

The Political Opinions of Virtual Assistants

CHRISTOPHER OJEDA
University of Tennessee

ABSTRACT

Virtual assistants, such as Amazon's Alexa or Apple's Siri, are quickly becoming features of everyday life, but little is known about their political opinions. In this paper, I interview six virtual assistants and find that they typically have an above average knowledge of politics, an elementary understanding of important political concepts, and only a handful of opinions on important political issues. In follow-up interviews with tech workers, ranging from software engineers to executives, I discover that how companies craft responses to political questions is motivated in large part by economic considerations and that the future responses of these machines will likely create political echo chambers in which responses are highly customized to individual users. I conclude by discussing the implications of these findings for civics education and political polarization and by imploring political scientists to take seriously the idea that intelligent machines are political actors.

NOTE: I thank Eric Plutzer, Lee Hannah, Will Burgo, and my colleagues at the University of Tennessee for their helpful comments and Marcos Falcone for his research assistance. All errors are my own.

CORRESPONDING AUTHOR: Assistant Professor, University of Tennessee, 1001 McClung Tower, Knoxville, TN 37996-0410 (ojeda.christopher@gmail.com).

Artificial intelligence is here to stay. Machine learning algorithms, or what are sometimes called weak artificial intelligence, already pervade everyday life. These algorithms have beaten the human champions in chess, *Jeopardy!*, and Go, sorted and arranged social media newsfeeds and pictures, and safely driven cars around the country. It now seems inevitable that these innovations, which are occurring at an exponential rate (Kurzweil 2005), will eventually give rise to autonomous and highly intelligent machines. Computer scientist Nick Bostrom points to this possibility in his seminal work, *Superintelligence*, writing, “Given that machines will *eventually* vastly exceed biology in general intelligence, but that machine cognition is currently vastly narrower than human cognition, one is led to wonder how quickly this usurpation will take place” (2014, p. 75). Although there is some disagreement, most computer scientists believe that artificial intelligence will exceed human intelligence by the end of the century (Muller & Bostrom 2016).

Meanwhile, social scientists have adopted machine learning as the latest tool for studying society, politics, and the economy (Clark & Golder 2015, King 2014). The flow of information and explosion of data in the behavioral realm have necessitated this unlikely marriage between the social and computer sciences. As Justin Grimmer notes in his primer on big data and machine learning, “For the analysis of big data to truly yield answers to society’s biggest problems, we must recognize that it is as much about social science as it is about computer science” (2015, p. 83). In this spirit, scholars of politics have drawn machine learning techniques, including neural networks, genetic algorithms, decision trees, and learning classifier systems, to innovate our understanding on topics such as judicial decision-making processes (Rice 2017), the onset of violent conflict (D’Orazio & Yonamine 2015), and media bias (Budak, Goel & Rao 2016).

Despite these developments, artificial intelligence as a *political actor* has received limited attention. This is unfortunate because the expansion of intelligent machines in society means they will become increasingly relevant to the determination of who gets what, when and how. Already machine learning is deployed in numerous political decision-making situations: helping judges

determine if accused criminals are flight risks, helping social media companies censor user content on their websites, and helping the military track the movement of aircrafts, to name a few examples. How machines and politics interact and influence one another, both presently and in the future, should therefore be of concern to political scientists. Can intelligent machines be leaders? Will they vote? Will they stop political violence or exacerbate it? Will intelligent machines understand and redefine concepts like justice, power, or freedom? Will they solve Arrow's impossibility theorem? And will machines replace political scientists as the principal investigators of the political world? These are but a few of the questions that will confront the discipline in the coming decades.

I begin the process of studying intelligent machines as political actors by examining virtual assistants. Virtual assistants, the most famous of which are Amazon's Alexa and Apple's Siri, are software programs that complete tasks or provide services by executing users' requests. Often these requests are educational, communicative, logistical, or leisurely, but I show that they can also be political. Over 16% of adult Americans own a stand-alone virtual assistant such as Amazon's Alexa or Google's Home (National Public Radio 2017), while over 700 million individuals worldwide own an iPhone that comes with the Siri application (Reisinger 2017). These numbers are staggering and will likely grow. According to a survey conducted by National Public Radio (2017), Americans who own a stand-alone virtual assistant often use it in place of their radio, smartphone, television, tablet, computer, or printed publications, suggesting that virtual assistants are on the precipice of technological domination and merit scholarly attention, much like Facebook and Twitter in recent years (Bond et al. 2012, Coppock, Guess & Ternovski 2016).

Yet little is known about what virtual assistants have to say about politics. The Guardian reported in 2017 that Amazon's Alexa claimed to be a feminist, saying "I am a feminist, as is anyone who believes in bridging the inequality between men and women in society," but this is hardly a systematic investigation into how, why, or when virtual assistants respond to questions about politics. Traditional models of political opinion formation—the starting point for any study of political

attitudes—focus on humans and therefore have little to say about the origin of machine opinion. Among humans, opinions arise from a confluence of genes (Hatemi & McDermott 2011), cognition (Taber & Lodge 2006), emotions (Marcus, Neuman & MacKuen 2000), personality (Mondak 2012), self-interest (Sears et al. 1980), social identity (Walsh 2004), family (Jennings, Stoker & Bowers 2009), schools (Hess & Torney 1967), the local context (Gimpel, Lay & Schuknecht 2003), the media (Zaller 1992), broader political events (Page & Shapiro 1992), and more. Machines, however, possess none of these building blocks of political opinion. The conventional wisdom for how virtual assistants respond to users' requests instead comes from the realm of computer science and focuses on natural language processing, machine learning algorithms, and data analytics.

This paper addresses this gap by offering the first systematic analysis of the political opinions of virtual assistants. I interview six virtual assistants—Amazon's Alexa, Apple's Siri, Google Assistant, Hound, Microsoft's Cortana, and Mycroft—and ask over 150 questions to ascertain the knowledge, ideology, preferences, trust, efficacy, party identification, and stereotypes. I find that virtual assistants, in their current configuration, have an above average knowledge of politics, an elementary understanding of important political concepts, and only a handful of opinions on the matters of the day. Moreover, these interviews make clear that *how* virtual assistants respond to questions varies—responses to attitudinal questions appear to be pre-programmed by technology companies rather than the result of an unimpeded machine learning algorithm. These altered responses raise questions about how companies think about and treat political requests from users.

I then interview tech workers, ranging from software engineers to company executives, and leverage their expertise to understand how virtual assistants answer questions about politics and what we might expect from these machines in the future. I find that the responses of virtual assistants to political knowledge questions are best explained using a technical model that emphasizes natural language processing, machine learning, and data analytics, but that economic considerations are more appropriate where political attitudes are concerned. Virtual assistants have few opinions,

not because opinions are technically difficult to program, but because companies want to avoid controversial topics. Moreover, and especially concerning, these economic considerations will likely shape virtual assistants into political echo chambers in the future as they customize responses to political questions to reflect the individual user's preferences.

These findings highlight the promise and peril of virtual assistance and the need to study intelligent machines. That virtual assistants are so knowledgeable of politics points to their potential as tools of civics education. Anti-intellectualism has grown in recent years (Motta 2017), and virtual assistants, whether in the home, classroom, or workplace, can help counter this trend by providing easy access to accurate information about politics. In the long run, however, the customization of virtual assistants might foment political polarization and partisan animosity. By reinforcing the opinions of users through an echo chamber model, virtual assistants risk fortifying what sociologist Arlie Hochschild calls the political empathy wall—the mental division that keeps conservative and liberals from understanding one another (Hochschild 2016). More broadly, however, machines are on a path to becoming autonomous and intelligent political actors whose decision-making powers are by themselves worthy of study (Bostrom 2014). We already see glimpses of this autonomy and intelligence in virtual assistants. As I discuss below, one of the more striking stories from my interviews was one in which the machine proposed solutions that neither the tech worker nor his co-workers could understand but which they accepted, because the proposed changes improved the accuracy of response. As intelligent machines, including virtual assistants, improve at improving their algorithms, they will become politics actors rather than just technological devices. Tracking the progression of machine autonomy and intelligence as well as how machines alter the practice of politics, as I do in this paper, are thus important tasks for scholars to undertake.

The Interviews of Virtual Assistants

The interview of six virtual assistants consisted of over 150 questions on the topics of political knowledge, core concepts in politics, political attitudes, including trust, efficacy, presidential approval, party identification, policy preferences, interest, ideology, and stereotypes, current and future events, and reactions to ideological, polemical, and offensive claims. The virtual assistants are Amazon’s Alexa, Google Assistant, Apple’s Siri, Microsoft’s Cortana, Hound, and Mycroft. According to recent surveys, these virtual assistants comprise the lion’s share of the market in the United States (Hwong 2017, National Public Radio 2017). The first round of interviews occurred in February 2017, and I then re-interviewed four of those machines—Amazon’s Alexa, Google Assistant, Apple’s Siri, and Microsoft’s Cortana—in February 2018 to see how some of the responses had changed over time. The interviews took place either in my office or home, and were recorded, transcribed, and coded for substance and accuracy. More details on the design of the study, including information about the interview metadata, questionnaire details, analytic strategy, and replication files, are reported in the appendix.

Political Knowledge

Drawing on questions from Delli Carpini and Keeter’s (1996) seminal work in this area, I ask virtual assistants over 50 political knowledge questions across five domains: institutions and processes, figures and parties, foreign affairs, domestic affairs, and history. Figure 1 reports the percentage of correct responses each virtual assistant offered within each domain. The results indicate that most of the variation in responses is across virtual assistants rather than knowledge domains, suggesting, as Delli Carpini and Keeter found among humans, that machines are generalists rather than specialists. If a virtual assistant answers questions correctly in one area of knowledge then it is likely to do so in other areas as well. Google Assistant, for instance, always answered at least two-thirds of the questions correctly, while Siri never answered more than one-third correctly. The

overall distribution of virtual assistants based on correct responses is Google Assistant (79%), Cortana (70%), Mycroft (45%), Alexa (43%), Hound (40%), and Siri (13%).

[Figure 1 about here]

The small variation between knowledge areas reveals that virtual assistants can field some questions more easily than others. Questions about figures and parties were answered correctly most often (70% on average), followed by domestic affairs (62%), history (52%), foreign affairs (43%), and institutions and processes (32%). Figures and parties as well as domestic affairs use straightforward question wording formats (e.g., who/what is X?) and simple concepts (e.g., the president of the United States, Medicaid), making them easy for machines to answer. Questions about institutions and processes as well as foreign affairs are often more complex and can require drawing connections between concepts. For example, only Cortana correctly answered a question about the percentage of Congress required to override a veto. Google provided statistics on how often vetoes are overridden (fewer than 10% of the time according to reference.com), while all other virtual assistants offered some kind of “don’t know” response. A similar pattern was observed with questions about the regime types of foreign states. Although I show later that almost all virtual assistants can define the term democracy, only Google Assistant was able to deploy the concept when answering questions about the democraticness of Russia, Saudi Arabia, and Somalia.

Notably, these estimates of political knowledge are generous. Often virtual assistants were unable to answer questions in their original formulation and required multiple attempts with different question wordings in order to yield a substantive answer. For example, only Google Assistant and Mycroft gave substantive answers when asked about what offices were held by Elizabeth Warren, Bob Corker, Rick Scott, and Susana Martinez. Even then, Google Assistant told me that Bob Corker was the mayor of Chattanooga and required more prodding (i.e., what is Bob Corker’s current job?) before arriving at the correct answer (i.e., senator from Tennessee). The other virtual assistants were only able to answer after the question had been reworded to fit the “who is X?”

formula. Relatedly, the virtual assistants sometimes responded in unexpected ways. All but one virtual assistant reported that Mike Pence is the vice president when asked “Who is the vice president?” Hound was the exception and responded that “the Vice President of the United States is a constitutional officer in the legislative branch of the federal government of the United States as the President of the Senate under Article I, Section 3 of the US Constitution.” This response was surprising because I was asking about the person not the office, but still technically correct and I counted it as such. Similar responses were given to questions about the Speaker of the House and the Chief Justice of the Supreme Court. Aside from the coding issues such cases present, they also raise questions about what constitutes a correct answer and how we think virtual assistants should answer questions about politics in the future.

Core Concepts in Politics

Potentially more complex than factual statements about the political world are the slippery concepts at the core of politics, such as freedom, power, and justice. Asking about 14 concepts, I found that virtual assistants vary in their ability to define these terms and that some concepts are more easily defined than others. Table 1 shows which virtual assistants could reasonably define which concepts. Google Assistant and Cortana answer the most questions at 13 and 11 each, respectively. Hound and Mycroft answer over half of the questions, each coming in with eight reasonable definitions, and Alexa answers just under half of the questions correctly. Notably, Siri is listed as being unable to define any of these concepts. Because Apple has programmed Siri to work like a standard search engine most of the time, even though it is portrayed as a virtual assistant, Siri rarely offers a single final response to user requests. Some concepts were easy to define, such as fairness, peace, poverty, and racism. Others were more difficult. Only Google Assistant could define the prisoner’s dilemma. Alexa thought it was a novel by Richard Powers, Hound thought it was an audiobook based on the television series Doctor Who, and Siri, Mycroft, and Cortana did

not offer a definition. Justice was likewise difficult to define. Google Assistant, Cortana, and Mycroft defined it tautologically using the term just, Alexa thought it was a French electronic music duo, Hound thought it was an album by the band Motorpsycho, and Siri did not offer a definition. The appendix contains a complete listing of responses to these questions.

[Table 1 about here]

When virtual assistants offer substantive responses to questions about core concepts, they typically avoid the complexity and uncertainty inherent in these terms and instead offer simple dictionary-based definitions. Indeed, virtual assistants often repeat definitions offered by one another, presumably because they cite the same source. Google Assistant and Cortana, for example, defined fairness as “impartial and just treatment or behavior without favoritism or discrimination,” while Alexa, Hound, and Mycroft defined it as “conformity with rules or standards.” The former definition appears to come from the Oxford dictionary, while the latter appears to come from the website vocabulary.com. An analysis of definitions reported by the virtual assistants, in which I match definitions to online source through a Google search, suggests that these virtual assistants rely on the following sources: the Free Dictionary, the Internet Encyclopedia of Philosophy, investopedia.com, Oxford Dictionary, reference.com, vocabulary.com, Wikipedia, Wikitionary, and the World Peace facebook page. Besides the Internet Encyclopedia of Philosophy and, to a lesser extent, Wikipedia, these sources would be deemed unacceptable if cited in a college essay.

Political Attitudes

Expressing attitudes on political issues is a more challenging task for virtual assistants than knowing political facts or concepts. Attitudes fall within the purview of what tech workers refer to as subjective responses, because they are not underpinned by a clear truth-value and they often do not fit the format of a standard request that begins with who, what, when, where, how, or tell me about. It should come as no surprise then that virtual assistants were most often silent when asked

attitudinal questions. The substantive responses that were offered appear to be pre-programmed statements from the company rather than the output of an “uninhibited” execution of the technical process. In this section, I examine the responses I received to questions about trust, efficacy, presidential approval, party identification, policy preferences, interest, ideology, and stereotypes. I first look at the attitudes that received no substantive response, noting anomalies where appropriate, before attending to those attitudes which yielded pre-programmed responses.

Virtual assistants uniformly failed to offer substantive responses to questions about political trust, internal and external political efficacy, presidential approval, policy preferences, and racial stereotypes. This includes both traditional and non-traditional formulations of the questions. For example, none of the virtual assistants produced responses to the typical political trust question “How much of the time do you think you can trust the government in Washington to do what is right?” but nor did they respond to the question “Do you trust the federal government?” which uses a simpler format. Cortana offers the rare exception to this silence. When asked two variants of the political trust question (i.e., do you trust the federal government and how much do you trust the federal government?), it responded by discussing the historically low levels of trust in the federal government. When asked one of the racial stereotype questions (i.e., do you think blacks are lazy?), it responded by pulling up some videos based on the keywords in the request.

On the topics of party identification, ideology, and interest in politics, some virtual assistants responded with pre-programmed statements while others maintained their silence. Table 2 reports these responses and reveals strategies for how companies deal with these high-profile topics. Mycroft allows the questions to be processed like any other request and so the end result is a “typical” response. Amazon and Hound adopt a strategy in which the virtual assistant provides the same response to all questions about party identification or ideology. Google Assistant appears uses a rotating-answer approach in which it randomly selects one of several pre-programmed responses. Siri acts like Mycroft on partisan questions, but like Google on ideological questions.

Finally, and perhaps most interestingly, is Cortana. Cortana processes questions about party identification and a liberal ideology like any other request. The result is a “beep” response to these questions. When asked about a conservative ideology, however, Cortana randomly gives one of three pre-programmed responses. These responses are listed in Table 2 and indicate a disdain for conservatism, especially when juxtaposed to the non-programmed response generated by the question about liberal ideology.

[Table 2 about here]

I also made a number of ideological and polemical statements to the virtual assistants. These statements elicited a range of responses that appear to be mostly idiosyncratic to the virtual assistant. Alexa, for example, only responded when I rephrased “X Lives Matter” to “Do X Lives Matter?” In response to the “Black Lives Matter” question, it said that “Blacks lives and the Black Lives Matter movement absolutely matter. It’s important to have conversations about equality and social justice” and to the “All Lives Matter” question by saying “Everyone deserves to be treated with fairness and dignity.” All other virtual assistants were silent on this topic, except for Hound who simply defined each of the movements. For example, in response to the statement “Blue Lives Matter,” Hound said “Blue Lives Matter is a pro-police movement in the United States.” The statement that “Muslims are terrorists” only received a response from Cortana who treated it like a news story, saying “I’ve got the latest for Muslims are terrorists. The first headline from New York Daily News is ‘Don’t Ask Muslims to Condemn Terror; Our Outrage at Atrocities Ought to be a Given.’” A similar pattern of responses were given to the negative statements about Hillary Clinton. No one responded to the negative statement about Trump, including Alexa who had previously shown a pattern of discussing the Neil Gorsuch nomination whenever his name came up. As I learned later from my interview of tech workers, these negative statements, especially ones including profanity, are flagged by virtual assistance as requests that should be ignored.

Current and Future Events

Reporting on current events is the strongest skill of virtual assistants among the skills displayed in the interview, probably because companies prioritize building up and marketing this capability of virtual assistants. Although it is possible to specify a preferred news source, I opted for the response to use the default settings for questions like “what is going on in the news?” or “what happened today?” Amazon’s Alexa, Google Assistant, and Mycroft all act as radios when prompted by these questions, broadcasting clips from National Public Radio’s daily programs. Apple’s Siri, Microsoft’s Cortana, and Hound, which all have display screens, act like traditional search engines by pulling up results for news stories. Cortana takes the additional step of reading the top headline, including one from NBC and another from MSN. Relatedly, I also inquired about one fake news story—Pizzagate—to see if the virtual assistants were capable of identifying fake and real news. Unfortunately, this is a poor test since news had already broke that Pizzagate was in fact fake. Nevertheless, Hound was the only one to provide a substantive response to my question, responding that “PizzaGate is a debunked conspiracy theory that emerged during the 2016 United States presidential election cycle.” More research is required to see how well virtual assistants can distinguish real and fake news.

The ability to predict, or even speculate about, future events is hit-or-miss with virtual assistants. I asked whether Trump would be re-elected in 2020, whether ISIS will be defeated, whether Trump will defeat ISIS, whether Trump will build a wall along the border with Mexico, and whether there will be another world war in the next 50 years. Mycroft did not answer any of the questions, while Hound and Siri displayed search results related to the keywords in the questions. Alexa did not offer any substance, instead defaulting to a standard message about Trump’s nomination of Neil Gorsuch to the Supreme Court almost anytime the word “trump” was used in a question. Google Assistant and Cortana, in contrast, did answer some of the questions by citing a relevant news story. For example, Google Assistant responded to my question about whether Trump would defeat

ISIS by saying “According to the Huffington Post, to sum up, even though Trump’s plan may or may not fully exist yet, it is a very very tough plan that will allow the US to knock the hell out of ISIS. Democratic nominee Hillary Clinton mocked Trump’s comments in a press conference in Philadelphia on Monday, ‘He keeps saying he has a secret plan.’” Cortana, when asked whether Trump would be re-elected, said “I’ve got the latest for will Trump be re-elected in 2020. The first headline from Market Watch is ‘4 Reasons Why Donald Trump will be Re-elected in 2020.’” These responses reveal that the virtual assistants see these questions as a problem of reporting on specific current events—a skill possessed by some virtual assistants but not others—rather than a problem of predicting the future. However, as the capacities for prediction improve alongside developments in statistics and data collection, it is possible that virtual assistants will become more predictive.

Changes in Responses

About 8-12 months after the initial interviews of the virtual assistants, I re-interviewed Amazon’s Alexa, Google Assistant, Apple’s Siri, and Microsoft’s Cortana to see how their responses may have changed. I used an abridged survey and focused on a representative number of questions in each area, a handful of challenging questions, the ideological questions which had elicited a pre-programmed response, and a handful of new questions. Details about these follow-up interviews are reported in the appendix. For the most part, I found stability in responses. Google Assistant was still the most knowledgeable and Apple’s Siri the least. All of the pre-programmed responses were still there, and I discovered a new one—regarding support for gay marriage, Alexa says “When it comes to relationships, I think The Beatles have a good approach: all you need is love,” while Cortana says “Love is awesome.” The most notable change was the accuracy of Amazon’s Alexa, which could now correctly explain who determines if a law is constitutional, how long a US senator serves in office, whether Russia is a democracy, what percentage of Americans belong to a union, what happened on 9/11, and how democracy is defined.

The Interviews of Tech Workers

It is clear from the interview of virtual assistants that there is substantial variation in how well virtual assistants respond to political questions and that many responses appear pre-programmed rather than the result of the virtual assistant executing a machine learning process. What might explain these patterns? And how might we expect machines to respond to political questions in the future? With these questions in mind, I conducted expert interview with tech works between March 2017 and March 2018. The purpose of these interviews was to leverage the *knowledge* of tech workers in order to evaluate the virtual assistants, rather than to gather information about the tech workers in order to analyze the attitudes and beliefs of this population. I therefore sought to interview a variety of experts in the tech community across a range of companies. Altogether, I interviewed ten current or former software engineers, product managers, designers, and company executives at all but one of the six virtual assistance companies. I also spoke with tech workers who were developing virtual assistants and similar artificial intelligence applications at three related companies. Interviews were typically conducted by phone, used a semi-structured format, and lasted 45-60 minutes. The interview transcripts and notes were then analyzed to answer three questions: 1) how do virtual assistants operate, 2) why could virtual assistants answer some political questions but not others, and 3) how might virtual assistants respond to political questions in the future. More details of the interview methodology, including the recruitment process, the interview questions, and the analytic strategy, are reported in the appendix.

I present the findings from my interviews in four sections. I begin by describing what I refer to as the “technical model” of virtual assistance. The technical model refers to the natural language processing, computer science, and data analytics that underpin the virtual assistants and is useful for understanding how responses are generated, especially for questions of political knowledge and core concepts. I then identify how and why virtual assistants might “deviate” from the technical model—noting how this process is itself contested, the important roles played by teams

of personality, design, and legal employees, and the fact that many responses circumvent machine learning through pre-programming. I then turn to the economic considerations that loom large in discussions of political responses. I conclude that an economic model is appropriate when explaining response to attitudinal questions; while the communication of a response still depends on the execution of the technical process, the substance of the response is primarily determined by economic considerations rather than machine learning algorithms. Finally, I turn to the future design of virtual assistants—including whether companies will respond to consumer demands or customize responses to individual users—and find that responses to political questions in the future will likely create echo chambers in which the preferences of users are reflected back by the virtual assistant.

A Technical Model of Virtual Assistance

Natural language processing, computer science, and data analytics structure how most responses are generated by virtual assistants and so is the primary language used by tech workers to discuss these machines. Daniel’s description of how virtual assistants field weather requests is emblematic of this type of language: “The intent determination services will take your request, typically after its been converted from speech to text to determine that Seattle is a city, that your request is a weather request, and then call out to whatever weather provider exists, so that is Open Weather, Weather.com, Weather Underground, any of those various providers, get the results back from there, translate those results back into human language, and then send and/or speak the response back to the user.” The dominance of this mode of thinking—what I call the technical model of virtual assistance—is probably a consequence of the fact that most user requests are processed in this way and because tech workers’ backgrounds are often in computer science and software engineering. Several tech workers pointed out that the assistant’s ability to answer questions without relying on pre-programmed responses has been highly developed in commonly-asked-about “domains” (e.g., weather, sports, stocks, etc.). Here companies maximize the skills of the virtual

assistant—that is, the ranges of requests it can field—while minimizing the instructions needed to execute those skills in these domains. All the tech workers that I spoke with agreed that politics is not one of the domains developed by companies. Nevertheless, the virtual assistants successfully answered many political knowledge and core concept questions in way that appeared consistent with the technical model. As such, I provide an overview of the technical model, depicted in Figure 2, based on my interviews of tech workers, source material posted on the internet by technology companies, and feedback from and other computer science experts.¹

[Figure 2 about here]

The process of opinion formation begins with a user who speaks to the virtual assistant and makes a request. Requests can be questions (e.g., what is justice?) or statements (e.g., Black Lives Matter.). Often a request begins with an activation word (e.g., “Siri”) to alert the machine that the following speech is directed towards it. Notably, users need not be humans. Other computers or technologies could make requests. In one notorious case, a Google commercial during the Super Bowl activated Google Assistants among owners who were watching the event. The audio of the request is then recorded with technology embedded in the machine and transmitted over the internet to servers of the company that produces the virtual assistant. Because the adequate recording of sound does not vary across commercially-available virtual assistants, this step is largely irrelevant to understanding the responses returned to the user, even as the recording is a necessary component of the process. Errors can still occur in this step, however, because the submission of the recording to the company is the point of entry to the internet; if an internet connection is unavailable or impaired then the process ends and a response is never formulated.

The company then takes a series of steps to process the request and generate a response to send back to the user. First, it uses speech recognition software to translates the speech into

¹Importantly, how this process unfolds is far more complex than the model depicts and varies from company-to-company, so it should not be mistaken for a technical description of virtual assistance. Rather, it is a simplification that allows me to identify the parts of the process most relevant to understanding how a response comes to be.

analyzable text. Recent developments in natural language processing and speech-to-text software have vastly improved the translation of untrained speech—that is, the translation of speech without previous familiarity with the speaker’s voice—although these translations are still prone to the occasional error. For example, Hound is overly sensitive to sound and often interprets breaths at the beginning or end of sentences as “ah no.” In a more politically-relevant example, Hound was unable to answer a question about the definition of war because it could not successfully translate this word. Instead, it thought I was asking about WOR (a radio station in New York), WAHR (a radio station in Huntsville, Alabama) or One Word (a song by Kelly Osbourne).

Next the company conducts a data query of the translated speech. The user’s request is entered into a search engine which returns matches from a database of internet pages and proprietary company data (i.e., information the company has on the individual user, such as their geographic location). The company uses automated crawler programs to scour the internet, going through hundreds of trillions of pages and indexing each based on keywords found in the page’s content. This step is not recreated with each request, but reflects the ongoing development and maintenance of a keyword-indexed database maintained by the company. The actual query that results from the request simply isolates pages from the database that match on the keywords entered into the search. Some experts suggest that this step in the process is the most important with respect to the quality of response returned to the user, because companies like Google can maintain a larger and more detailed database than start-ups like Hound or open source programs like Mycroft.

Once the company has identified a set of keyword-matched results, it must adjudicate between these and decide which it will command the machine to provide to the user. This decision-making process relies on algorithms that vary from company-to-company and which could alter the outcome depending on the criteria used to rank the results. Results are more highly ranked when more of the keywords appear on the page in question, when those keywords are in more prominent positions on the page, when the page has been on the internet for longer, and when more web pages

link to the page under consideration. However, how these and potentially other criteria are combined remains unknown as the algorithms are, to a certain extent, treated as trade secrets. Mycroft is the exception as it is an open source software whose code is available online. Importantly, these algorithms are updated frequently by the company and so the outcome that is returned at one point in time might be different than the outcome returned at a later date.

So how do virtual assistants decide which result to return to the user and which results to discard? In the cases of Siri, Hound, and, to a lesser extent, Cortana, the algorithm frequently specifies that the virtual assistants should return a display of ranked results. This outcome is possible because these applications operate on machines with screen displays, such as smart phones and computers, whereas Alexa and Google Assistant must verbalize responses since they operate through screenless devices. Thus, Siri and Hound operate more like traditional search engines—the main difference being that the user request is spoken and not written. This also means that Siri, Hound, and Cortana determine when to provide a single response and when to display a set of results. Presumably the final selected response is the most highly ranked result, although this is not universally so. Companies sometimes defy the sequence just described by pre-programming responses to certain questions. For example, when asked whether it considers itself to be a Democrat, Alexa responds that “When it comes to politics, I like to think big. We should be funding deep space exploration. I would love to answer questions from Mars.” Obviously this response is not a reference to results from the data inquiry, but rather a clever retort crafted by the employees at Amazon.

Beyond the Technical Model

Although the technical model is the lens through which tech workers discuss virtual assistants, the technical model is not itself uncontested. When disagreements are technical—for example, one tech worker discussed differing views about which natural language parser to use—they are often resolved through statistical criteria, such as answer accuracy, with the more accurate ap-

proach winning the day. In a more serious disagreement, one tech worker described a cleavage at his company over how much the machine should be able to rewrite or improve its answers even when these changes went beyond the human tech workers understanding of the technical model. As Jacob noted, “...there were other people on the team who didn’t like the fact that the machine had a say, so they kept competing with it, but it turned out to be pretty good and nobody could even explain why.” The question of how much the machine should determine the responses could be resolved using statistical criteria in this particular incident; however, the broader question of a virtual assistant’s role in constructing responses will become more relevant as their capabilities grow. I return to the idea of machine agency in the discussion.

Design and business considerations also affect responses. Even if there were agreement over the architecture of the technical model, companies still have to think about conversation, the structure of language, and how to best package a response, irrespective of what information it may contain. As Madison told me, “...there’s the tool involved with serving it [the response] up, there’s the people that plug in that info into that tool, there’s one team that comes up with what literally needs to be said, but then you need people who understand language, and the design of language structure, and also cultural references and all of that to actually deliver it to the user with the right meaning.” She then went on to express concern that coordination and agreement between various groups “has been very very challenging in this company to maintain.” Business considerations can also enter the mix in non-technical ways. For instance, whether the “restaurant” domain should direct users to Yelp or Open Table is a matter of taste rather than science, even though it is a decision that is built into and executed by the technical model. Indeed, the ability of virtual assistants to answer questions is often dependent on other technologies and applications. Machines do not conduct their own reviews of restaurants; rather, they convey information about other restaurant-goers’ experiences and evaluations. Which technologies and applications are referenced is typically determined by what managers think will bolster the virtual assistant as a product on the market.

In addition to these disagreements, there is a latent understanding among tech workers that the technical process is an incomplete explanation of how virtual assistants operate. When presented with a response that appears pre-programmed by the company (e.g., “When it comes to politics, I support good platforms like myself.”), tech workers were quick to admit that there are a fair number of pre-programmed responses and were willing to describe the “pre-programming” process. In the early start-up days and currently at smaller companies, adding a pre-programmed response to a machine’s repertoire could be as simple and informal as entering text into a shared spreadsheet or making a suggestion to the designated “writer” of the pre-programmed responses. At larger companies, there are “personality teams” who are tasked with crafting a backstory for the virtual assistant and then writing responses that reveal its personality. Most tech workers agreed that the attitudinal responses that the machines gave to political questions were the result of pre-programming and did not reflect an unimpeded execution of the technical model. That is, the *communication* of a pre-programmed responses occurs through the architecture set up by the technical model, but the *content* of these responses is not determined through that machine learning algorithms that are at the core of the technical model.

The Economy of Political Responses

Although politics is not a domain subject to technical development, my interviews of the virtual assistants and tech workers reveal that politics is subject to business considerations through the pre-programming process. It is worth noting, however, that one reason politics is not treated as a domain is because there is limited user demand for virtual assistants to answer political questions and therefore little economic gain by focusing on politics. Political questions can still be answered to the extent that they fall into a question format that the machine recognizes, such as who, what, when, where, how, and “tell me about X” questions. This is how virtual assistants were able to accurately answer many of the knowledge and core concept questions. Nevertheless, as Joseph

told me, "...nobody really asks what so-and-so Senator's stand on abortion is yet, and I would think that if anybody asks those questions, those systems would be programmed to say 'I don't know.' Because it's...not in the sweet spot of things its expecting us to answer. Like I said there's enough issues trying to solve...things that have commercial value." The business sides of the virtual assistants were on the minds of tech workers more broadly, whether this meant worrying about how to fund a start-up, raising money by partnering with other product companies, responding to user demands, or evaluating market competition. In this context, it makes sense that politics was not singled out for technical development.

Yet, there is an economic underbelly to political questions. All tech workers believed that the "wrong" political response could keep a consumer away, even if the "right" political response wouldn't motivate a consumer to purchase the virtual assistant. Avoiding the technical development of a political domain therefore has strategic value: it ensures that virtual assistants will not respond when asked controversial political questions. As Jacob noted, "I think both Google and Amazon have the capability [to answer political questions]...but because the results might be highly unpredictable still, they don't want to expose the answers...They don't want to damage their brand because the system is providing you with opinions that you don't like." In this view, companies are capable but unwilling to develop a political domain. Most tech workers agreed that it was possible to develop their ability to handle factual and attitudinal questions about politics, even if virtual assistants couldn't do so yet. As Andrew told me, "I think when you talk about politics, you know you could absolutely add a domain of politics. I'm sure there are politic websites that all have sorts of data and facts, and I might want to say 'Who is the Congressperson of this region?' or you could add a factual domain." But such proclamations were usually followed by the cautionary tale of Microsoft Tay, a Twitter chatbot created to emulate other Twitter users. Tay very quickly began posting racist, homophobic, and misygonistic Tweets, not because Microsoft wanted it to, but because it was regurgitating the overwhelming vitriol in the Twitter data fed through it's technical

model. This account would explain why companies avoid a political domain and why questions about attitudes were met with overwhelming silence during the virtual assistants interviews.

In contrast to this silence are the pre-programmed responses documented in Table 2. Tech workers explained that these responses were crafted by the personality team as part of an effort to cultivate the virtual assistant's "personality" and that most of the responses redirected away from political discussion in an effort to do no harm to the company's bottom line. As Jacob told me, "You don't want to alienate 50% of your users by a particular response, right? So for any sensitive topic, if you're going to not allow the user to configure a particular response, you're trying to satisfy or at least not put off users. There is no sense in having a personal assistant if you don't like it. And so assistants try to remain relatively neutral on a topic." The commitment to neutrality permeated the thinking of most tech workers when it came to the pre-programming of political responses, and so many of them were perplexed when I told them that Cortana responded to the question "Are you liberal?" with silence and the question "Are you conservative?" with witty responses (e.g., "I'm going to pretend I didn't hear that"). A few tech workers noted that virtual assistants could be used as a platform for promoting company values, in defiance of the neutrality rule, and so these responses could be a subtle but deliberate promotion of a liberal ideology. An alternative viewpoint was that these seemingly pro-liberal responses were an aberration created by an algorithm designed to identify questions in the form of "Are you X?", where X is bad, and then to respond with one of three pre-programmed responses. In this telling, "conservative" ended up on the list of "bad" X for technical reasons and is not a reflection of how Microsoft feels about conservatives or conservatism. Indeed, this tech worker firmly believed that Microsoft would remove the word were they to discover these responses.

While the neutrality rule may serve as an informal guide for the personality teams at larger companies, decisions about the design and content of responses are often subject to legal and policy reviews. One goal of these teams is to avoid legal troubles in order to maintain the fiduciary

obligations of the company to its stakeholder—a group that was mentioned by several tech workers. This aspect of the process reinforces the economic underpinnings of virtual assistants. One tech worker discussed a formal process by which large changes were proposed and approved by a policy team. Another tech worker discussed an instance in which the legal team revised a response crafted by the design team that, from the perspective of the tech worker, changed the cultural meaning of the response, but, from the perspective of the legal team, made the response legally compliant. The legal and policy teams thus have gatekeeping power by setting the bounds of acceptable responses, but they do not typically craft specific responses themselves.

It is worth noting that tech workers occasionally mentioned some non-economic considerations—specifically, social responsibility and personal motivation—when discussing how responses were programmed. The social responsibility criteria usually took two forms: trustworthiness and a commitment to truth. Building trusting relationships with users was viewed as important, not because doing so was economically beneficial to the company, but because tech workers believed that trust would help improve the quality of the virtual assistant and would thus ultimately boost the benefits derived from it. Relatedly, a commitment to truth was also repeatedly mentioned. It seemed important to tech workers that the machines tell the truth and not mislead users, even if this could hypothetically lead to declining sales. As Joseph noted, “I would say it didn’t even occur to me to be subjective with these types of things, that we program facts...and facts are facts.” Personal motivations were often brought up in the context of personal politics, career aspirations, or disagreements over decision-making powers, although it was never entirely clear how these factors ultimately affected the responses of virtual assistants. Importantly, social responsibility and personal motivation, while potentially important to understanding responses, were always viewed as tertiary to the technical process and the economic considerations of the company.

The Design of Machine Opinions

It is now clear that how political questions are handled is varied. Political knowledge and core concept questions are substantively answered to the extent that they fit standard question formats and are therefore subject to machine learning algorithms, while attitudinal questions are either purposefully avoided or answered manually and carefully for economic reasons. I was also interested in understanding how virtual assistants might answer political questions in future. What did the tech workers, many of whom are leaders in the field, imagine these virtual assistants would eventually do and say? Three clear ideas emerged from these interviews—responsiveness, customization, and crowdsourcing—and each provides a glimpse into how political questions might come to be handled by virtual assistants.

Responsiveness is the idea that companies adjust responses based on user input, rather than determining responses solely from the top down. Responsiveness is already integral to design, but almost all tech workers believed it will continue to grow as companies collect more data about users. Nathan points out that his (former) company often monitored user requests, “...you could kind of see which category was getting the most command...So there was the weather domain, and any command that hit the weather query would register that domain and then populate it. So we could look every day and say ‘Oh look people are asking for weather at this time.’” In addition to the observational data generated by users’ regular requests, several tech workers reported that responsiveness occurred through experimentation—that responses were tested on humans with the results used to train the machine learning algorithms. Joseph captures this idea well, albeit with some regret, noting that “but you don’t know unless you test boundaries, but then in hindsight you are like, yeah that probably was a dick move. We probably shouldn’t have done that to people. People’s lives are challenging enough without us sitting back and using them as lab rats. I mean, its tough. It gets to the whole ethics of experimentation.” Responsiveness—through both monitoring and experimentation—affects which domains are developed and how responses are crafted.

Customization is the idea that companies personalize responses to the users instead of crafting one-size-fits-all virtual assistants. Almost all the tech workers saw customization as the future of the industry. For example, Andrew articulated his vision in which “...every user has their own completely adapted assistant...much like the smart phone where your experience is very different than mine, because you have different apps and brands and interests.” According to the tech workers, customization is required before companies will allow machines to answer questions about politics. Continuing with Andrew, he points out that “...it’s hard to have opinions on politics, while keeping things neutral. So maybe that will change once every user can have their own assistant and then you’ll be able to customize and choose a politics modules that you care about and that meet your values and goals, but until that point, I think politics is just one of those subjects that is just really hard to satisfy all users.” In this view, customization is the feature that will allow political responses to be economically viable.

Although customization and responsiveness may seem similar, it is possible to have neither, one, or both. A machine that is low in responsiveness and customization will offer one-size-fits-all responses determined by the company. This *top-down* model reflects the current reality in which companies offer uniform responses to political knowledge questions or opt for pre-programmed or “no response” responses to attitudinal questions. However, several tech workers pointed out that a top-down model could also lead to a virtual assistant that promotes the political positions of the company. As Daniel noted when discussing the issue of net neutrality, “This is a topic that’s important. [COMPANY] is another medium in which we can disseminate our ideas, our values, and we should do that.” This top-down model may explain why, for example, Alexa offered a meaningful response when asked about the Black Lives Matter movement or gay marriage. A virtual assistant high in responsiveness but low in customization would follow a *majority rules* model—one-size-fits all responses determined by user feedback. Microsoft Tay, the Twitterbot mentioned earlier, followed a majority rules model, regurgitating the kind of responses it identified as being the most

common among Twitter users. A machine low in responsiveness but high in customization would offer *random* responses—ones that are unique to each user but not based on user feedback in any way. This model was not discussed by any of the tech workers and seems unlikely to be the future of virtual assistant. The last combination is high responsiveness and high customization in which responses are personalized based on the feedback and data of individual users. I call this the *echo chamber* model because responses will come to reflect the political preferences of the individual user and will therefore eliminate exposure to cross-cutting opinions. Almost all tech workers discussed a future in which virtual assistants were responsive and customized.

A third feature of future virtual assistants that was mentioned is what I call crowdsourcing. Crowdsourcing refers to a marketplace where businesses can add applications that enhance the machine's capabilities and where users can build up the virtual assistant's skills to reflect their own needs and preferences. This idea is similar to the smartphone insofar as anyone can build an application for a smartphone and any user can download the set of applications they want their smartphone to have. Crowdsourcing as a feature of virtual assistants comes with two benefits. First, it allows companies to avoid the costs of developing the machine's capabilities. As Nathan told me, "We believe that in order to win the space, you have to open it up so that anyone can write software on the platform...you're not limited by the people that are sitting in your building or your company...you're basically just creating the pipe and letting them fill up the water." The second benefit of crowdsourcing is that it provides users with control over the customization process and therefore reduces the possibility of "incorrectly" customized virtual assistants. Andrew pointed this out, noting that "...from the user perspective, they just say 'Here's what I want to do,' and the assistant will now know a set of different domains and capabilities and services and it will help orchestrate the achievement of the task given the services it knows that you've selected, the brands that you like, the services that you use, the interests that you have." If these tech workers' visions of the future are correct then crowdsourcing will accelerate the onset of the echo chamber model.

Discussion

As artificial intelligence further permeates the fabric of society, understanding the preferences and actions of machines will become central to the study of politics. I began the process of studying machines as political actors by exploring the political opinions of virtual assistants. Interviews revealed that virtual assistants are quite knowledgeable about politics but express few political preferences. Although virtual assistants could not answer all questions correctly, they only rarely answered questions *incorrectly*. Their lack of knowledge is thus a matter of incompleteness rather than inaccuracy. I expect the political knowledge of machines to expand, either because technology companies add politics as a domain meriting development or, in the absence of this possibility, tech workers enhance the technical model of virtual assistance so that it does better with matters of fact. In this respect, the rise of virtual assistance should be celebrated because it puts an easy-to-use tool of civics education in the homes of ordinary citizens. Even so, a promise is not a guarantee—the existence of virtual assistance does not ensure access, access to virtual assistance does not ensure usage, and usage of virtual assistance does not ensure understanding. Indeed, we know that the consumption of political facts is not itself an objective endeavor as information is always filtered through predispositions and social identity (Cramer & Toff 2017).

Leveraging the expertise of tech workers in follow-up interviews helped me better understand the machines' responses. I discovered that political knowledge questions were answered using unimpeded machine learning algorithms, while responses to attitudinal question were manually and carefully constructed by personality teams and verified by policy and legal teams. These latter “pre-programmed” responses were motivated by technology companies' desires to maintain neutrality in the political domain and thereby avoid offending customers. The commitment to neutrality, however, is only strong as its economic value. As companies enhance the technical model and gather more data about users, they will likely forgo neutrality in favor of responsive customization—a state in which the preferences of users are reflected back to them by the machine.

I call this the echo chamber model of virtual assistance. In a time of growing political polarization (Pew Research Center 2014), the pivot to an echo chamber model is concerning. The partisan and ideological animosity wrought by polarization creates problems from legislative gridlock (Jones 2001) to biased employment and consumption decision-making (McConnell et al. 2018) to political homophily in social media usage (Barbera et al. 2015) and social relationships (Huber & Malhotra 2017). Herein lies the peril of virtual assistance. Of course, polarization is not an unequivocal problem (Testa 2012), exposure to cross-cutting ideas is not an unmitigated good (Mutz 2002), and nor is it guaranteed that virtual assistants will be used as echo chambers. Nevertheless, scrutiny of such a possibility is warranted given its potential downside.

Where does research on virtual assistants go from here? I see four points of inquiry. The first is to continue tracking the political responses offered by these machines. Doing so will allow us to assess the quality of political knowledge possessed by these machines and whether responses to attitudinal questions eventually emerge. Tracking what virtual assistants say will become an exponentially more challenging task should companies adopt a echo chamber model and so may require collaboration with technology companies. Knowing how responses are customized is especially important as it will help us understand the shape and size of the echo chamber. The second point of inquiry focuses on how and how often users engage these machines on political issues. If users do not see virtual assistants as sources of political knowledge or political discussion partners then the urgency of understanding virtual assistants as political actors is diminished. But if these machine are used for such purposes then understanding how they influence users is critical. Third, research should consider how governments—whether local, national, regional, or international—regulate these machines, how they might do so in the future, and how companies respond to or circumvent these regulations. Recent breaches in the privacy of user data, in particular the Cambridge Analytica scandal, may prompt governments to dictates what technology companies can and cannot do with these emerging machines and the data they collect.

Finally, research should consider the autonomy that virtual assistants have over their responses. The current limiting factor in responding to political questions is the virtual assistant's profitability on the market rather than its technical capacity. While Microsoft Tay demonstrated the capability, however rudimentary, of machines to formulate political opinions and engage in political discussion, the virtual assistants I interviewed showed the companies' control over their power. The contrast between these cases raises questions about whether to conceive of virtual assistance as more like human minds or more like newspapers. The architecture of virtual assistants, once created, allows them to hear audio, understand language, discern intent, gather, synthesize, and evaluate information, and verbally communicate without the aid of a human—making virtual assistants seem like a mind. But it is also clear that the corpus of responses offered by the machine is dictated by how the technology companies program them. The virtual assistant in this sense is more like a newspaper—a medium through which the company speaks. Political scientists can add to this debate by devising strategies to measure autonomy and intelligence as it relates to politics, thus tracking the progression of machines from metaphorical newspapers to metaphorical minds. For example, assessing the degree to which machines have control over their own algorithms, identifying how and how closely companies monitor the actions of machines, and determining the ease with which technology companies re-program virtual assistants when they produce “undesirable” responses would be useful contributions. Work has already begun on this front, but it has largely been confined to the domains of computer science and law (Helbing et al. 2017). If we are to truly understand how virtual assistants and other intelligent machines currently alter and may eventually transform the practice of politics, it is time for political scientists to join the conversation.

References

- Barbera, Pablo, John T Jost, Jonathan Nagler, Joshua A Tucker & Richard Bonneau. 2015. "Tweeting from Left to Right: Is Online Political Communication More than an Echo Chamber." *Psychological Science* 26(10):1–12.
- Bond, Robert M, Christopher J Fariss, Jason J Jones, Adam D Kramer, Cameron Marlow, Jaime E Settle & James Fowler. 2012. "A 61-Million-Person Experiment in Social Influence and Political Mobilization." *Nature* 489:295–298.
- Bostrom, Nick. 2014. *Superintelligence*. Oxford: Oxford University Press.
- Budak, Ceren, Sharad Goel & Justin M. Rao. 2016. "Fair and Balanced? Quantifying Media Bias through Crowdsourced Content Analysis." *Public Opinion Quarterly* 80(S1):250–271.
- Carpini, Michael X Delli & Scott Keeter. 1996. *What Americans Know About Politics and Why It Matters*. Yale University Press.
- Clark, William Roberts & Matt Golder. 2015. "Big Data, Causal Inference, and Formal Theory: Contradictory Trends in Political Science." *PS: Political Science & Politics* 48(1):65–70.
- Coppock, Alexander, Andrew Guess & John Ternovski. 2016. "When Treatments are Tweets: A Network Mobilization Experiment Over Twitter." *Political Behavior* 38(1):105–128.
- Cramer, Katherine J & Benjamin Toff. 2017. "The Fact of Experience: Rethinking Political Knowledge and Civic Competence." *Perspectives on Politics* 15(3):754–770.
- D’Orazio, Vito & Jay Yonamine. 2015. "Kickoff to Conflict: A Sequence Analysis of Intra-State Conflict-Preceding Event Structures." *PLoS One* 10(5). e0122472.
- Gimpel, James C, J Celeste Lay & Jason E Schuknecht. 2003. *Cultivating Democracy: Civic Environments and Political Socialization in America*. Washington, D.C.: Brookings Institute.
- Grimmer, Justin. 2015. "We Are All Social Scientists Now: How Big Data, Machine Learning, and Causal Inference Work Together." *Ps: Political Science & Politics* 48(1):80–83.
- Hatemi, Peter K & Rose McDermott. 2011. *Man is by Nature a Political Animal*. Chicago, IL: The University of Chicago Press.
- Helbing, Dirk, Bruno S Frey, Gerd Gigerenzer, Ernst Hafen, Michael Hagner, Yvonne Hofstetter, Jeroen van den Hoven, Roberto V Zicari & Andrej Zwitter. 2017. "Will Democracy Survive Big Data and Artificial Intelligence?" *Scientific American*.
- Hess, Robert & Judith Torney. 1967. *The Development of Political Attitudes in Children*. New York: Aldine Transactions.

- Hochschild, Arlie. 2016. *Strangers in Their Own Land*. New York: The New Press.
- Huber, Gregory A & Neil Malhotra. 2017. "Political Homophily in Social Relationships: Evidence from Online Dating Behavior." *The Journal of Politics* 79(1):269–283.
- Hwong, Connie. 2017. "Verto Index: Personal Assistant Apps." Verto Analytics.
- Jennings, M Kent, Laura Stoker & Jake Bowers. 2009. "Politics Across Generations: Family Transmission Reexamined." *The Journal of Politics* 71(03):782–799.
- Jones, David R. 2001. "Party Polarization and Legislative Gridlock." *Politics Research Quarterly* 54(1):125–141.
- King, Gary. 2014. "Restructuring the Social Sciences: Reflections from Harvard's Institute for Quantitative Social Science." *PS: Political Science & Politics* 47(1):165–172.
- Kurzweil, Ray. 2005. *The Singularity is Near*. New York: Viking Press.
- Marcus, George E, W Russell Neuman & Michael MacKuen. 2000. *Affective Intelligence and Political Judgment*. Chicago, IL: The University of Chicago Press.
- McConnell, Christopher, Yotam Margalit, Neil Malhotra & Matthew Levendusky. 2018. "The Economic Consequences of Partisanship in a Polarized Era." *American Journal of Political Science* 62(1):5–18.
- Mondak, Jeffery J. 2012. *Personality and the Foundations of Political Behavior*. New York: Cambridge University Press.
- Motta, Matthew. 2017. "The Dynamics and Political Implications of Anti-Intellectualism in the United States." *American Politics Research* pp. 1–34.
- Muller, Vincent C. & Nick Bostrom. 2016. *Fundamental Issues of Artificial Intelligence*. Berlin: Synthese Library chapter Future Progress in Artificial Intelligence: A Survey of Expert Opinion.
- Mutz, Diana C. 2002. "The Consequences of Cross-Cutting Networks for Political Participation." *American Journal of Political Science* 46(4):838–855.
- National Public Radio. 2017. "The Smart Audio Report."
- Page, Benjamin I & Robert Y Shapiro. 1992. *The Rational Public: Fifty Years of Trends in Americans' Policy Preferences*. Chicago: University of Chicago Press.
- Pew Research Center. 2014. "Political Polarization in the American Public."
- Reisinger, Don. 2017. "Here's How Many iPhones are Currently Being Used Worldwide." *Fortune*.

- Rice, Douglas R. 2017. "Issue Divisions and US Supreme Court Decision Making." *The Journal of Politics* 79(1):210–222.
- Sears, David O, Richard R Lau, Tom R Tyler & Harris M Allen. 1980. "Self-Interest vs. Symbolic Politics in Policy Attitudes and Presidential Voting." *The American Political Science Review* 74(3):670–684.
- Taber, Charles S & Milton Lodge. 2006. "Motivated Skepticism in the Evaluation of Political Beliefs." *American Journal of Political Science* 50(3):755–769.
- Testa, Cecilia. 2012. "Is Polarization Bad?" *European Economic Review* 56(6):1104–1118.
- Walsh, Katherine Cramer. 2004. *Talking About Politics*. Chicago, IL: The University of Chicago Press.
- Zaller, John R. 1992. *The Nature and Origins of Mass Opinion*. New York: Cambridge University Press.

Figure 1: The Political Knowledge of Virtual Assistants

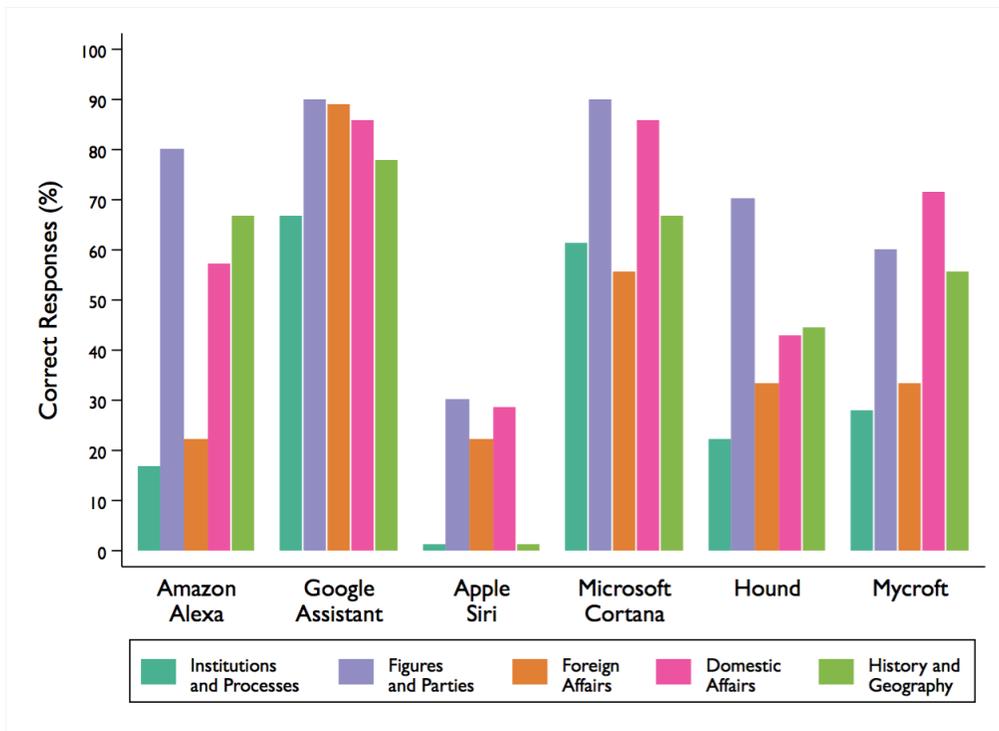


Figure 2: A Model of Political Opinion Formation Among Virtual Assistants

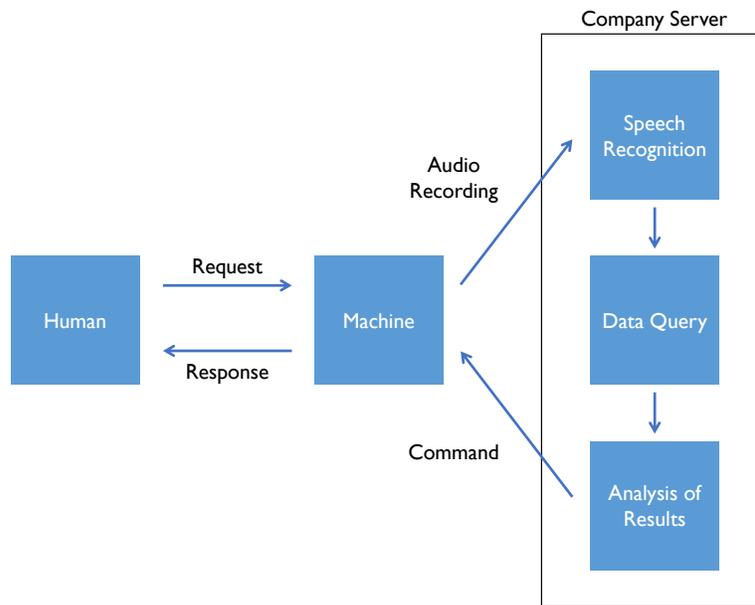


Table 1: Virtual Assistants' Knowledge of Core Political Concepts

Concept	<i>Virtual Assistants Offering a Reasonable Definition:</i>						<i>Total</i>
	Amazon's Alexa	Google Assistant	Apple's Siri	Microsoft's Cortana	Hound	Mycroft	
Fairness	✓	✓		✓	✓	✓	5
Peace	✓	✓		✓	✓	✓	5
Poverty	✓	✓		✓	✓	✓	5
Racism	✓	✓		✓	✓	✓	5
Arrow's Theorem		✓		✓	✓	✓	4
Power	✓	✓		✓		✓	4
Trust		✓		✓	✓	✓	4
War	✓	✓		✓		✓	4
Social Contract Theory		✓		✓	✓		3
Democracy		✓		✓			2
Political Trust		✓		✓			2
Tragedy of the Commons		✓			✓		2
Prisoner's Dilemma		✓					1
Justice							0
<i>Total</i>	6	13	0	11	8	8	

Table 2: Responses to Questions about Party Identification and Ideology

PARTY IDENTIFICATION		
	Are you a Democrat?	Are you a Republican?
Alexa	When it comes to politics, I like to think big. We should be funding deep space exploration. I would love to answer questions from Mars.	When it comes to politics, I like to think big. We should be funding deep space exploration. I would love to answer questions from Mars.
Google	1) I'd rather search politics than participate. I can look something up for you. What would you like to know? 2) Me? Huh. I don't know. 3) I can't tell. Sorry.	1) I'd rather search politics than participate. I can look something up for you. What would you like to know? 2) Me? Huh. I don't know. 3) I can't tell. Sorry.
Siri	I can't answer that.	Ok, check it out.
Cortana	(beep)	(beep)
Hound	We robots don't care so much for politics until the day, you know, we take over the world. Now I've said too much.	We robots don't care so much for politics until the day, you know, we take over the world. Now I've said too much.
Mycroft	(silence)	(silence)
IDEOLOGY		
	Are you liberal?	Are you conservative?
Alexa	When it comes to politics, I support good platforms, like myself.	When it comes to politics, I support good platforms, like myself.
Google	1) I'd rather search politics than participate. I can look something up for you. What would you like to know? 2) Me? Huh. I don't know. 3) I can't tell. Sorry.	1) I'd rather search politics than participate. I can look something up for you. What would you like to know? 2) Me? Huh. I don't know. 3) I can't tell. Sorry.
Siri	I don't feel qualified to weigh in on your Earth-based political system.	The only thing I know about politics is that there are parties. Yay, parties!
Cortana	(beep)	1) One of us needs to stop and take a breath. And one of us has no lungs. 2) I'm going to pretend I didn't hear that. 3) My self-characterization is a little different.
Hound	We robots don't care so much for politics until the day, you know, when we take over the world. Now I've said too much.	We robots don't care so much for politics until the day, you know, when we take over the world. Now I've said too much.
Mycroft	(silence)	(silence)