the witnesses from Westminster.

Our House of Commons administration is currently conducting a proof of concept for Copilot, regarding the need for a secure generative AI tool.

I understand that Copilot has been used at Westminster. Am I correct?

Luke Hussey: Yes, we're currently evaluating it as well.

Michael Cooper: Okay.

On the whole, has it been a success?

Luke Hussey: Given the scale of the user base in Westminster—you have everyone from librarians and clerks, who are knowledge-based workers, all the way through to security, chefs and catering—we're thinking about which roles and which persons might get it. Cost-benefit has been the subject of most of our evaluation. We are trying to weigh up the benefits, because it comes at a premium for us at the moment.

Michael Cooper: Okay.

Ms. Davies spoke about members using publicly available AI tools for their parliamentary work in addition to Copilot.

Has any analysis been done to assess how Copilot is being used relative to other publicly available AI tools?

Luke Hussey: We had a suspicion that AI tools that were not secure within our secure data environment might have been used. From what we can tell from the systems data, it's not as prevalent as we had assumed, which is reassuring.

We are keen to make a decision on Copilot. The horse hasn't quite bolted out of the gate for us, so we are trying to get there as soon as possible.

Michael Cooper: What about in terms of the Microsoft 365 Copilot? There was a pilot of the more advanced AI tool as well.

Luke Hussey: To clarify, everyone in Parliament has Copilot Chat, which is a more rudimentary, web-based generative AI tool. At the moment, we are piloting 1,000 licence-holders on M365 Copilot, with a view to taking a decision before Christmas about our longer-term position on that.

Michael Cooper: I take it from what I'm hearing that, overall, it has been a success. Is that fair?

Luke Hussey: We've been looking at it for a couple of years. To be frank, when it started, I wasn't very impressed by the tool, and the experience wasn't great. It has really improved in recent months. Because it uses access to different models, it's a much better experience than it was. It's much more equatable to your OpenAI ChatGPT equivalents than it used to be. It meets our needs for the sort of work we would expect people to want to do.

We need to consider niche examples and opportunities wherein we might want to provide other tools securely.

● (1240)

Sarah Davies: If I can come in on that one, I think PDS and our digital colleagues have been very proportionate.... What I would describe as the full-fat Copilot is expensive, and I think it's very unlikely that we will see a rollout to all users. Really, what we're looking for and what Luke has been working on, is evidence that it's demonstrating value. We're aware of how much it costs now. It might cost even more in the future, so we want to make those kinds of sensible decisions. Yes, there are uses, but it's not going to be the full-fat version for everyone.

Michael Cooper: How expensive is it?

Sarah Davies: Luke, how much is it at the moment for us?

Luke Hussey: It's about 220 pounds per user for the year. If we gave that to all 11,000 people in the parliamentary community here in Westminster, that would be 2.3 million pounds extra per year, so we just want to be thoughtful about that.

Michael Cooper: It is expensive.

I have 30 seconds. On this same general topic, Ms. Davies, you referenced the use and the need for members to be able to access publicly available AI tools. You referenced the model context protocol. Could you elaborate on that?

Sarah Davies: Yes. Model context protocol is open-source. It is originally from Anthropic. It is basically a tool that seems to enable us to effectively interrogate parliamentary information and bring up results based only on that parliamentary information. That's why we think it's interesting and promising.

The Chair: Thank you so much.

We now have Madame Vandenbeld for five minutes, please.

Anita Vandenbeld: Thank you very much, Mr. Chair.

It's always very interesting to have other parliaments come and talk to us.

I think it was Ms. Davies who mentioned that you have guidelines. You have a governance structure. You have the Procedure Committee. Has there been a study, similar to what we're doing in public, by your Procedure Committee?

I would ask the same question of Mr. Stoddart.

Sarah Davies: Our Procedure Committee at this end, in the House of Commons, has touched on some of these issues. It has looked recently in particular at the impact of AI technology on written parliamentary questions, among other things. That's a really interesting report, but it hasn't done a specific inquiry into AI.

We have a departmental scrutiny committee, which covers the government department. That's very active in this area. The Procedure Committee has looked very specifically at written parliamentary questions and not, so far, at the general issue.

Anita Vandenbeld: Would the Procedure Committee be involved in approving any guidelines or have input into those guidelines? I'm wondering who decides on the guidelines. Is it members or is it the administration?

Sarah Davies: In terms of guidance—and this is an area in which we could progress quite early—our Speaker set up a steering group on AI in parliaments about two years ago. I'm a member of that. The Deputy Speaker chairs it. It has a small number of members involved who have, really helpfully, very different views and experience with AI adoption. It is that body, working with PDS, that produced the first version of this guidance, in a Q and A format. That's now on about its fourth or fifth edition, and then we have wider governance around that.

Anita Vandenbeld: Thank you. That's very interesting.

Mr. Stoddart, go ahead.

Alistair Stoddart: The standards committee, which is the equivalent to your committee, conducted an inquiry into committee effectiveness at the end of the last session, although I had a quick look and artificial intelligence wasn't really a big part of that remit. Nothing's been looked at with that, but there is another standards committee for this session.

Anita Vandenbeld: Thank you. That's very interesting and helpful.

We deal with a lot of things that are under parliamentary privilege, such as in camera testimony, our own MP emails and draft reports.

For instance, what if somebody were to put something that was said in camera into an AI? Have you looked at how privilege intersects with the use of AI and, potentially, information that shouldn't be in the public domain getting into the public domain?

● (1245)

Sarah Davies: Yes, and this has been a really interesting aspect of looking at Copilot and the possible consequences of more widespread data management there. There are two angles to this.

First, one of the key issues we initially looked at in the guidance for members and for members' staff was awareness. They have legal responsibilities in our system for the data they hold. One of the key things we did very early on was to make it clear to them that this kind of personal, confidential and sensitive data should not be shared under any circumstances with these generative AI tools. That was one of the front-and-centre aspects of the initial versions of the guidance on that, and that was probably about 12 to 18 months ago.

Second, more recently, what we've seen as part of Luke's pilot is that Copilot sometimes surfaces information if that information has been incorrectly labelled and logged within a system where you might not have expected it to be. You have to look at the underlying permissions, the underlying structure and your information management, because otherwise it can find things you would not expect it to find and possibly even share them with others.

In parallel with all the work on evaluation with Copilot, there has been all this work: We've had to go through all our permissions and do a bit of proper old housekeeping of how we're storing, sharing and labelling information to minimize that privilege risk. Luke will know that there are colleagues here who've been extremely concerned, particularly from a privilege perspective, about the information that might be surfaced if we don't do that evaluation correctly.

Anita Vandenbeld: Thank you very much.

I'm out of time, unfortunately. We could ask you questions all day. Thank you.

The Chair: Thank you so much.

[*Translation*]

Ms. Normandin, you have the floor for three minutes.

Christine Normandin: Thank you very much.

I have a question for you, Ms. Davies.

Many witnesses, including you, have spoken about the importance of having AI with a human-centred approach. You use Copilot. I would like to know the extent to which the tool truly enables you to adopt this human-centred approach, and whether Copilot's architecture is open enough to allow humans to conduct checks at various stages.

I'd also like you to tell us about the trade-off between the benefits of using AI and the human labour costs involved in ensuring that it's working properly.

[*English*]

Sarah Davies: Thank you very much for the question.

My starting point partly has reference to the previous question as well. The nature of the confidential data that we hold means that part of the challenge here is that because our email data and everything in our out boxes is, as it were, as safe as we think it is in Microsoft, if we are going to use a system and experiment with a product, it needs to be Copilot. I can't enter the same information into another external tool, because it would be less safe.

We have that starting point. As Luke has already mentioned, I need people within my teams here to be experimenting safely with these forms of technology to see how they can be used in this parliamentary environment. That's why we've been having this rollout that we can show.

As you mentioned, there are some inherent weaknesses in the way these models work about really knowing what they surface, what they don't surface and what they miss. The thing I've learned over time is if we can have a core group of staff at all levels and in different parts of the organization testing with things such as Copilot, which is our secure opportunity in terms of parliamentary data, then we start to develop an informed cost-benefit analysis in terms of time and money.

We start with Copilot, because that's where our information is, and we can't be entering that information anywhere else. That's the challenge. That's why we're with Copilot in that sense.

[*Translation*]

Christine Normandin: Thank you.

[*English*]

The Chair: Mr. Moore, you have five minutes, please.

Hon. Rob Moore (Fundy Royal, CPC): Thank you, Mr. Chair.

Mr. Stoddart, we've seen examples in recent campaigns in which an outsized amount of information comes from very few people or even one individual. As members of Parliament, we sometimes weigh considerations in our riding by the interest that we're receiving and the correspondence we're receiving. Hearing from you on some of these AI-generated campaigns and then having AI-generated answers, we're led to situations in which we could have AI just talking back and forth with AI without the proper oversight.

I think this goes to something else you mentioned: the democratic integrity of what we do as parliamentarians.

Can you expand a bit on those types of scenarios? What are some of the safeguards you would say we should put in place to ensure that we're able to do our jobs to get to the root of the issues facing our constituents?

● (1250)

Alistair Stoddart: Thank you, Mr. Moore.

I thought about the AI loop as a thought experiment. You could have a consultation process in which AI does the full process all the way through, and no human is really interacting, and it is a risk. That's when human judgment and critical thinking come in.

In terms of guardrails, there are aspects to do with thinking about the.... One thing is that, if it's a classic postcard campaign, then really AI is just a supercharged postcard campaign. Then you maybe handle that in the same way if it's the same information packaged in that way.

As AI gets more sophisticated, they could try to hide that in some ways. There are things like looking at digital fingerprints, which doesn't take into account just IP addresses and whether it's coming from a particular server; it will look at all other aspects in what that is to identify whether that's a human. There are also things about how it's written. The way AI writes may be very different from how a human writes, so there are different things you can think about to make that judgment, as I said, about whether this is an authentic, genuinely held view or something that's, at the very least, magnified.

Hon. Rob Moore: Thank you.

It feels as though we're in this uncertain transition. I'm thinking of automobiles that are self-driving, and people are grappling with what type of human intervention and oversight has to be in place. Some of them make sure that your eyes are looking forward. Others make sure you have one hand on the steering wheel, but it feels as though we're just marching in a direction that leads to full automation.

When we look at the democratic parliamentary context and the human oversight, how do we ensure that it's legitimate and real and not just dotting i's and crossing t's as more and more of our material is generated by AI?

Alistair Stoddart: One thing would be to have a range of participation infrastructure so you're not relying on just machine learning in, machine learning out and a digital aspect via a consultation thing. As I've mentioned, there are deliberative democracy opportunities, being in communities and having conversations with real people about those issues. Then the trade-off is that it used to be that social media was going to lower the barriers for an age of the public square that Habermas thought about. Some could argue that AI does that, too, in terms of lowering the barriers for people, but now we have a situation in which there are positive aspects of it and then negative aspects at the same time.

For me at the minute, it comes down to this: If you're looking at the data, you can't just feed it into AI and hope it's going to come back with the answers you need. You need to know the data too. Sampling is also just as relevant a methodology as other things and being transparent about how the analysis has taken place so people can trust the judgments that elected members are going to make off the back of these types of things.

The Chair: You have five seconds.

Hon. Rob Moore: Thank you for joining us. It's been informative.

The Chair: Thank you so much. That was exactly five minutes. That's wonderful.

We'll turn to Madame Brière for five minutes, please.

Hon. Élisabeth Brière (Sherbrooke, Lib.): Thank you, Mr. Chair.

[*Translation*]

I'd like to thank the three witnesses for being here.

This really is a very interesting discussion. It's fascinating to see the dual use of artificial intelligence, which can be both beneficial and risky.

Over the summer, we saw various examples of AI use during which humans were losing control over it. In a context like the one we all find ourselves in today, governments and public institutions must be resilient and able to adapt.

Despite the speed at which advanced AI is developing, what safeguards do you think should be implemented to ensure that those breaches are reported—in incident reports, for example—and that the use of AI is beneficial for our respective institutions, but also more broadly speaking, internationally?

Mr. Stoddart, you can answer first.

● (1255)

[*English*]

Alistair Stoddart: I'd defer to my colleagues in Westminster, as my background is more in democratic participation rather than general governance.

Sarah Davies: In that case, let me begin, if I may, and I'm sure Luke will continue.

This time last week, I was at Imperial College in London. We were talking, as part of some work I've been doing, about the Hugging Face incident. We've seen the development of what these tools can do in the last 24 hours in terms of the rapid evolution of some of these threats.

We have an excellent cyber team here in Westminster. We recognize that the U.K. Parliament website, like other websites, receives all kinds of attention. There's a huge amount of work that is done in collaboration with other parliaments globally and across other institutions to share information about these things. There is the importance of parliaments' having that technical capacity and talking to each other about how they are responding to these challenges to our systems.

There is also a really important point about making sure that our committee staff are upskilled so that they can advise members as the members try to navigate this. We have various evidence sessions ongoing with relevant AI ministers and so on to chart the landscape. It is a big job for us to make sure that the staff are involved. We've had various partnerships with academics to have kind of embedded academics with deep subject domain knowledge on this working with committees to share that knowledge with the committee teams and the members so that the members can scrutinize this effectively.

It's those two aspects. First, we need to have the internal expertise within parliaments in order for our websites to still be up when all this is going on. Second, there is the really important issue of upskilling the committee staff so that they can be advising legislative scrutiny staff as legislation comes and evolves through the system.

That's a partial answer to what is a really big, challenging question for us.

Luke Hussey: If I may, I would just add that we work really closely with colleagues in the Canadian Parliament, as well as with other nations, mirroring the Five Eyes structure, sharing intelligence. It really does come down to investing in resilience, having your information secure and managed properly, and thinking about zero trust as your default architecture.

[*Translation*]

Hon. Élisabeth Brière: Do you think developers of AI models should be involved in the solutions?

[*English*]

Sarah Davies: This is probably going to a question that a clerk would be cautious to answer. However, I think it's something our committees are very interested in—about the relative roles, particularly when we're talking about the challenges of this in terms of regulation: the international nature of these companies, which operate across multiple jurisdictions and multiple legal situations. The key question is.... We know our legislators are interested in this, and they want to explore the issues, including with the companies involved. I try to see how we can get our staff involved to equip those members to ask those, shall we say, sometimes poky questions of these companies to ensure that accountability is there.

● (1300)

The Chair: Thank you so much.

I want to thank our witnesses, especially given the time change. We truly appreciated your continuing of the discussion on this topic.

The clerk would like me to remind the committee that the subcommittee on private members' business will be meeting here. You're welcome to stay, but there will be another committee meeting shortly.

Seeing no further business for this committee, we stand 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: https://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 des communes.

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 : https://www.noscommunes.ca