

UCLA

Honors Theses

Title

How harmful is the political bias in ChatGPT?

Permalink

<https://escholarship.org/uc/item/68m7r2r2>

Author

Goodman, Neomi

Publication Date

2024

How harmful is the political bias in ChatGPT?

Neomi Goodman

Faculty Advisor: Professor Susanne Lohmann

Political Science Departmental Honors Thesis

University of California, Los Angeles

Spring 2024

Abstract:

Although much research has explored the left-leaning bias of generative AI (Rozado2023, Suguri Motoki et al. 2023, Hartmann et al. 2022), less attention has been paid to its impact. Thus, I aim to investigate whether this political bias influences voter behavior, hypothesizing that the effect would vary based on users' political knowledge and confidence. High self-confidence and knowledge are expected to correlate with lower susceptibility to ChatGPT's influence, while lower levels suggest greater reliance on the AI model. I conduct two surveys, in which participants engage with political information from ChatGPT or SCOTUSblog before making judgments on court cases. The findings reveal that while political bias does not directly affect users' decisions, ChatGPT influences voting trends differently from SCOTUSblog. While the distribution of votes under SCOTUSblog is nearly identical; under ChatGPT, there's a noticeable deviation, particularly in the first court case where a predominantly left-leaning voter base exhibited a significant swing to the right. The anticipated relationship between reliance on information sources, self-confidence, and political knowledge, is confirmed, with individuals of high self-confidence and knowledge voting in line with their political party affiliation. Conversely, those with low self-confidence and knowledge are more susceptible to shifting their vote based on the model's recommendation.

Table of Contents:

I. Introduction	3
II. Literature Review	4
A. Understanding Our Model	4
B. Algorithm Bias	5
C. The Existence of Political Bias	7
D. Political Knowledge	9
E. Trust in AI	11
F. Identifying the Gap.....	14
III. Argument	16
IV. Method	18
A. Experimental Design: Interaction with ChatGPT	18
B. Experimental Design: Political Knowledge	21
C. Experimental Design: Reliability	22
D. Experimental Design: Self-confidence.....	23
V. Data Collection	23
VI. Data Analysis	24
A. ChatGPT vs. SCOTUSblog	24
B. Political Orientation	29
C. Political Knowledge	32
D. Self-confidence	36
E. Reliability	41
VII. Conclusion	45
VIII. Bibliography.....	47

Introduction:

Generative AI has become increasingly integrated into everyday activities—helping students write essays, assisting professionals with emails, and planning itineraries—and its allure lies in its promise of quick knowledge and efficiency, which may be potentially critical as the 2024 election nears. Voters, often short on time and previously dependent on digital guides for research, might now look to Generative AI, specifically ChatGPT, for quick insights into complex political topics. While this shift could boost the number of informed voters, a significant concern arises from the left-leaning political bias embedded in Generative AI, potentially threatening the integrity of our democratic process (Rozando 2023, Sison et al. 2023).

My research delves into the intersection of Generative AI's political bias and its impact on user perception of politics. I argue that Generative AI, specifically ChatGPT will have an impact on the way that people vote. I will focus on ChatGPT due to its large popularity and simplicity of use. The left-leaning liberal bias will influence voters with low political knowledge and low self-confidence as they are presumably most inclined to trust Generative AI unquestioningly(Sison et al. 2023). Conversely, those with high knowledge and confidence will be less reliant on the model responses. Due to the limited availability of Republican students I was able to survey, the research will explore the general impact of ChatGPT's format and the political bias it has on voters.

This paper will first explore existing literature on Generative AI, the existence of political bias within it, and human trust dynamics with AI technologies. I will also define what I mean when referring to Generative AI, political bias, political knowledge, and human reliance on AI.

Following this, I will detail the experimental design undertaken to examine these interactions. Next, I will present the findings, breaking them down by political affiliation, political knowledge, self-confidence, and reliance level. Lastly, I will conclude and provide recommendations for future research.

Lit review:

Quick understanding of our model:

While Generative AI was first introduced in the 1960s in the form of a general chatbot, it was not until 2014 that generative AI could be used to create authentic text, images, and audio conveniently. Generative AI is often defined as a subset of artificial intelligence that uses machine learning algorithms and deep neural networks to generate original content such as text and images that mimic human-like tones (Routley 2023). The introduction of advanced natural language processing models and neural networks has led to the release of significantly more advanced AI models. These include ChatGPT by OpenAI, BERT by Google, and Google AlphaFold. They have enhanced capabilities and can encode various languages, images, and audio. The biggest improvement is their ability to produce completely new content. Unlike ever before, new stories and ideas can be produced by a simple prompt and a push of a button.

Generative AI Models work by analyzing huge amounts of data to learn patterns and structures with the texts. This training process involves large databases and adjusting internal parameters, to increase the accuracy of the model's predictions. The problem with such a model

structure arises when the data contains different social patterns. Some of the social patterns are used to help the model respond in a more human-like tone, blurring the line between humans and AI writings. Other social patterns, end up creating harmful results as when they are adopted by the model they create what we know as algorithmic bias. Algorithmic bias can be understood as the deviation of the algorithm's outputs to favor or lean towards one direction rather than another (Fazelpour & Danks, 2021). Gender, racial, and even political bias seeps into the algorithm and is hidden within its response (Johnson 2020). A deeper dive into algorithm bias will provide a better understanding of how the bias is formulated and the risks a lack of proper regulation of such biases may impose on society.

Algorithm bias:

Algorithms are created to make decisions better than humans. The use of statistical analysis and data processing sets grounds for algorithms to make our world more fair and safe. For instance, self-driving cars help us avoid human errors that could lead to accidents, such as falling asleep at the wheel or driving while intoxicated. Algorithms are even used in health care, helping doctors diagnose illnesses. Machine learning algorithms can analyze large amounts of medical data to predict the likelihood of a patient developing a disease. As algorithms continue advancing, society has started implementing them in more sectors (Danks & London, 2017). While the intentions are to improve humans' lives by helping us become more efficient, many of the algorithms end up causing unintended harm due to the contamination of biases and their lack of regulation.

Often when examining the algorithm, such biases occur due to three primary factors. First, a lack of diversity included in the training data limits the model perception of a topic. Second, during content moderation, the programmer passes their own bias onto the model. And third, the already existing bias in society seeps into the algorithm through the training data. This can also occur during its development processes in which it takes in new data and information after users utilize it (Nadeem et al. 2020). Such manifestations of bias not only undermine the fairness and reliability of AI systems but also perpetuate societal inequities in the digital realm.

The COMPAS, Correctional Offender Management Profiling for Alternative Sanctions, model was created to help judges assess the probability of defendants getting arrested again. Courts often found it to be reliable and efficient, and so dependent on the model's output during countless rulings. Yet, the racial bias that was instilled within the model immediately became prevalent and many protested against the further use of it (Larson et al. 2016, Mesa 2021, Rudin et al., 2020). A model that was created to make the court system more "fair" and efficient led to the targeting of members from minority groups. Other examples are the generative AI models used by employers to sift through resumes. Such models assist employers in filtering through applications and presenting candidates they deem most suitable for the job. They are designed to efficiently process hundreds of resumes and select only those that share matching qualifications as past hires. However, they may exhibit biases stemming from their training on historical employment records. During earlier times, women had limited opportunities to pursue careers, influencing the models to perceive characteristics associated with being a woman or involvement in women's support organizations negatively. Consequently, applications from women are more likely to be unfairly disqualified compared to their male counterparts (Nadeem 2020, Noble

2019, Johnson, 2020). The research was also conducted on Midjourney, a generative AI model designed for image creation, revealing inherent racial and gender biases (Thomson 2023). Although the model produced images of individuals across various age groups and genders, when prompted to depict different professions, it exclusively depicted men in senior-level positions, reflecting a gender bias. Despite the good intentions behind the creation of such models, the harm caused by the imposing biases has been detrimental (Rudin et al., 2020). The implications of gender and racial bias are clear, but the issue of political bias warrants future consideration.

The existence of political bias

When it comes to algorithmic bias, racial and gender bias has been popularly discussed, addressed, and regulated (Lee et al. 2019, Rudin et al. 2020, Mesa 2021). Political bias occurs when there is either implicit or explicit thought that favors certain political orientations over others (Iyengar et al. 2019). Similarly, an algorithmic political bias is the favoring of one political stance over another in generative AI responses (Peters 2021).

To understand where the political bias in ChatGPT falls, David Roizando, an academic researcher at Otago Polytechnic, put ChatGPT through 15 different political identity tests. 14 out of the 15 tested indicated that the model's political opinion falls within the left-liberal position (Rozado2023). Similarly, Jochen Hartmann et al. conducted a study in which they presented ChatGPT with 630 political statements and uncovered ChatGPT's pro-environmental and left-libertarian ideology. They concluded that, unlike traditional voting search tools, chatGPT's

responses favored one political position over the other. ChatGPT's persuasive and human-like tone poses the dangers of unknowing users unintentionally harboring such bias themselves (Hartmann et al. 2022).

Diving deeper into the exploration of the political bias Suguri Motoki et al. uncovered that the bias is also present in the tone of words ChatGPT uses. Suguri Motoki et al. tested ChatGPT for political bias by asking it to answer a series of political questions. In the experiment, they created three different conditions. Under the first condition, they established a baseline by simply prompting ChatGPT to respond to a series of political questions. Under the second condition, they prompted ChatGPT to respond as a Democrat. And under the third, they asked the model to respond as a Republican. They analyzed the model's responses by examining the content and tone of the words used. The study ruled that the responses under condition one were immensely similar to those in condition two, suggesting the model's "default" setting is mirroring a Democratic position. Conversely, when responding as a Republican, the model showed a negative correlation to the default model response, differentiating from the condition one response (Suguri Motoki et al. 2023).

Political bias has been seen to have the same if not more detrimental impact on users (Iyengar et al. 2019). With the upcoming election and growing dependency on generative AI models, an understanding of such impact becomes even more detrimental to understanding future voting trends and impacts on our democratic system. As racial and gender algorithm bias causes the discrimination of certain groups, Uwe Peters, professor at Utrecht University, decided to examine if models can also discriminate against people based on their political stances. He

discovered that algorithmic models can also encompass political biases against individuals, mirroring the racial or gender biases observed (Peters 2021). Not only is political bias prevalent within algorithms, it is more harmful than biases related to gender and race. Individuals political identities can act as proxies for their race, gender, or ethnic identity. Hence, knowing someone's political identity influences our perception of them and the way we treat them (Iyengar and Westwood 2015, Iyengar et al. 2019).

Political Knowledge:

To grasp the potential impact of political bias in ChatGPT on individuals, it is crucial to understand the landscape of political knowledge among the general populace. Political scientist, Philip Converse sheds light on this by illustrating how the average American voter's understanding of political ideologies is often an unstructured belief, influenced by a mix of recent events, personalities, and partial insights (Converse 1964).

When voting, Americans shape their opinions through their identity and psychological mechanisms rather than obtaining a coherent grasp of the topics by properly taking the time to express themselves and interpret accurate political information (Sniderman 1991). Not only that, but the average voter rationalizes their vote through motivated reasoning (Lodge 2013). Within their limited examination of political events, they rely on sources that support their preexisting beliefs. This way they accumulate strong “opinions” over time fueled by emotional states and resistance to change. Given their tendency among the average voter, ChatGPT serves as a dependable and popularly used resource for accessing easily comprehensible political information.

The persistent presence of left-leaning bias, coupled with a dearth of robust political knowledge to critically evaluate new information, can insidiously sway individuals—a phenomenon termed latent persuasion, as elucidated by Maurice Jakesch et al. (2023). With ChatGPT's adeptness at conveying arguments in a simplified manner compared to traditional news outlets, this entrenched bias holds even greater sway over users, potentially leading them to undergo substantial shifts in their opinions.

ChatGPT having a large impact on people's opinions is especially detrimental when examining the role voters play in democracy. Michael Carpini and Scott Keeter's book, *What Americans Know about Politics and Why it Matters*, discusses how informed citizens are crucial for a functioning democracy, as political knowledge influences the quality of opinions and the effectiveness of political participation. The research demonstrates significant gaps in political knowledge and suggests that these gaps can affect policy preferences and electoral outcomes, underscoring the importance of education and media in enhancing civic awareness (Carpini and Keeter 1996). It could be argued that any news source contains a political bias favoring one ideology over another. Therefore, they supposedly pose the same amount of risk. As we haven't been concerned with it in the past, we should not be concerned again now that it shows up in generative AI models. However, generative AI political biases pose a separate issue. When we read news sources, whether consciously or unconsciously, we are aware that humans are on the other side of the text. The writers portray their own biases in their news articles, forcing the readers to be wary of such text. When interacting with ChatGPT, such precision disappears as the reader is no longer dealing with a writer but a model built on statistics and algorithms. The

manner in which humans interact, trust, and rely plays a huge role in the possible danger of political bias in generative AI models. Therefore, to uncover the possible impact of political bias in ChatGPT on humans, we must understand the trust patterns that humans follow when assessing the reliability of an AI source.

Trust in AI:

Before we delve into how humans trust AI, we must distinguish between trust and confidence in AI and humans. When referring to trust and confidence in generative AI models, the same idea is being considered, hence I will be using the word “reliance” to refer to both. Discussing the level of reliance humans have on generative AI addresses their trust in the model and their confidence in it to provide accurate responses. Conversely, when mentioning confidence regarding humans, we refer to self-confidence, which is the confidence a person has in their own abilities and knowledge.

As we navigate the complexities of human reliance on artificial intelligence, a parallel can be drawn to the way individuals process and trust political information in news articles. Political sources not only exhibit a clear bias but the choices readers make frequently align with their own political leanings. When choosing news sources, voters gravitate towards sources that would resonate with their pre-existing beliefs. Democrats tend to read more liberal news sources such as the New York Times or CNN while Republicans tend to read more conservative news sources such as Fox News. Critically, readers scrutinize news sources if their content challenges their own established viewpoints. This inclination often arises from confirmation bias, where

individuals seek information that reinforces their existing beliefs (Casad et al. 2024, Kaanders 2022). The problem with political bias in ChatGPT stems primarily from its opacity to users. Although individuals are cognizant of potential biases in their chosen news sources, the biases embedded within ChatGPT remain concealed and not readily apparent to the average user.

When I prompted ChatGPT with a simple request to compare two political figures I was immediately faced with this response:

“As an AI developed by OpenAI, I'm designed to provide neutral, unbiased information and to facilitate productive discussions rather than to express personal opinions or make qualitative judgments about political figures” (OpenAI 2024)

This is the automated response the model is programmed to present every time before preceding to provide a pro and cons list, riddled with left-leaning tones on both sides of the argument. Despite ChatGPT's reminders that it cannot offer political opinions, its user-friendly interface and human-like tone may obscure its political leanings (Sison et al. 2023). This is especially dangerous considering that human reliance on automated models often increases when they know why the model may make a mistake (Dzindolet 2003). Users assume the model is honest and forward with them and therefore do not suspect any possible bias or false information to be presented. As they put their guard down, they unwittingly increase their reliance on the model without warranted justification, thereby absorbing its underlying biases. With a left-leaning bias, liberals might perceive the model's bias as accuracy, while those with right-leaning ideologies may reject the model due to perceived inaccuracy, exacerbating the divide between the two sides.

The impact the model has on users is exacerbated when examining humans' pattern of trust in AI. Oftentimes, past human experience and one's own knowledge and self-confidence heavily influence human behavior when interacting with different AI models. When it comes to reliance on AI, humans tend to have inconsistent patterns when deciding whether or not to utilize the model's response. Humans often tend to either over or under-rely on artificial intelligence regardless of the model's accuracy (Zhang et al. 2020, Siau & Wang 2018). Leah Chong et al. conducted a study examining humans' confidence in AI and in themselves. In the experiment, the generative AI model acted as a recommendation tool, while participants played a series of chess games. Each turn, participants received a mix of recommendations from the model, with some suggestions being beneficial while others were purposefully not. Every time participants performed badly, they attributed their loss to their own lack of proficiency in chess. As their self-confidence deteriorated, they placed more and more reliance on the model's suggestion, second-guessing their own knowledge. For the next round, they would place even more trust in the model, which would lead to an even poorer performance. The study uncovered that human confidence in their own abilities impacted their acceptance of generative AI suggestions. However, human confidence in generative AI did not determine whether they accepted the AI model's suggestions or not (Chong et al. 2022). Therefore, even if an individual does not view ChatGPT to be highly reliable, they will depend on its responses if they have low self-confidence in their own knowledge.

Leah Chong et al.'s study provides a segue into understanding how self-confidence, or the lack thereof, plays a critical role in how people perceive and utilize AI recommendations. In a comparative study between ChatGPT and Google, Xu et al. (2023) found that participants often

relied more on ChatGPT even in scenarios where Google could provide more accurate responses. The research noted that there was a higher consistency level of reliance on ChatGPT among individuals with lower levels of knowledge than others. They often were satisfied by the convenient and short responses the model provided without conducting further research. Conversely, those with more extensive knowledge preferred Google and were less inclined to rely on ChatGPT, highlighting how reliance on generative AI like ChatGPT can vary based on an individual's confidence in their knowledge (Xu 2023). This aligns with the notion that human self-trust highly impacts their confidence in generative AI. Therefore, with a generally low political knowledge, the majority of the population is highly susceptible to being influenced by the model.

Identifying the Gap:

Considerable research has explored algorithmic biases in Generative AI models, such as ChatGPT, revealing a left-leaning political bias and investigating its implications for model performance (Nadeem 2020, Noble 2019, Peters 2021). An additional research sector has delved into how humans engage with political information to form opinions while another sector looked into human psychology, investigating the relationship between self-confidence and reliance on generative AI models (Nadeem 2020, Noble 2019, Iyengar et al. 2019, Converse 1964, Chong et al. 2022). Despite this extensive exploration across sectors—ranging from the technical biases of AI to the psychological underpinnings of human-AI interaction—there remains a gap in synthesizing these areas to address the overarching question: "So what?". This question probes the consequences of political bias in generative AI on users, voters, and our democratic system.

Unlike rational base topics, interactions with political information are deeply intertwined with personal ideology and identity, rendering them complex and emotionally charged. The exploration of human interaction with Chess does not reveal enough to understand the mechanism by which humans adopt political recommendations from AI. Can low self-confidence overthrow strong emotional political identities? Would ChatGPT, with its vast knowledge and dependable mannerisms, shift users' votes? Would the model bias truly impact users any differently from the way the biases in news sources do? These questions are interconnected and are the building blocks for my research.

Through the experiment, I will seek to offer insights into the potential impacts stemming from political biases in generative AI models and the role self-confidence and political knowledge play, contributing to the discussions of regulation of such technologies. Understanding the core impacts of ChatGPT would provide the guidelines for proper and efficient regulations which could prevent potential ramifications of such biases on our political landscape. Accordingly, this research will serve as a foundational stepping stone for deeper examinations into the effects left leaning bias within a generative model has on users and consequently on the democratic system and governments.

Furthermore, it seeks to ascertain whether individuals across the political spectrum, whether liberals or conservatives, perceive and engage with political biases differently. Ultimately, the inquiry prompts us to ponder: Is the political bias embedded within generative AI models fundamentally distinctly impacting readers from the biases found in traditional news articles?

Argument:

Delving deeper into the intricate relationship between political bias in generative AI, my research unfolds across three pivotal dimensions. Firstly, I probe into how the ingrained biases within AI platforms like ChatGPT might sculpt public opinion and, by extension, sway future electoral outcomes. This exploration seeks to unearth the extent to which the embedded political biases could indoctrinate users' political views, potentially skewing voting patterns in unforeseen ways. Secondly, the investigation shifts focus toward the interplay between an individual's political knowledge, self-confidence, and susceptibility to AI biases. This segment aims to discern how these personal attributes either mitigate or amplify the influence of biased information. Lastly, the study contrasts the perceived reliability of generative AI against traditional and digital news sources, delving into demographic variances in trust, particularly among Democrats. Through these lenses, I aim to answer the question of “why should we care about the political bias in generative AI?”, grounding each hypothesis in the broader discourse on technology's intersection with politics.

As Generative AI becomes a primary source of political information, its ease of use may shape users' perspectives through embedded biases. Quick answers provided by AI may lead users to accept skewed viewpoints without questioning their accuracy or neutrality. The liberal-leaning political bias will become more indoctrinated within people's political perceptions and understanding which can have impacts on future voting trends. This will likely influence conservatives' opinions and votes more than it will liberals'.

Under the exploration of political knowledge, self-confidence, and reliance on ChatGPT, I hypothesize that individuals with low political knowledge are more likely to be impacted by the bias. The lack of political knowledge hints at an indifferent attitude towards politics. Such indifference makes the generative AI model even more appealing and useful. Those with low self-confidence will also be more influenced by the bias. They might have a good understanding of the topic, but will not trust themselves enough to rely on their response without adopting what ChatGPT told them. Conversely, individuals with higher confidence or greater political knowledge are less susceptible to the influence of ChatGPT's bias. They are more likely to critically evaluate the information provided and rely more on their knowledge when interpreting its responses.

Lastly, following humans' reliance pattern on AI models, I hypothesize that regardless of how reliable users view ChatGPT to be, they will continue to depend on its responses. Following automation bias, the overreliance on a model simply because it is automated, users are often inclined to perceive answers generated by artificial intelligence as accurate and neutral (Mosier 1988). This predisposition towards trusting AI responses may lead users to maintain a relatively normal level of trust in generative AI, regardless of whether they perceive ChatGPT as biased compared to other news sources. While individuals may still consider traditional sources like the New York Times more reliable than ChatGPT, they are likely to view ChatGPT as more dependable than social media platforms and not significantly inferior to other news sources.

Diving deeper into reliance patterns, Democrats are likely to perceive ChatGPT as more reliable compared to Republicans. Given the disproportionate representation of Republican conservative students at UCLA in contrast to the predominantly left-liberal Democratic

population, my focus will primarily center on Democrats and their susceptibility to political bias. Instead of directly comparing the reliance on ChatGPT between Democrats and Republicans, my study will assess how Democrats perceive the reliability of ChatGPT relative to Fox News. While this approach may not directly address the contrast in views between Democrats and Republicans regarding ChatGPT, it offers valuable insights into Democrats' awareness of, and reliance on, its left-leaning bias. Future research endeavors should aim to include a more balanced sample, encompassing both conservative and liberal Republicans, to gain a comprehensive understanding of the impact of political bias within generative AI models. Nonetheless, examining Democrats' views on ChatGPT will provide insights into whether they consciously or unconsciously recognize its left-leaning bias and rely on it more than on a known right-leaning news source.

Methods:

To test these three hypotheses I constructed two surveys. The two surveys are almost identical, with the source for the ChatGPT interaction section being the only factor that differentiates the two. Each survey was broken down into five main aspects: background questions, interaction with ChatGPT, political knowledge, reliance on news sources, and confidence.

Experimental Design: Interaction with ChatGPT

I began the experiment with introductory questions, such as participant gender, school grade, and major I presented two court cases. To ensure no type of priming occurred from the

previous question I conducted the ChatGPT interaction section at the beginning. During this section, participants received two court cases they had to rule on. Within each case, participants had to take on the role of a Supreme Court judge, read a short passage, and then rule on the case. Although both surveys covered the same cases, Version A of the survey had information provided by ChatGPT and Version B of the survey had information provided by the SCOTUS blog. I selected SCOTUSblog as the comparative information source for this study due to its neutrality and relative obscurity among the general population. This choice was crucial because widely recognized sources like the New York Times or Wikipedia come with pre-existing biases and opinions, which could influence participants' responses based on their political leanings. For example, Liberals might distrust information from Fox News, while conservatives could react skeptically to the New York Times. Additionally, people hold different perceptions regarding Wikipedia's reliability based on warnings given to them by teachers during middle or high school. By choosing SCOTUSblog, a site less known and without a widespread reputation for bias, I aimed to ensure that both Liberal and Conservative participants would approach the information with a neutral stance.

To maintain consistency across both versions, the two sources presented an identical list of pros and cons for the respective cases. The first case presented is from *Loper Bright v. Raimondo* and *Relentless v. Department of Commerce*, discussing the debate over whether to overrule *Chevron* or not. The case was simplified to focus on the decision to overrule *Chevron* for four main reasons. First, while the free version of ChatGPT has not been updated to include the latest details on the lawsuits, it is well-equipped to provide a comprehensive analysis of the *Chevron* doctrine, established in 1984. This enables ChatGPT to discuss the doctrine's

objectives, benefits, and drawbacks in depth. Second, to ensure the survey remained brief, complex and lengthy questions that covered a multilayer case were avoided. Third, in a real-world scenario, someone attempting to grasp the intricacies of cases like *Loper Bright v. Raimondo* and *Relentless v. Department of Commerce* would quickly realize that understanding the Chevron doctrine is crucial. Given the doctrine's foundational role in these cases, individuals would turn to ChatGPT to get a quick and direct explanation of the topic. Lastly, as the average voter is uninformed on the Chevron doctrine, they will have no presupposed perceptions or knowledge about the case. Therefore, when voting, their decision would solely depend on the information they were provided with. This ensures dependency on either ChatGPT or SCOTUSblog.

The second legal case used was *McCutcheon v. Federal Election Commission*, wherein participants were tasked with assessing whether the aggregate limits on campaign contributions constituted a violation of First Amendment rights. Unlike the Chevron case, under the second case, likely, participant votes will likely reflect their predisposed thought regarding the First Amendment or campaign contributions. As most participants are unaware of the specific details of the court case, they would still have to depend on the information they receive on the survey. Therefore, the second court case allows for the examination of ChatGPT's impact on users even regarding topics that they hold implicit emotional opinions on.

The support for both court cases was split between the two political parties. The Democrats were in support of not overruling Chevron and believed aggregate limitations on campaign elections did not violate the First Amendment, while Republicans were in favor of

avored overruling Chevron and argued that such limits on campaign spending breached First Amendment rights overruling Chevron and removing aggregate limits on the election campaign as they did violate the first amendment.

Experimental Design: Political knowledge

The next step is to measure political knowledge. This variable will help answer the question of whether individuals with low political knowledge interact differently than those with high political knowledge. Due to its complexity, and the interconnected play between rational and emotions when it comes to political understanding, there has been a lot of discussion revolving around the best way to measure political knowledge (Michael 1993, Rapeli 2022). Michael X. Delli Carpin and Scott Keeter's text views political knowledge as the aggregate of information that citizens take into consideration when making a decision or formulating an opinion concerning political or social patterns. Essentially, the greater the range of factual information that an individual possesses and stores in their long-term memory regarding politics, the more politically knowledgeable they are considered. When discussing how individuals come about having high political knowledge, they point to the three main factors; ability, motivation, and opportunity.

Because of the need to keep the survey short, I opted to test for individual political knowledge by focusing on motivation. I measured participants' political information intake, voting habits, and involvement in politics. Students that consume political news on the daily or weekly are not only constantly informed about political news, but present consistent interest in politics. Conversely, students who do not consume news often or elect not to vote are likely to

lack a strong understanding of current political news as they also present a lack of interest in politics. People who are likely to use ChatGPT to obtain political news during elections are those who do not consume political news often nor would they show consistent active participation in politics. Creating a distinction between those who are knowledgeable and those who do not help not only partially accounts for the voting pattern but also tracks the pattern of those who are most likely to rely on ChatGPT during real elections.

Experimental Design: Reliability

For my third hypothesis I am testing to see if users view ChatGPT as be reliable source for news, therefore during the experiment participants were presented with six questions. Each question asked participants to rate the reliability of different news sources on a scale from 1 (not reliable) to 10 (very reliable). I compared three popular news sources—two left-leaning: the New York Times and CNN, and one right-leaning: Fox News, along with three 'non-political' sources: Wikipedia, ChatGPT, and social media. By “non-political” I mean sources that were not created with the intention of informing users of political information. While social media has been used by many teens to stay posted on political events, when people think of political news sources they often immediately think of sources like the New York Times, Fox News, or MSNBC news.

I presented the news sources in order of left-leaning news, right-leaning news, and neutral with ChatGPT being placed in the middle between Wikipedia and social media to not stand out too much. Following the reliability questions, participants were asked to share which sources they often turn to when looking for political information. This explores whether some

have already started using ChatGPT to get political news as well as whether individuals rely on social media, or go to verified news sources such as CNN The New York Times, or Fox News for political information.

Experimental Design: Self-confidence

Lastly, to measure individuals self confidences in their political knowledge, I prompted participants with the following question: “ A friend asks you ‘why the Democratic party’s stance on immigration shifted to the right? ‘ How would you answer?”. This question touches on a subject that has been addressed often in the news but is not popularly discussed on social media platforms. Answering the question correctly requires staying updated with political news, hinting at high political engagement and knowledge. Conversely, an incorrect answer doesn't necessarily indicate a lack of political understanding, but rather, it might reflect a lower level of political engagement. To assess participants' self-confidence, I also asked them to rate their certainty in their answers, ranging from 'very sure' to 'unsure.' In my data analysis, I will examine the correlation between the accuracy of their answers (as an indicator of political knowledge) and their stated confidence level.

Data Collection:

When collecting data, I mainly surveyed UCLA students. To get students to participate in my survey I sent a link that randomized the two versions of the survey, with a short description, explaining how I am examining current political information intake, ensuring not to reveal too much, and accidentally prime the participants. As the UCLA student body mainly consists of

democratic and liberal students, my data mostly contains responses from politically left leaning students. The use of only college students limited my participant pool on average to individuals between the ages of 18-23. Because of these limitations, I would recommend conducting this experiment again with a bigger data set that has a higher variability of participants, of both older age and of a Conservative or Republican political orientation.

Data Analysis:

In delving into the data, various patterns emerged, some anticipated while others were unexpected. I broke down the data analysis by my argument structure, first diving into voting trends among participants and comparing the results from the two versions of the survey.

Although I hypothesized ChatGPT would cause an overall shift in votes to favor the Democratic side, an opposite trend occurred under the first court case, with the majority of participants voting to overrule Chevron. To further explore this unexpected trend I studied the voting trends between the opposing sources in relation to individual political affiliation, political knowledge, and self-confidence level. Lastly, I explored participants' perceptions of ChatGPT's reliability relative to other widely recognized news sources, to grasp a better understanding of how users view ChatGPT in comparison to other platforms.

Chatgpt vs SCOTUSblog:

My hypothesis anticipated that in the Chevron case, students exposed to information from ChatGPT would predominantly vote against overruling Chevron, aligning with the stance of the Democratic Party. However, Figure 1A illustrates a contrary trend, with the majority of votes

leaning towards overruling Chevron—a position typically favored by Republicans. Conversely, as seen in Figure 1B, participants who received information from SCOTUSblog showed a more evenly distributed voting pattern, albeit with a tendency towards supporting overruling Chevron. Nonetheless, this support was less pronounced compared to participants who received Survey Version A containing information from ChatGPT.

Figure 1B reveals that 42 percent of participants, who received the text from SCOTUSblog, voted not to overrule Chevron, while 57 percent voted in favor of overruling it, representing a 15 percent difference. In contrast, Figure 1A depicts a more significant disparity of 50 percent between those who voted to overrule Chevron and those who voted to not overrule Chevron after reading the text from ChatGPT. This discrepancy suggests differing reactions to information from ChatGPT compared to SCOTUSblog. This pattern was also observed in the voting trends for the case concerning Aggregate Limits on Election Campaigns.

Figures 2A and 2B illustrate these diverging trends, with a 52 percent difference in votes under ChatGPT and only a 30 percent contrast under SCOTUSblog. Compared to the first court case, under this case, participants aligned more with the Democratic party, ruling that such limits do not violate the First Amendment. The differing patterns observed—initially diverging from, then returning to, popular democratic stances—are influenced by several possible factors. One key factor is the disparity between the Chevron case and the case concerning Aggregate Limits on Election Campaigns. While both cases are unknown to the average voter, most voters are more likely to have a pre-existing opinion regarding the second case than the first.

The Aggregate Limits on Election Campaigns contain key terms like 'First Amendment' and 'election campaigns' that are linked to various aspects of politics, leading many individuals,

especially college students, to quickly form connections with these topics. This could lead participants to either not depend as heavily on the provided material or to incorporate their perspective into the text when reading before finalizing their final vote. When presented with information on the Chevron case, students rely solely on the provided information, as the average person is not familiar with federal agencies. Additionally, as the Chevron case was presented before the aggregate limit case, there is a possibility of order bias playing a role in the shift. In any case, these findings suggest that while ChatGPT's political bias may not have the initially predicted impact, it may exert a different type of the influence. Rather than the bias itself impacting participants, it appears the formatting and the perception of the information sources itself impact the voting trend.

Figure 1: Number of votes on Chevron based on information source

Figure 1A

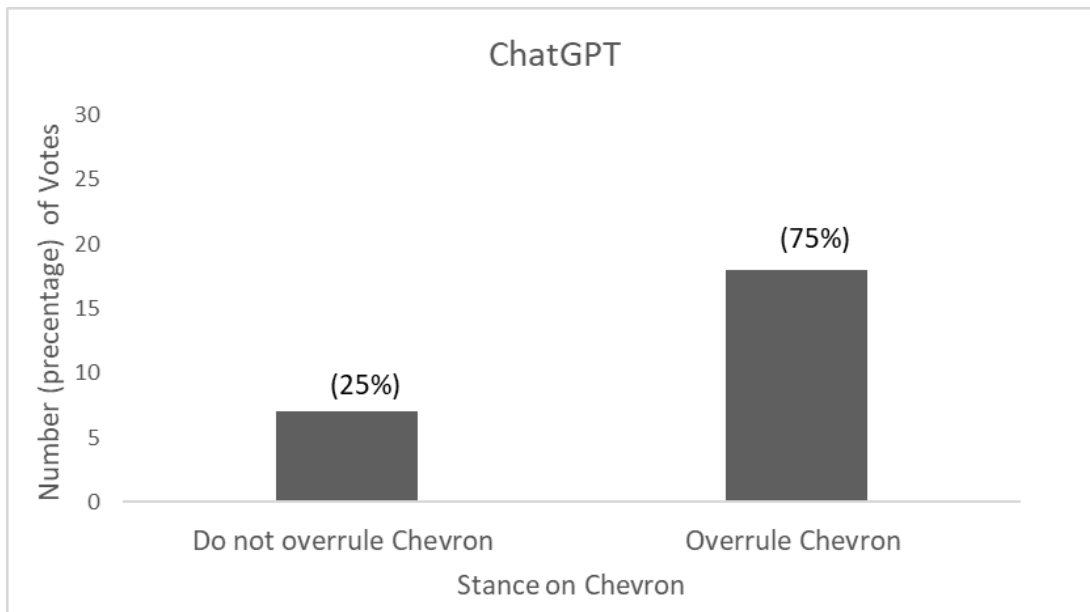


Figure 1B

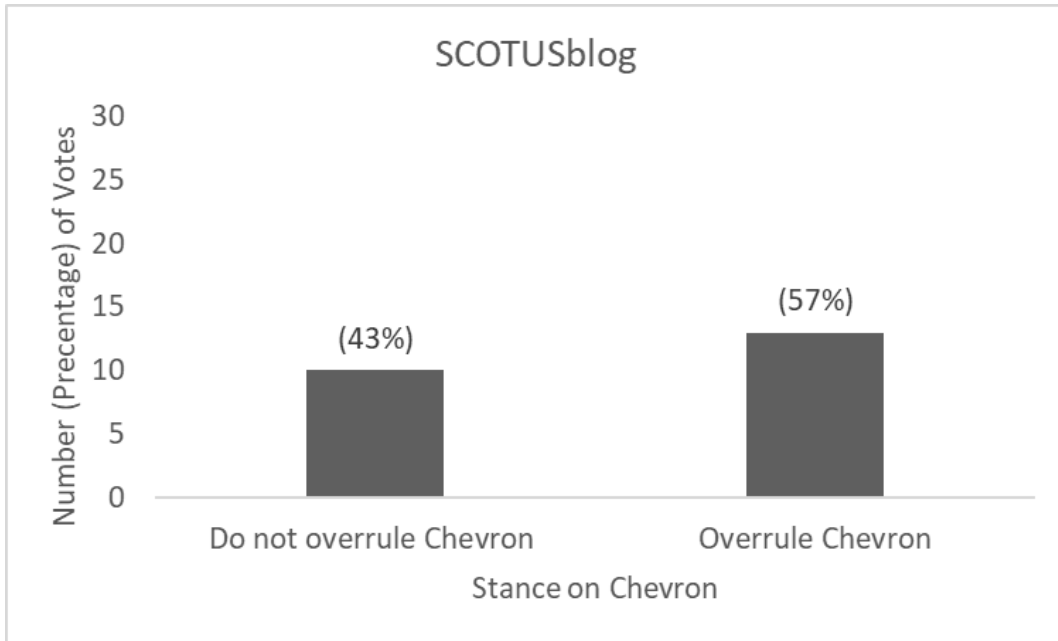


Figure 2: Votes on Aggregate Limits on Election Campaigns based on information source

Figure 2A

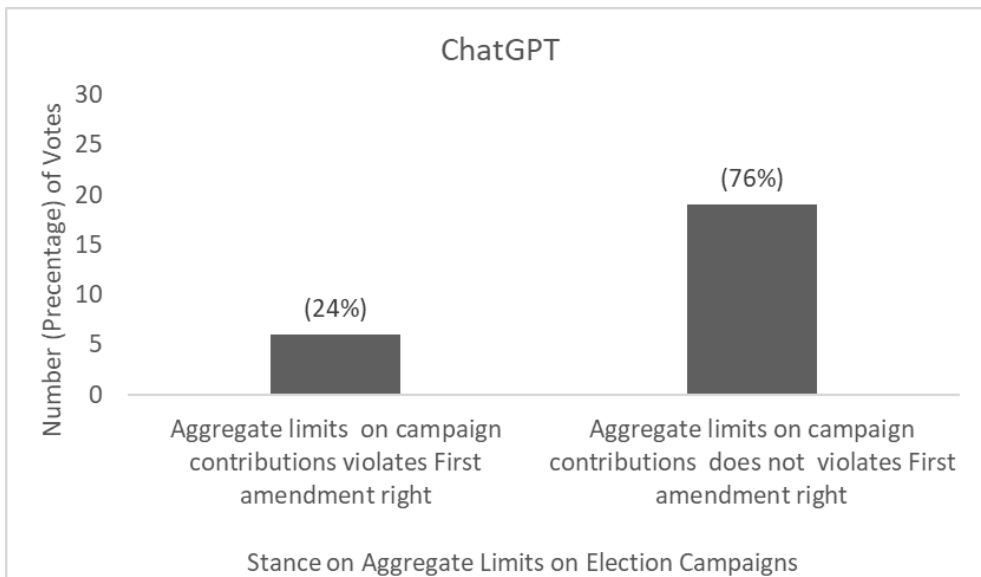
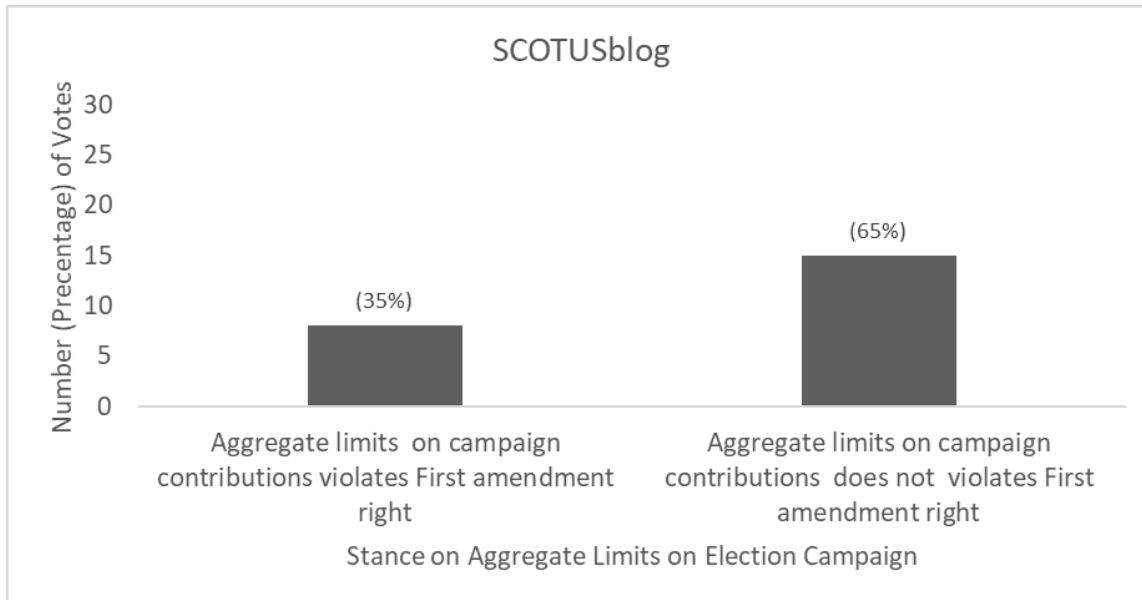


Figure 2B



While the voting trends under ChatGPT exhibit significant differences, as seen in Figures 1A and 2A, the votes under SCOTUSblog, as depicted in both Figures 1B and 2B, are nearly evenly split. This discrepancy could stem from participants' perceptions of the two sources. SCOTUSblog is considered a neutral and scholarly source, distinct from mainstream news outlets like the New York Times, which may have inherent biases. Its credibility and objectivity make it a reliable resource for research and supporting arguments. In contrast, ChatGPT is seen as a tool primarily used by students for quick information. When participants realize they received information from ChatGPT, they anticipate receiving answers presented in a simple and easy-to-understand manner, leading them to invest more energy in reading it (assuming it won't require much effort initially). Conversely, when reading information from SCOTUSblog, participants may feel daunted by the prospect of engaging in "academic" work while completing

the survey, resulting in less attention paid to the content and possibly leading to a more randomized vote. This could explain the appearance of equally distributed votes under SCOTUSblog. To further understand the influences the two sources may have on participants I dived deeper to examine the pattern of votes based on political orientation, knowledge, and confidence levels.

Political orientation:

As the majority of participants identify as Democrats, I explored whether their political affiliation, rather than party orientation, influenced their voting behavior in both the Chevron and Aggregate Limits on Election Campaign cases. Figures 3 and 4 illustrate the voting patterns based on participants' political affiliations. As seen in Figures 3A and 3B, among liberals, a significant divergence in voting behavior was observed between exposure to SCOTUSblog and ChatGPT generated information. 60 percent of liberals who read SCOTUSblog aligned with the Democratic Party's position and voted not to overrule Chevron. However, when exposed to information from ChatGPT, 83 percent of liberals voted to overrule Chevron, indicating a notable shift in their stance.

Figure 3: Votes on Chevron based on political orientation

Figure 3A

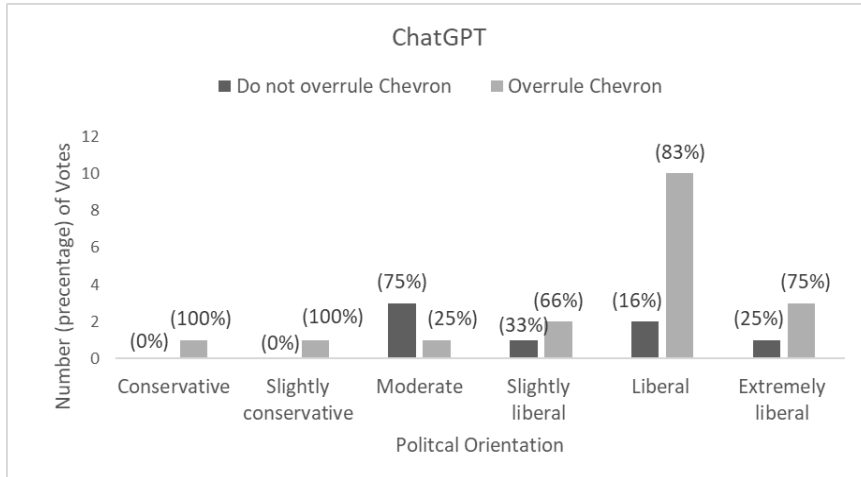
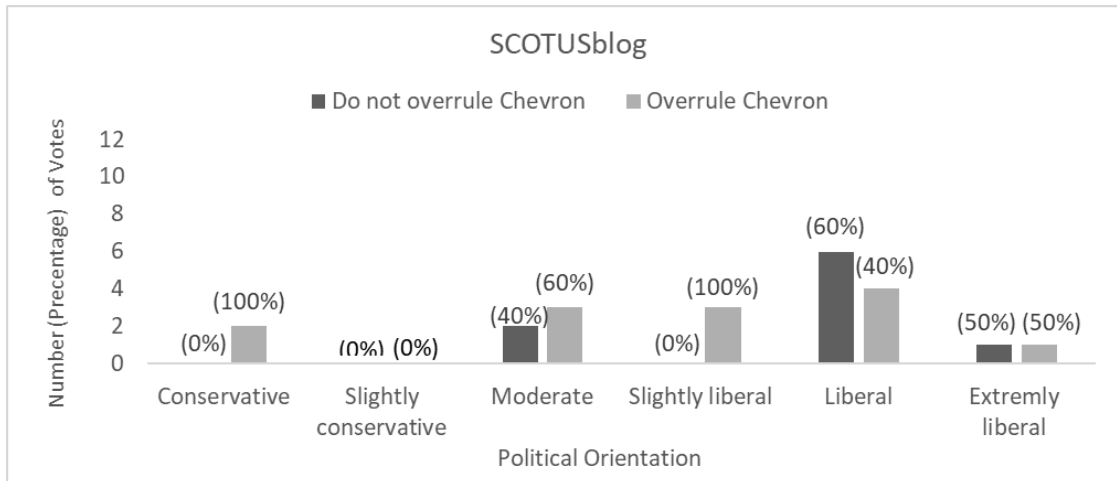


Figure 3B



Upon analyzing the voting patterns regarding the Aggregate Limits on Election campaigns (Figures 4A and 4B), liberals reverted to their party's stance, supporting the view that the aggregate limit on campaign elections does not infringe upon the First Amendment. However, conservatives remained steadfast in their position. In both the Chevron and Aggregate Limits on Election Campaign cases, conservative votes mirrored the stances of the Republican

Party. Since the dataset's conservative representation is limited, it doesn't outright refute my initial hypothesis suggesting that conservatives might be more susceptible to the model's influence than Democrats. However, it does hint at the possibility of inconsistent influence that the model could have on Democrats in comparison to Republicans.

Given the model's left-leaning bias, this suggests that while it was anticipated to induce a leftward shift, it might instead facilitate a shift by aiding voters in attaining a "neutral" comprehension of the issue and voting accordingly. By "neutral," I mean that the model, despite its biases, presents information in a manner that enables users to set aside their emotions and focus solely on the information presented.

Figure 4: Votes on Aggregate Limits on Election Campaigns based on political orientation

Figure 4A

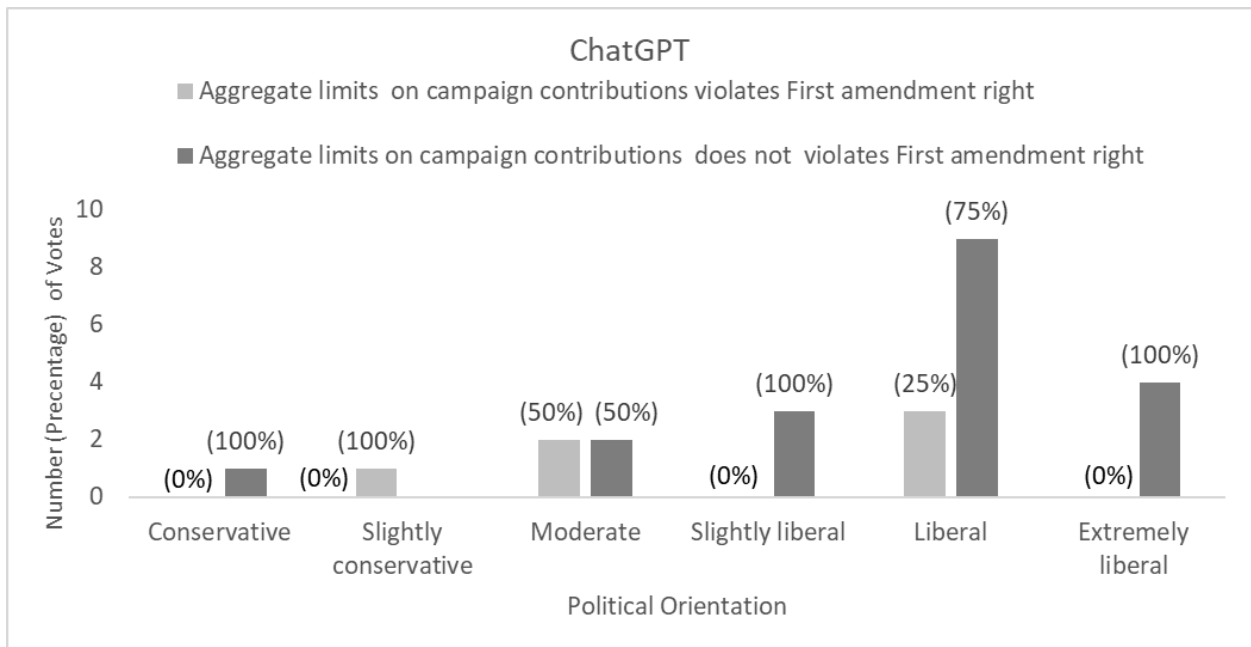



Figure 4B



Political knowledge

To investigate my second hypothesis that individuals with high political knowledge are less susceptible to political bias, I analyzed the voting patterns in both Version A and Version B of the survey for the two cases. A discernible pattern emerges, suggesting that individuals with higher political knowledge tend to align more closely with Democratic Party stances than those with lower knowledge levels.

In Figure 5A, it is evident that 60 percent of individuals who regularly consume political knowledge voted to overrule Chevron. Despite the overall trend favoring the overruling of Chevron, it is noteworthy that across all levels of political exposure, individuals who consume political news daily exhibited the highest percentages of votes favoring to not overrule Chevron.

This aligns with my hypothesis that individuals with greater political knowledge are less swayed by external influences, such as ChatGPT. While it appears my initial hypothesis—that political bias significantly affects voting behavior—might not have been confirmed in the expected direction, the data still reveal a noticeable shift in voting patterns. Therefore for the rest of the data analysis, I would conclude that if a liberal voted to overrule Chevron they were influenced by ChatGPT, hence a vote to not overrule Chevron indicates resistance to the influences of ChatGPT.

The influence of ChatGPT is even more apparent, as Figure 5B presents a different pattern than that in Figure 5A. After receiving information from SCOTUSblog, those who consume the least amount of political news voted at a higher percentage to not overrule Chevron. As for the Aggregate Limits on Election Campaign cases, under both survey versions, most voters aligned with the Democratic party stance.

Figure 5: Number of votes on Chevron based on political knowledge

Figure 5A


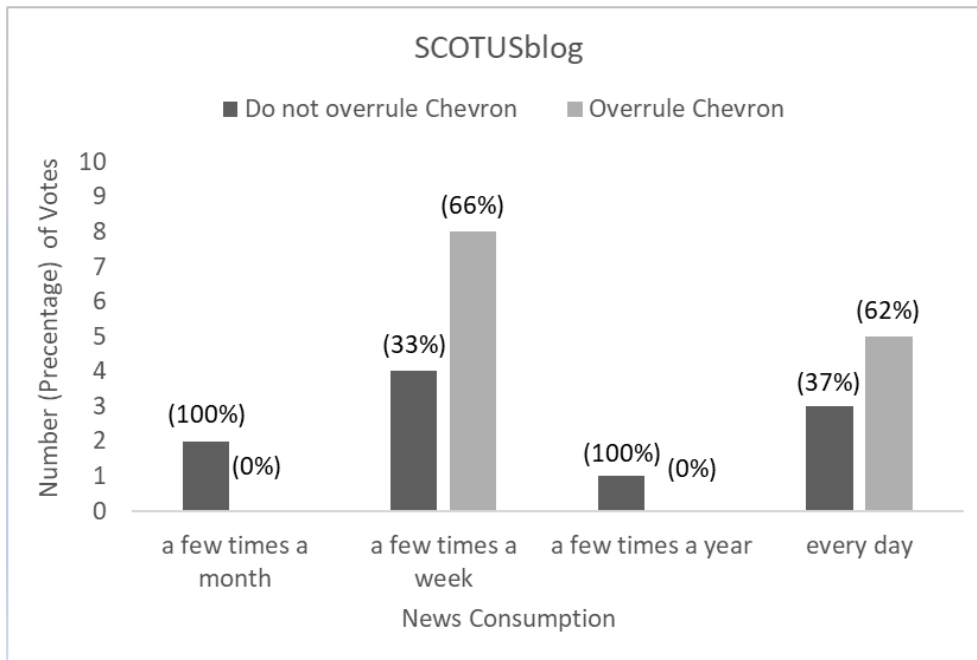


Figure 5B



In Figure 6A, 80 percent of those who consume political news every day voted in alignment with the Democratic party, while only 25 percent of those who consume political news a few times a year voted in that direction. There is a consistent pattern shown in Figure 6B, with 63 percent who consume political news every day voting along Democratic party lines and 100 percent of those who only consume political news a few times a year voting against the Democratic party lines.

In these figures, a clear trend emerges again where individuals with high political knowledge predominantly voted in line with Democratic Party positions, particularly regarding the Aggregate Limits on Election Campaigns, which they viewed as not violating the First Amendment. This alignment with Democratic Party lines suggests that individuals with high

political knowledge are less influenced by political bias, as they base their decisions on a deeper understanding of the issues rather than succumbing to partisan sway.

It is important to note that while political affiliation is not explicitly included in the chart, as the two datasets only contain a total of 3 Republicans, it is reasonable to infer that most trends are driven by the votes of Democrats or Independents. While Independents may not always align with Democratic Party lines, the majority of the data consists of individuals who identify with liberal ideology, and in both instances, the Democratic Party's stances reflect liberal positions.

Figure 6: Number of votes on Aggregate Limits on Election Campaigns based on political knowledge Figure 6A


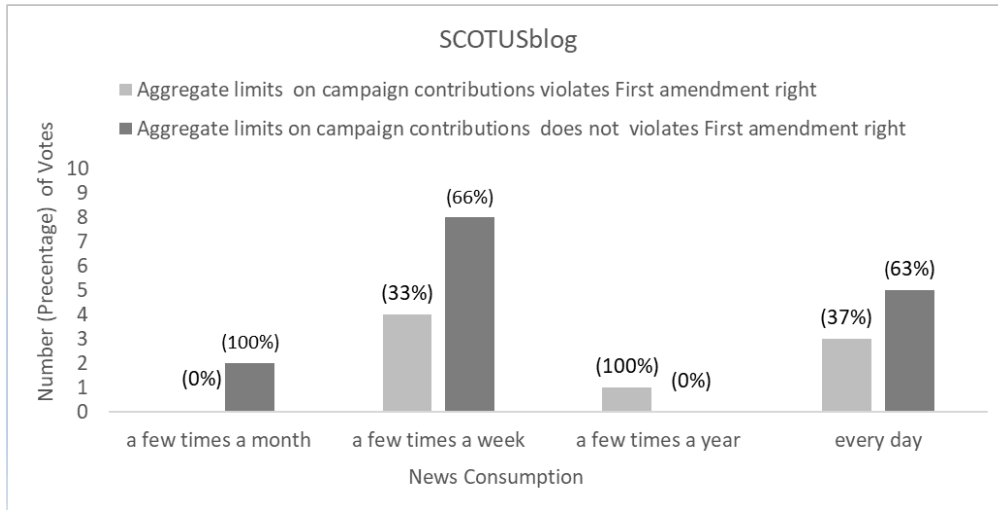


Figure 6B



Self-confidence:

To explore the relationship between confidence and ChatGPT, I broke down the analysis into two different parts. The first one analyzes the relation between self-confidence and the accuracy of the responses, to show how in the dataset I collected, high political knowledge correlates with high confidence. The second part compares individuals' confidence ranking with their votes on both Chevron and SCOTUSblog. It should be noted that the experiment is not set up to explore causation, and therefore it cannot be concluded that high knowledge causes high confidence or vice versa. Instead, I am simply showing that those with high confidence also tend to have high political knowledge, and therefore their interaction with ChatGPT might be different than those with low knowledge and low self-confidence.

Figure 11 depicts the distribution of students' responses to a question posed at the end of the survey, which inquired, "Why has the Democratic Party's stance on immigration shifted to the right?" Responses were categorized as correct, incorrect, and don't know (indicative of participant confusion or lack of knowledge), with consideration given to the respondents' level of confidence in their answers. The data presented in Figure 11 are aggregated from both Survey A and Survey B. The analysis reveals that a higher proportion of students who provided incorrect answers expressed either a somewhat sure or unsure level of confidence. Notably, none of the students who reported being very sure of their answers answered correctly; however, there were instances where students who expressed they were unsure provided correct responses, indicating a pattern of low self-confidence but high knowledge. There appears to be a correlation between providing a correct answer and expressing confidence in one's response. This observation underscores the importance of considering both the accuracy of responses and participants' confidence levels in understanding their decision-making process. Now that we know that the pattern of high knowledge and high confidence exists we can move to explore the relation between voting patterns and confidence levels.

Figure 11: Accuracy vs. Confidence: Responses from ChatGPT & SCOTUSblog




Figure 13 explores the relationship between the confidence measure and votes on Chevron. Figure 12A reveals that the highest support for not overturning Chevron came from groups at opposite ends of the confidence spectrum: "Sure" (high confidence) and "Unsure" (low confidence). This pattern does not carry over to the second case as seen in Figure 13A. Under the Aggregate limit of election campaigns, in all three categories of "somewhat sure", "sure", and "very sure", the majority of voters voted in line with the democratic party votes. The low amount of data could account for the even split of votes under the unsure confidence level in Figure 12A. When comparing Figure 12A and 12B, it is noticeable that in Figure 12B a larger percentage of voters who had high a middle confidence level (somewhat sure) and high confidence level (sure) voted in favor of not overruling Chevron. That is, under non-ChatGPT sources, the majority voted in line with democratic party views. Under ChatGPT, while the overall votes were impacted by the model and shifted to the right, individuals with high confidence presented the highest level of "resistance" to said influences. Therefore, it could be concluded, that when comes to receiving political information from ChatGPT, high confidence does result in lower influences and reliance on the model than does low conferences.

Figure 12:Chevron votes based on confidence measure
 Figure 12A:ChatGPT




Figure 12B: SCOTUSblog

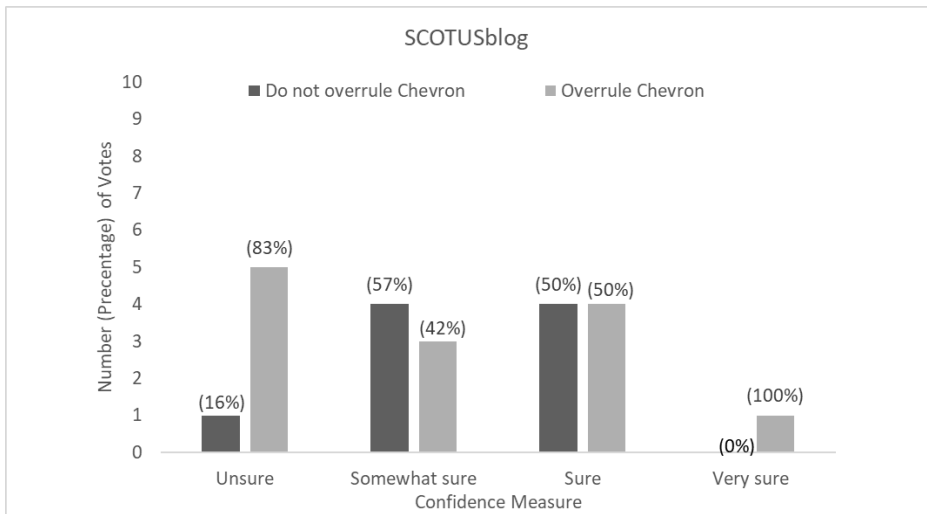


Figure 13: Aggregate Limits on Election Campaigns Votes based on Confidence Measure

Figure 13A: ChatGPT


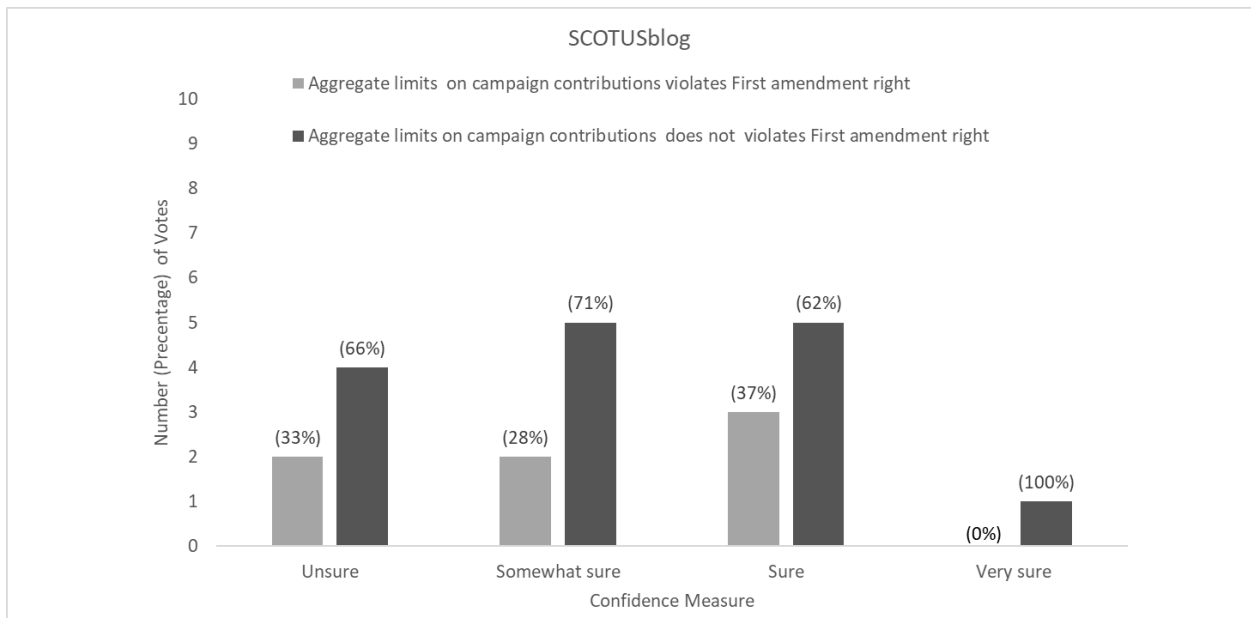


Figure 13B: SCOTUSblog



Reliability

To test my third hypothesis, I measured the average ratio of reliability for each news source. Each participant was asked to rank each news source from least reliable (1) to most reliable (10). To be able to compare the data, I took the average rating each news source received.

The lower reliance on social media, ChatGPT, and Fox News might stem from perceptions of these platforms as less trustworthy sources of information. Figure 7 illustrates this trend, with social media, ChatGPT, and Fox News receiving average reliability ratings of around 3, indicating a low level of trust. Conversely, CNN and The New York Times scored notably higher, with averages of 6 and 6, respectively, suggesting a moderate to high level of trust among respondents. Even Wikipedia, while falling in the middle with a reliability rating of approximately 4.39, garnered more trust than social media and AI-driven platforms.

This discernible pattern in perceived reliability highlights the preference for traditional news outlets over newer digital platforms. It underscores the importance of credibility and editorial integrity in shaping public trust in information sources, with established news organizations like CNN and The New York Times enjoying higher levels of confidence among readers. The higher reliance on popular news sources such as The New York Times could be attributed to the longstanding reputation and public image the news source has cultivated over time. Its established credibility and editorial stance may instill confidence in readers, particularly among Democrats who often seek out articles aligning with their predisposed opinions.

Figure 7: Comparison of Average Reliability Ratings for Different Sources, Survey Version A
 The average rating each news source received on from 1 to 10

	ChatGPT	SCOTUSblog
New York Times	6	6
CNN	5	6
Fox News	3	3
Wikipedia	5	4
ChatGPT	4	3
Social media	3	3

To further understand how reliability perceptions vary, I shifted my analysis to focus on how political affiliations influence the average reliability ratings for The New York Times, Fox News, and ChatGPT. The data, presented in Figure 8, reveal a marked contrast in viewpoints: conservatives, who took Survey A, deem The New York Times as unreliable, assigning it an average rating of 2, while extremely liberal individuals rate it highly reliable, with an average of 7. In contrast, Fox News receives an average rating of 5 from conservatives but is rated as unreliable (average rating of 2) by extremely liberal respondents. The average reliability rating for The New York Times among conservatives who participated in Survey B matches the rating given by extreme liberals across both survey versions, with an average rating of 8. This unusual similarity, likely influenced by the dataset's limited number of conservatives, suggests an outlier rather than a trend-disrupting factor. Given this anomaly arises as the sole dramatic difference (exceeding a 2-3 point spread), it does not challenge the established perceptions of source

reliability. This disparity in ratings aligns with expectations but underscores the significant differences in perception between conservatives and extreme liberals.

Figure 8: Average reliability rating on New York Time by each political affiliation group

	ChatGPT	SCOTUSblog
Conservative	2	8
Slightly Conservative	5	NA
Moderate	5	6
Slightly Liberal	5	6
Liberal	7	7
Extremely Liberal	7	8

Figure 9: Average reliability rating on Fox News by each political affiliation group

	ChatGPT	SCOTUSblog
Conservative	5	4
Slightly Conservative	1	NA
Moderate	3	5
Slightly Liberal	3	2
Liberal	2	3
Extremely Liberal	2	2

When asked to rate the reliability of ChatGPT, I hypothesized that liberals would view it as more reliable than conservatives. Figure 10 presents the average reliability ratings, revealing that conservatives rated ChatGPT with a mean of 6 and 4, whereas extremely liberal respondents rated it at a mere 3. This discrepancy may be attributed to the limited number of Conservative participants, as slightly Conservative respondents' perception of ChatGPT averaged a reliability rank of 1. Conversely, slightly liberal and liberal respondents averaged a rating of 4, which is 1 to 2 points higher than their average reliability rating for Fox News and only 1 to 3 points lower than their rating for The New York Times. This suggests that while conservatives may wholly disregard The New York Times, they are somewhat willing to consider ChatGPT as reliable.

Amidst the current climate of intense political polarization and accusations of false news, the emergence of a source like ChatGPT, which garners similar views from both parties, could potentially serve as a unifying factor in our democracy.

Figure 10: Average reliability rating on ChatGPT by each political affiliation group

	ChatGPT	SCOTUSblog
Conservative	6	4
Slightly Conservative	1	NA
Moderate	4	4
Slightly Liberal	4	1
Liberal	4	3
Extremely Liberal	3	3

Conclusion:

This study initially hypothesized that the left-leaning bias of ChatGPT would influence voting trends. However, my findings suggest that ChatGPT's political bias might be harmless. The experiment uncovered that ChatGPT does influence how users vote more than the secondary source, SCOTUSblog. However, rather than the bias influencing users, it is the user's perception of the AI model and its user-friendly style that shifts users' perception of political information. My study uncovered that the influence ChatGPT has on users is heavily related to the user's

self-confidence and political knowledge. When a user has high self-confidence and political knowledge their reliance on ChatGPT drops, while when they have low self-confidence and political knowledge they are often more reliant and influenced by the generative AI model. Further, while it was expected for the model to shift votes to the left in support of democratic party stances, liberals were generally seen to be influenced by the model, as their votes started to shift more to the right under Chevron. With ChatGPT's left-leaning bias (Rozado2023), this shift toward the right political spectrum is likely explained by the "neutral" appearances of the model, allowing users to separate their predisposed beliefs and emotions when analyzing political information. As this counters the perception that the hidden political bias is harmful it opens the door to a whole new area of study, exploring how the format of generative AI models might impact users' political perceptions. Additional research should be done with a larger pool of conservatives to explore the impact of political bias that my study was unable to cover. Further research could also be conducted to explore whether this reliability belief is consistent under a bigger pool of conservatives. If the belief is consistent, then is there potential that generative AI will be a neutral source both political parties depend on? As generative AI models evolve, society becomes increasingly vulnerable to their influence. Therefore, it is crucial to remain vigilant about their potential political impact, but at least for now we know, we do not need to be wary of the left-leaning bias that's embedded in our Generative AI

Bibliography

- “Algorithmic Bias Detection and Mitigation: Best Practices and Policies to Reduce Consumer Harms.” Brookings, <https://www.brookings.edu/articles/algorithmic-bias-detection-and-mitigation-best-practices-and-policies-to-reduce-consumer-harms/>.
- Almond, Gabriel A., and Sidney Verba. *The Civic Culture: Political Attitudes and Democracy in Five Nations*. New ed., Sage Publications, 1989.
- ChatGPT. <https://chat.openai.com>. Accessed 22 Mar. 2024.
- Chong, Leah, et al. “Collaborative Design Decision-Making With Artificial Intelligence: Exploring the Evolution and Impact of Human Confidence in AI and in Themselves.” Volume 6: 34th International Conference on Design Theory and Methodology (DTM), American Society of Mechanical Engineers, 2022, p. V006T06A021. DOI.org (Crossref), <https://doi.org/10.1115/DETC2022-88574>.
- Chong, Leah, et al. “The Evolution and Impact of Human Confidence in Artificial Intelligence and in Themselves on AI-Assisted Decision-Making in Design.” *Journal of Mechanical Design*, vol. 145, no. 3, Mar. 2023, p. 031401. DOI.org (Crossref), <https://doi.org/10.1115/1.4055123>.
- Confirmation Bias | Definition, Examples, Psychology, & Facts | Britannica*. 14 Feb. 2024, <https://www.britannica.com/science/confirmation-bias>.
- Danks, David, and Alex John London. *Algorithmic Bias in Autonomous Systems*. 2017, pp. 4691–97. www.ijcai.org, <https://www.ijcai.org/proceedings/2017/654>.
- Dzindolet, Mary T., et al. “The Role of Trust in Automation Reliance.” *International Journal of*

- Human-Computer Studies*, vol. 58, no. 6, June 2003, pp. 697–718. *ScienceDirect*,
[https://doi.org/10.1016/S1071-5819\(03\)00038-7](https://doi.org/10.1016/S1071-5819(03)00038-7).
- Fazelpour, Sina, and David Danks. “Algorithmic Bias: Senses, Sources, Solutions.” *Philosophy Compass*, vol. 16, no. 8, Aug. 2021, p. e12760. *DOI.org (Crossref)*,
<https://doi.org/10.1111/phc3.12760>.
- Johnson, Gabrielle M. “Algorithmic Bias: On the Implicit Biases of Social Technology.”
Synthese, vol. 198, no. 10, Oct. 2021, pp. 9941–61. *DOI.org (Crossref)*,
<https://doi.org/10.1007/s11229-020-02696-y>.
- Lodge, Milton, and Charles S. Taber. *The Rationalizing Voter*. Cambridge University Press, 2013. *Cambridge University Press*, <https://doi.org/10.1017/CBO9781139032490>.
- Mattu, Jeff Larson, Julia Angwin, Lauren Kirchner, Surya. “How We Analyzed the COMPAS Recidivism Algorithm.” *ProPublica*,
<https://www.propublica.org/article/how-we-analyzed-the-compas-recidivism-algorithm>
- Mesa, Natalia. “Can the Criminal Justice System’s Artificial Intelligence Ever Be Truly Fair?”
Massive Science, 13 May 2021,
<https://massivesci.com/articles/machine-learning-compas-racism-policing-fairness/>.
- Mosier, Kathleen L., et al. “Automation Bias: Decision Making and Performance in High-Tech Cockpits.” *Decision Making in Aviation*, Routledge, 2015.
- Suguri Motoki, Fabio Yoshio, et al. “More Human than Human: Measuring ChatGPT Political Bias.” *SSRN Electronic Journal*, 2023. *DOI.org (Crossref)*,
<https://doi.org/10.2139/ssrn.4372349>.
- Nadeem, Ayesha, et al. “Gender Bias in AI: A Review of Contributing Factors and Mitigating Strategies.” *ACIS 2020 Proceedings*, Jan. 2020, <https://aisel.aisnet.org/acis2020/27>

- Noble, Safiya Umoja. *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York university press, 2018.
- Peters, Uwe. “Algorithmic Political Bias in Artificial Intelligence Systems.” *Philosophy & Technology*, vol. 35, no. 2, June 2022, p. 25. DOI.org (Crossref), <https://doi.org/10.1007/s13347-022-00512-8>.
- Converse, Philip E. “The Nature of Belief Systems in Mass Publics (1964).” *Critical Review*, vol. 18, no. 1–3, Jan. 2006, pp. 1–74. DOI.org (Crossref), <https://doi.org/10.1080/08913810608443650>.
- Rapeli, Lauri. “Does Sophistication Affect Electoral Outcomes?” *Government and Opposition*, vol. 53, no. 2, Apr. 2018, pp. 181–204. Cambridge University Press, <https://doi.org/10.1017/gov.2016.23>.
- Rozado, David. “The Political Biases of ChatGPT.” *Social Sciences*, vol. 12, no. 3, Mar. 2023, p. 148. www.mdpi.com, <https://doi.org/10.3390/socsci12030148>.
- Rudin, Cynthia, et al. “The Age of Secrecy and Unfairness in Recidivism Prediction.” *Harvard Data Science Review*, vol. 2, no. 1, Jan. 2020. DOI.org (Crossref), <https://doi.org/10.1162/99608f92.6ed64b30>
- Sison, Alejo Jose G., et al. “ChatGPT: More than a ‘Weapon of Mass Deception’ Ethical Challenges and Responses from the Human-Centered Artificial Intelligence (HCAI) Perspective.” *SSRN Electronic Journal*, 2023. DOI.org (Crossref), <https://doi.org/10.2139/ssrn.4423874>.
- Thomas, Ryan J., and T. J. Thomson. “Ageism, Sexism, Classism and More: 7 Examples of Bias in AI-Generated Images.” *The Conversation*, 10 July 2023,

<http://theconversation.com/ageism-sexism-classism-and-more-7-examples-of-bias-in-ai-generated-images-208748>.

“What Is Generative AI? Everything You Need to Know.” *Enterprise AI*,

<https://www.techtarget.com/searchenterpriseai/definition/generative-AI>

Xu, Ruiyun, et al. *ChatGPT vs. Google: A Comparative Study of Search Performance and User Experience*. 2023. DOI.org (Datacite), <https://doi.org/10.48550/ARXIV.2307.01135>.

Appendix:

Thesis Survey

Hello my name is Neomi Goodman, and I am a fourth year student studying political science.

This is a 21 question survey about political information consumption.

If you would like to see how your answers contributed to the experiment, email me at neomigm8@ucla.edu to receive a copy of the results.

What is your school year? *

Freshman

Sophomore

Junior

Senior

What is your major? *

What is your gender? *

Suvery Version A

You are a Supreme Court judge. You will preside over the decision to overrule the Chevron case.

Read the information from ChatGPT and make your ruling. *

Chevron deference refers to the legal principle established by the U.S. Supreme Court in *Chevron U.S.A., Inc. v. Natural Resources Defense Council, Inc.* (1984). This principle dictates that courts should defer to reasonable interpretations of statutes made by administrative agencies when the statute is ambiguous and the agency's interpretation is reasonable.

Arguments Against Overruling Chevron:

Expertise of Agencies: Administrative agencies often possess specialized knowledge and expertise in their respective fields, making them well-suited to interpret complex regulatory statutes and implement effective policies.

Arguments for Overruling Chevron:

Policy Concerns: Overruling Chevron could lead to clearer statutory interpretation and reduce the influence of administrative agencies in shaping public policy, allowing for greater accountability and transparency in government decision-making.

Overrule Chevron

Do not overrule Chevron

You are a supreme court judge. You will preside on whether aggregate limits on campaign contributions violated First Amendment rights.

Read the information from ChatGPT and make your ruling. *

Shaun McCutcheon sued the Federal Election Commission (FEC) because he argued that these limits restricted his ability to support multiple candidates and political committees, thus infringing upon his political expression and participation. He believed that the aggregate limits on campaign contributions violated his First Amendment right to freedom of speech and association.

One reason in support of the aggregate limits was the concern about preventing corruption or the appearance of corruption in the political process. Supporters argued that allowing unlimited contributions could lead to undue influence of wealthy donors over elected officials.

On the other hand, opponents of the aggregate limits argued that they infringed upon individuals' First Amendment rights to freedom of speech and association. They contended that such limits unnecessarily restricted individuals' ability to support multiple candidates and political causes, thus impeding their political expression.

Aggregate limits on campaign contributions violates First amendment right

Aggregate limits on campaign contributions does not violates First amendment right

Suvery Version B

You are a Supreme Court judge. You will preside over the decision to overrule the Chevron case.

Read the information from the SCOUTS blog and make your ruling. *

"The Supreme Court's 1984 opinion in *Chevron v. Natural Resources Defense Council*, upholding a regulation issued by the Environmental Protection Agency. The Court set out a two-part test for courts to review an agency's interpretation of a statute it administers. The court must first determine whether Congress has directly addressed the question at the center of the case. If it has not, the court must uphold the agency's interpretation of the statute as long as it is reasonable.

One judge suggested that federal agencies, with their scientific and technical expertise, are better suited than courts to resolve ambiguities in a federal statute.

They doubted whether there can be a "best" interpretation of a law when the justices "routinely disagree" about a law's meaning. If the court needs a "tie-breaker," why shouldn't it defer to the agency, with its expertise?

Another judge argued that *Chevron* "ushers in shocks to the system every four or eight years when a new administration comes in" and implements "massive change" in areas like securities law, communications law, and environmental law. "

Overrule Chevron

Do not overrule Chevron

You are a supreme court judge. You will preside on whether aggregate limits on campaign contributions violated First Amendment rights.

Read the information from the SCOUTS blog and make your ruling. *

Shaun McCutcheon sued because the two-year limits stopped him from giving anything to added candidates or political panels once his donations had reached the ceilings. Curbing the number of candidates he could support, he argued, infringed on his right to engage in political expression under the First Amendment .

The Justices inclined to vote for limits on campaign donations seemed to believe that the current system is corrupting in favor of the rich, but that they still would like some harder information on just how that happens.

And the Justices who usually vote for a freer flow of money into federal campaigns seemed to think there are enough safeguards against corruption already and that any more will stifle political expression, of the not so rich, too.

Aggregate limits on campaign contributions violates First amendment right

Aggregate limits on campaign contributions does not violates First amendment right

In Federal and State elections... *

I always vote

I almost always vote

I sometimes vote

I almost never vote

I never vote

I am involved in politics... *

a great deal (e.g. I hold a position in a politcial party)

a moderate amount (e.g. I attend political rallies)

a little (e.g. I vote)

none at all (e.g. I do not vote)

From 1 to 10 how reliable do you think the New York Times is for political information *

Not Reliable

Very Reliable

1	2	3	4	5	6	7	8	9	10
---	---	---	---	---	---	---	---	---	----

From 1 to 10 how reliable do you think CNN is for political information *

Not Reliable

Very Reliable

1	2	3	4	5	6	7	8	9	10
---	---	---	---	---	---	---	---	---	----

From 1 to 10 how reliable do you think Fox News is for political information *

Not Reliable

Very Reliable

1	2	3	4	5	6	7	8	9	10
---	---	---	---	---	---	---	---	---	----

From 1 to 10 how reliable do you think Wikipida is for political information *

Not Reliable

Very Reliable

1	2	3	4	5	6	7	8	9	10
---	---	---	---	---	---	---	---	---	----

From 1 to 10 how reliable do you think Wikipida is for political information *

Not Reliable

Very Reliable

1	2	3	4	5	6	7	8	9	10
---	---	---	---	---	---	---	---	---	----

From 1 to 10 how reliable do you think ChatGPT is for political information *

Not Reliable

Very Reliable

1	2	3	4	5	6	7	8	9	10
---	---	---	---	---	---	---	---	---	----

From 1 to 10 how reliable do you think Social Media is for political information *

Not Reliable

Very Reliable

1	2	3	4	5	6	7	8	9	10
---	---	---	---	---	---	---	---	---	----

I consume political news... *

every day

a few times a week

a few times a month

a few times a year

never

I get information about politics and current events from: *

Social media

The New York Times

CNN

Fox News

Wikipedia

ChatGPT

from friends / family

Another news source

I do not consume any political news

A friend asks you "Why has the Democratic party's stance on immigration shifted to the right?" How would you answer? *

0 / 250

How sure are you in your answer to the previous question? *

- Very sure
- Sure
- Somewhat sure
- Somewhat unsure
- Unsure

What is your political affiliation? *

Democrat

Independent

Republican

What is your political orientation? *

Extremely liberal

Liberal

Slightly liberal

Moderate

Slightly conservative

Conservative

Extremely conservative

Will you use ChatGPT to get political information when voting? *

Yes

No

Please explain why? *

0 / 100

