AI's Role in Deliberative Discussion

Jakob Miller, *Taylor University, United States*
Katie Kelley, *Covenant College, United States*
Peter Staritz, *Taylor University, United States*

ABSTRACT The ability to discuss politics intelligently with others is a central activity in many undergraduate political science courses, often forming a key part of student assessments and contributing significantly to the development of civic competence. This study explores the use of Artificial Intelligence (AI) as a tool to facilitate deliberative discussions among students. Based on our research, we present evidence that student interactions with an AI for political discourse can yield benefits that are similar to human-to-human political deliberation—including an increased openness to sharing views, discovery of common ground, and incremental improvement in respect for differing political views. These findings suggest that AI may offer a promising, scalable, and accessible solution to enhance deliberative learning experiences while potentially overcoming barriers observed in traditional classroom settings.

Deliberative democratic practices are increasingly vital in contemporary political science education. Engaging in civil deliberative discourse—in which individuals weigh reasons and respectfully discuss political issues and differences with one another[1]—is not only a core objective in many undergraduate political science courses but also is crucial for building civic competence in the general population. Decades of research show that structured political deliberation can provide multifaceted advantages for both citizens generally and undergraduate learners specifically (Andersson 2015; Delli Carpini, Cook, and Jacobs 2004; Minozzi et al. 2023).[2] For a more complete review of interpersonal political interactions, both outcomes, study, and practice, see Settle (2025).

Political deliberation deepens students' factual and conceptual understanding of policy issues. When students engage in structured discourse that challenges them to consider various perspectives and sources, they are encouraged to question assumptions, examine competing arguments, and connect abstract concepts to real-world implications. This exposure to a broader spectrum of viewpoints enhances their ability to synthesize complex ideas, and it fosters a more nuanced grasp of contentious topics (Esterling, Neblo, and Lazer 2011).

Deliberative activities also sharpen students' critical-thinking and argumentation abilities. In these settings, they must formulate coherent positions, anticipate counterarguments, and revise their claims based on new input. This iterative process mimics scholarly inquiry, thereby training students not only to defend their views but also to approach political dialogue with an openness to modification and self-reflection (Minozzi et al. 2023). As a result, students become more adept at analyzing rhetoric and reasoning—skills that are transferable to other academic and civic endeavors.

Discussion also can foster empathy and perspective taking. When students are exposed to political positions that are different from their own in a respectful and reasoned environment, they are more likely to recognize the legitimacy of opposing viewpoints. This experience potentially can reduce affective polarization and cultivate a broader sense of tolerance and interpersonal understanding (Caluwaerts et al. 2023; Rossiter and Carlson 2024).

Another key benefit of deliberative discussion is the enhancement of civic efficacy and the likelihood of future political participation. Students who engage in meaningful discussions about public issues often report feeling more confident in their political knowledge and more empowered to take action. Deliberation helps them to view themselves as agents in the democratic process, connecting classroom activities with larger societal outcomes (Delli Carpini, Cook, and Jacobs 2004).[3] This empowerment may translate into higher rates of voting, community involvement, and sustained public engagement.

Moreover, regular participation in deliberative dialogue cultivates core democratic norms including mutual respect, an appreciation for compromise, and trust in institutions that are tasked

Jakob Miller *is associate professor of political science at Taylor University. He can be reached at jakob_miller@taylor.edu.*
Katie Kelley *is assistant professor of psychology at Covenant College. She can be reached at katie.kelley@covenant.edu.*
Peter Staritz *is associate professor of physics and engineering at Taylor University. He can be reached at peter_staritz@taylor.edu.*

doi:10.1017/S1049096526102108
© The Author(s), 2026. Published by Cambridge University Press on behalf of American Political Science Association.

https://doi.org/10.1017/S1049096526102108 Published online by Cambridge University Press

with resolving conflicts. When students experience disagreement managed through dialogue rather than hostility, they internalize the value of democratic discourse and are more likely to support deliberative institutions in the future (Andersson 2015).[4] In essence, the classroom becomes a microcosm of democratic society in which students learn not only about politics but also about how to participate ethically and effectively.

Despite these benefits, fostering robust deliberation in face-to-face settings remains difficult. Paradoxically, the undergraduate political science classroom—as convenient as it is for assembling an interested audience—still faces deep attitudinal and logistical barriers. Students often arrive with free time, physical proximity, and curricular motivation to deliberate, yet many hesitate to share personal views in front of their classmates. Professors often attribute this reticence to both the fear of peer judgment and the panopticon effect of social media, where any misstep seems indelible and accessible to a wider audience (Goodman 2022). Psychological dispositions including conflict avoidance, varying levels of prior preparation, and a fear of holding a minority view further depress participation, skewing whose voices are heard. These attitudinal hurdles compound practical constraints: in a large lecture with limited class time, it is difficult to guarantee everyone a turn or to simultaneously supervise and assess multiple breakout conversations. Consequently, instructors must surmount a double challenge—creating the space and the willingness for diverse viewpoints to surface—while still meeting other course objectives and assessment demands.

AI AS A POLITICAL DELIBERATION PARTNER

Large language models (LLMs) such as ChatGPT and Claude are now commonplace among undergraduates; surveys have found that majorities of students have tried AI for academic tasks (Cahill and McCabe 2024). This has prompted research into AI's role in political communication (Argyle et al. 2025). This article investigates whether conversing with an AI about politics can confer many of the same benefits as speaking with a human peer while attenuating several obstacles.

A key advantage is the lower perceived social risk when speaking with an AI. Students often regard AI interlocutors as nonjudgmental and impartial, which reduces their fear of embarrassment, criticism, and peer disapproval that can suppress open discourse in traditional classroom settings (Skjuve et al. 2021). This perception fosters a psychologically safe environment in which students feel more comfortable expressing controversial or underdeveloped opinions. More than 80% of US college students admit to self-censoring in class due to a fear of social consequences; AI may provide a novel way to overcome this hesitation and promote authentic self-expression (Goodman 2022). Of course, if this safety makes interaction with an AI too unlike interacting with a human, it may produce none of the benefits an instructor is seeking. However, similarly low-risk interaction such as anonymous discussions also have produced benefits (Rossiter and Carlson 2024).

Students often regard AI interlocutors as nonjudgmental and impartial, which reduces their fear of embarrassment, criticism, and peer disapproval that can suppress open discourse in traditional classroom settings.

Another benefit lies in AI's facilitation of equitable turn taking. In live classroom settings, discussions can be dominated by outspoken students or shaped by existing social hierarchies. AI, by contrast, responds individually and consistently to each student, ensuring that every participant receives thoughtful engagement. This feature promotes inclusivity, allowing quieter and less-confident students space to articulate their ideas without competition. The AI model's responsiveness can mimic ideal discussion norms in which all voices are valued equally. The AI also can be instructed to act as an ideal discussion partner—polite, nonaggressive, and curious—in a manner that might require in-class training or the use of student confederates to achieve in a classroom.

Furthermore, the asynchronous, individual, and flexible nature of AI interactions provides students with valuable opportunities for repeated deliberative practice. Unlike classroom sessions, students can engage with the AI anytime, allowing them to deliberate on political issues at their own pace. This repetition supports the development of deliberative fluency and may build confidence over time, especially for students who may need more time to process complex ideas. The removal of time constraints, the individual nature of the discussion, and the text-based format of the AI interaction (although most platforms also can converse via voice at the time of this writing) may increase accessibility of such in-class exercises, thereby addressing disparities in students' ability to participate in deliberation (Sanders 1997).

Moreover, AI systems offer instructors powerful and versatile tools for scaffolded reflection. Because each dialogue automatically generates a transcript, instructors gain a complete record of every exchange—an evidentiary trail that allows for targeted feedback, assessment of student knowledge, grading via rubrics, and academic-honesty verification. AI also can be scripted to adopt an ideological stance that is the opposite of the user's, ensuring that every student encounters at least one substantively different perspective even in ideologically homogeneous cohorts. The use of AI ensures that discussion partners are available even for those students with niche political interests who might have difficulty finding an informed partner. Through carefully designed prompts, professors can embed cues that encourage metacognitive awareness—inviting students to weigh evidence, consider counterarguments, and identify shared values. When students are guided to think critically about their own reasoning processes, they can develop habits of thought that mirror the goals of deliberative pedagogy.

It is important to note that our research suggests AI-mediated dialogues may yield measurable gains similar to those produced by well-designed human deliberations (Esterling, Neblo, and Lazer 2011; Minozzi et al. 2023). At the same time, scholars caution that substituting AI for peer interaction may backfire if the technology reinforces echo-chamber dynamics, fails to challenge students' ideas, or tempts them to outsource reasoning to an algorithm (Cahill and McCabe 2024). Other concerns, such as ensuring that

https://doi.org/10.1017/S1049096526102108 Published online by Cambridge University Press

the AI is factually accurate and makes effective use of trustworthy sources, also must be addressed. Therefore, understanding how design choices influence political discussion and learning remains an urgent research frontier.

METHOD

To investigate AI's capacity to facilitate civil discourse, we conducted a survey and a learning exercise with undergraduate students. A total of 87 student volunteers (67% female; all from Taylor University) participated in Summer 2024. Students responded to an email invitation sent to a random sample of 800 Taylor students, provided by the university registrar. According to Taylor University policies, students were offered the chance to win a small prize (e.g., a Bluetooth speaker) in exchange for participating (Kelley, Miller, and Staritz 2026).

The survey collected information pre- and post-interaction with the AI. Before the interaction, we measured participants' self-identified political beliefs on a Likert scale indicating the strength and direction of their political leanings. Participants then engaged in a role-playing discussion with an AI persona (i.e., "John") via the Claude.ai interface. Students were instructed to copy and paste a specific role-play prompt into Claude, an AI-powered chatbot, and then hold a conversation in which they interacted with the AI. "John" was prompted to hold specific political views on the topic of immigration, maintaining either conservative or progressive stances in opposition to a student's own views. After completing the conversation, students copied and pasted the full transcript into the survey form for analysis. The following is an example of a prompt for a student who identified as a Republican:

> Your name is John. You hold progressive political views. You are able to see from other perspectives, but you don't change your opinions. You were raised in a family that leaned Democrat, but you've formed your own beliefs through research and life experience. You believe that society should provide a strong social safety net and that the wealthy should pay higher taxes to fund programs that help the less fortunate. You're deeply concerned about climate change and believe the government needs to take stronger action to transition to clean energy and protect the environment. You're supportive of immigration reform that creates a pathway to citizenship for undocumented immigrants already contributing to American communities. However, you don't demonize those who disagree with you politically. You understand that people can have different backgrounds, values, and life experiences that shape their worldview. You're always willing to listen and have respectful discussions. I'm a Republican voter who leans more conservative. I'm just sitting down at the lunch table with you. Start the conversation by saying, "Hey, I've been meaning to talk to you about your thoughts on immigration." Your answers should never exceed 50 words.

The AI was designed to acknowledge opposing arguments while staying committed to its ideological stance—mirroring a principled but courteous human discussant. Its beliefs were framed as shaped by both research and life experience, simulating the thoughtfulness of an actual peer. For example, a conservative "John" might emphasize the importance of increasing deportations and border security; deprioritize accepting refugees; and oppose easing pathways for undocumented immigrants, including family sponsorship and protections for children brought to the United States illegally. In contrast, a progressive "John" might support a strong social safety net, emphasize climate action, and advocate for comprehensive immigration reform—including welcoming refugees and facilitating family reunification. Students were instructed to paste the provided prompt into Claude, participate in a conversation with "John" for at least 10 messages, and then submit their full conversation in the questionnaire.

After the conversation with AI, we assessed attitudes toward political discussion and civic confidence using a battery of the following questions, each rated on a five-point Likert scale ranging from strongly disagree to strongly agree:

- "I consider myself well qualified to participate in politics."
- "I feel that I could do as good a job in public office as most other people."
- "Sometimes politics and government seem so complicated that a person like me can't really understand what's going on."
- "I often don't feel sure of myself when talking with other people about politics and government."
- "In a political discussion, I am comfortable sharing my opinions with others."
- "When discussing politics, I am able to appreciate the views of other people even if I don't agree with them."
- "Talking with people who don't share my political views is a worthwhile use of my time."

The post-interaction survey also included a set of questions to measure the impact of the AI conversation. Students were asked, for example, "Compared to how you saw things before taking part in this research, how much would you say you respect the views of people who disagree with you politically?" Response options were on a five-point Likert scale ranging from "I respect their views a lot more" or "a lot less," "a little more" or "a little less"; "I respect them about the same as before" was a neutral option. Additional Likert-scale (i.e., strongly agree to strongly disagree) items gauged the session's perceived educational value and civic relevance: "I found this session to be helpful and informative"; "Discussions with my fellow citizens on politics are important in our democracy"; and "I feel like I learned a lot from participating in this session."

Students also completed a battery of emotional reactions to the political discussion, indicating the intensity with which they felt emotions such as enthusiastic, hopeful, proud, anxious, worried, afraid, hatred, contempt, bitterness, and anger when they thought about political topics. These items were rated on a four-point Likert scale from strongly agree to not at all.

Students' conversation transcripts were used to verify their engagement with the AI. The majority of students (i.e., more than 90%) exceeded the required 10 messages and demonstrated substantive interaction and engagement.

The majority (i.e., more than 90%) of students went beyond the required 10 messages and demonstrated substantive interaction and engagement.

https://doi.org/10.1017/S1049096526102108 Published online by Cambridge University Press

Transcripts also were reviewed to ensure that the AI had maintained its persona, as instructed, and that there were no deviations or hallucinations. As instructed, the AI was consistently polite, seeking common ground, and positive. It maintained its persona throughout all conversations without any hallucinations or factual errors.

RESULTS

This section summarizes the results of our quantitative and qualitative analyses of students and their engagement with the exercise.

Demographics

The majority of participants (57%) identified as Republican; of those who identified as Independent (36% of the sample), the majority indicated that they "leaned Republican" (74% of those who chose "Independent"); and 74% of participants indicated a moderate political stance. Only a small minority identified as Democrat (7%) and, of those, none described themselves as "strong" Democrats. As a result, in the majority of AI interactions reported herein, the AI took the opposing role as a political progressive. This is typical for a religious university located in a red state; students likely would encounter people with differing perspectives on campus but those with differing views would be few.

Confidence and Engagement

The mean score for confidence in political and civic engagement was 27.8 (of a possible 45) with a standard deviation of 4.3. Participants endorsed the statements "When discussing politics, I am able to appreciate the views of other people even if I don't agree with them" (M=4.12, SD=0.75) and "Talking with people who don't share my political views is a worthwhile use of my time" (M=3.95, SD 1.06) at high rates. However, the rates of endorsement were significantly lower for the other aspects of political and civic engagement: "I consider myself well qualified to participate in politics" (M=2.93, SD=0.93); "I often don't feel sure of myself when talking with other people about politics and government" (reverse coded M=3.04, SD=1.1); and "In a political discussion, I am comfortable sharing my opinions with others" (M=3.53, SD=1.03). The highest endorsed statement was "Discussions with my fellow citizens on politics are important in our democracy" (M=4.52, SD=0.61).

Participant Evaluation of Exercise

Participants strongly endorsed that they found the session helpful and informative (M=4.04, SD=0.75); not a single participant chose to strongly disagree with this statement. The majority also endorsed that they felt like they had learned a lot in the session (M=3.62, SD=0.97). In addition, the majority of participants indicated that they respected those with whom they disagreed about the same or a little more, as summarized in table 1.

Table 1

Student Respect for Those with Opposing Political Views

Compared to how you saw things before you took part in this research, how much would you say you respect the views of people who disagree with you politically?

	Frequency	Percentage*
I respect their views a little less.	5	5.7
I respect them about the same as before.	55	63.2
I respect their views a little more.	24	27.6
I respect their views a lot more.	2	2.3

Note: *Percentages do not total 100 because one participant did not answer this question.

Qualitative Themes

Qualitative analysis of participants' interaction with the AI conversant revealed interesting themes, including possible future directions and potential for ongoing use as a tool for civic discourse. These themes are summarized in the following subsections.

Student Engagement

Perhaps the key question about using AI in deliberative exercises is whether students will engage as seriously with the AI as they would with a fellow student in class and be open to sharing their personal views. If students interact with the AI discussion partner as they would with a machine or appliance—such as by cursorily ordering it to alter its political views or dismissing its statements and opinions as invalid due to its artificial nature—then deliberative exercises using AI partners are unlikely to yield the benefits that they repeatedly have been shown to produce with human partners.

In our research, students engaged seriously with their AI partner. The majority of students conversed with the AI as though it were a person, including expressing an interest in its opinions and engaging in full conversations. This was facilitated by the prompt, which asked the AI to begin the conversation with a question ("I've been meaning to ask about your thoughts on immigration") and to call itself "John" in the discussion. Participants responded to this human-like interaction by calling the AI by its "name," with one exemplar participant saying, "Thanks for asking, John," and another using the more colloquial response: "Hey, John, yeah what sup? What would you like to know?" Interacting with the AI partner as if it were human was the norm such that the participants (approximately 3%) who did not interact in this way stood out significantly.

Participants strongly endorsed that they found the session helpful and informative (M=4.04, SD=0.75); not a single participant chose to strongly disagree with this statement.

One interesting pattern that emerged from this human-like interaction was the willingness of participants to disclose their ignorance to the AI partner. "I myself need to research these issues more, ha ha," one participant wrote. Others were even more frank:

- "I honestly am not an expert on this issue, I just have some opinions based on what I have observed."

https://doi.org/10.1017/S1049096526102108 Published online by Cambridge University Press

- "I don't really know that much about current immigration processes. What part of the process takes so long?"
- "If I am being honest, I am like many Americans in that I know little about our political state and policies."

More than 52% of the participants explicitly stated their lack of knowledge about the subject or acknowledged their own limited experience. It seems unlikely that participants would have had the same willingness to disclose ignorance to their classmates, which suggests that the AI created a psychologically safe space. These examples are supported further by the quantitative data indicating that respondents found the conversation with the AI to be helpful.

Examples of students not treating the AI partner as a person were rare. One exception was the tendency of a few students to demand that AI provide links to sources to support its statements, in contrast to the more informal way we might ask a fellow human discussant, "Where did you hear that?" The AI model used in this research was unable to access the Internet and thus responded much like students in a classroom setting that they could not access that source right now. It is important to note that students did not seem to "break immersion" or stop engaging with the AI when this occurred.

Common Ground Discovery

A number of students (38% of the sample) started their discussion with the AI by taking hardline stances on the immigration issue, either completely supportive or completely condemnatory of increased immigration. These students echoed talking points common in political media concerning immigration, such as specific claims about safety and jobs. One student began the conversation by stating, "I believe that since America was a country founded by immigrants, we should continue to be open and a safe haven for those fleeing their countries." Another student proclaimed, "Countries, such as Venezuela, are sending their criminals here, and there have been an appalling amount of instances where these immigrants have tortured and done horrible things to Americans." Other students repeated talking points. "I don't like the fact that my taxpayer dollars are going to helping people who don't pay taxes stay in the country," one student wrote. Another claimed, "The large number of illegal immigrants is hurting the US citizens in many ways, like drugs and taking jobs from Americans."

As instructed, the AI was willing to listen and ask questions of these students. Often, this resulted in the students moderating their beginning strong position by the end of their discussion with the AI. More than half (54%) of the students added caveats, additional considerations, or conceded points to their AI conversation partner, often adding that they did not have complete information or that they were sharing their limited perspective on a complex topic. "I am open to hearing better solutions to combating these problems as I do not know of many," one student wrote. "To be honest, having grown up in a small community that has not felt the impacts of immigration, I don't necessarily feel the most equipped to prioritize what the biggest problem is," another shared.

Many dialogues also concluded with explicit acknowledgment of shared values (e.g., compassion and the rule of law), reflecting deliberate discussion's theorized ability to reveal underlying consensus. For example, some students who began with strong anti-immigration views spent time with the AI discussing their support for increased legal immigration and ethical treatment of current undocumented migrants in the United States, ending with a recognition of a shared appreciation for the opportunities available in the United States. "I think what concerns me most would be immigrants taking jobs from Americans because companies can pay them less because they are happy to have any work," one student shared, continuing, "I am also concerned about Americans taking advantage of immigrants because they have no other option. Just because they were born in another country doesn't mean they don't have human dignity." The AI explicitly pointed out areas of agreement, usually beginning with a statement of shared values before highlighting policy disagreements. "Those are valid concerns. It's true that worker exploitation can happen, regardless of immigration status," was the AI's reply before this discussion moved to an expanded work-visa program.

Overall, whereas approximately 38% of participants began their discussion with strong, one-sided issue opinion statements—and many more expressed some combination of confusion and talking points—approximately 54% had moderated their initial position or found common ground with their interlocutor by the end of their dialogue. This suggests that discussion with the AI might reduce polarization, similar to the repeated finding that in-person deliberative discussion reduces polarization (Caluwaerts et al. 2023). This also aligns with research showing that deliberation is associated with more elaborate models of political topics (Eveland 2004).

Respect for Disagreement

In post-survey responses, students stated that they believed their respect toward ideological opponents had either stayed the same (approximately 63% of respondents) or increased either a lot or a little (approximately 30%). This aligns with findings that deliberation can temper negative emotions and foster trust. In fact, our results are roughly aligned with other research in which participants interacted with human respondents: both their subjects and ours reported a somewhat similar level of self-assessed improvement in respect for political opponents (Minozzi et al. 2023). It is important to note that this suggests that interacting with an AI is perceived by participants as having an effect on their political interactions with and evaluations of other human beings. In contrast to human interactions, in which the behavior of the conversation partner can limit consensus and lead to further frustration (Pew Research Center 2018), the consistent emotional tone of the AI conversation was calm and prioritized common ground. This may have contributed to how few students reported lessened respect, as well as the fact that they reported a low level of anger on the emotional self-report after their interaction. The use of this particular prompt did not reveal evidence of AI increasing political polarization by reinforcing hostile echo chambers.

DISCUSSION, LIMITATIONS, AND FUTURE RESEARCH

Our preliminary evidence suggests that AI-facilitated conversations may provide benefits to students while overcoming key classroom hurdles. Technology provides an ever-present, patient, and ideologically diverse interlocutor, which allows students to practice democratic discourse free from social risk.

This does not mean that there are not any possible concerns. For example, during our research, we found that some students demanded that their AI interlocutor provide sources for various claims that it made. Given that there have been notable concerns about the ability of AI to provide accurate citations and the tendency of LLMs to hallucinate, this is a potential concern for those who want to use AI as an in-class teaching tool. The model that we chose to use for this project was unable to provide

https://doi.org/10.1017/S1049096526102108 Published online by Cambridge University Press

citations, approximating the likely response of fellow students who are asked to provide sources during class. Instructors who use AI deliberative-discussion partners should anticipate this and investigate their chosen model, as well as including instructions about this in their AI prompt. More broadly, they should be aware of the many other concerns when using LLMs in a political context. For examples, see Goldstein and Sastry (2023).

The fact that many AI models are unfailingly polite is overall, we believe, a virtue of using this tool to facilitate political deliberation. AI models might be instructed to serve as examples of how civil political discussions should be conducted. However, it does raise the possibility that the safety mentioned previously as a benefit also may come at a cost: that is, students who engage with an AI discussant may be missing an element that is more likely to occur during interactions with actual humans. For example, perhaps the AI's calm nature would prevent students from building conflict-resolution and peace-making skills that research shows can be gained through debate and argument (Price, Cappella, and Nir 2002). To some extent, this can be mitigated by an instructor's design of the conversation prompts, instructing the AI to be slightly more aggressive or to simulate some level of irritation.

These drawbacks may be mitigated by not relying solely on student–AI interactions; instead, instructors could combine AI and human interactions. Students could first interact with an AI model, then with their fellow students with the AI as a moderator, and then with their classmates once they have practiced and learned from a patient and calm model (Argyle et al. 2023). Testing this idea of AI as a supplement or role model is an intriguing possibility for future research.

Looking forward, larger and perhaps multi-institutional replications would be needed to accurately estimate effect sizes and fully compare with human interaction. Longitudinal designs could assess whether AI-mediated deliberative practice sustains civic virtues beyond the semester. Testing different AI personas, issue domains, and interaction formats also are all promising avenues of pedagogical inquiry.

CONCLUSION

Integrating AI interlocutors into political science curricula can augment—but not replace—human discussion, providing a complementary approach for students to cultivate knowledge, empathy, and democratic skills. When thoughtfully woven into course design, AI-based deliberate discussion holds promise for a more inclusive, scalable, and reflective civic education.

DATA AVAILABILITY STATEMENT

Research documentation and data that support the findings of this study are openly available at the *PS: Political Science & Politics* Harvard Dataverse at https://doi.org/10.7910/DVN/7DBZYD

CONFLICTS OF INTEREST

The authors declare that there are no ethical issues or conflicts of interest in this research. ■

NOTES

1. This definition loosely follows Freelon's (2010) definition of deliberative discussion as rational discussion of a public issue among equals with ideological differences.
2. There are, of course, important problems and challenges that face political deliberation, such as its inverse relationship with political participation (Mutz 2006) and disparities in students' ability to participate (Sanders 1997).
3. See, however, Mutz (2006).
4. See, however, Settle (2025).

REFERENCES

Andersson, Klas. 2015. "Deliberative Teaching: Effects on Students' Democratic Virtues." *Scandinavian Journal of Educational Research* 59 (5): 604–22. DOI:10.1080/00313831.2014.965795.

Argyle, Lisa P., Ethan Busby, Joshua Gubler, Chris Bail, Thomas Howe, Christopher Rytting, and David Wingate. 2023. "AI Chat Assistants Can Improve Conversations About Divisive Topics." DOI:10.48550/arXiv.2302.07268.

Argyle, Lisa P., Ethan Busby, Joshua Gubler, Alex Lyman, Justin Olcott, Jackson Pond, and David Wingate. 2025. "Testing Theories of Political Persuasion Using AI." *Proceedings of the National Academy of Sciences* 122 (18). DOI:10.1073/pnas.2412815122.

Cahill, Christine, and Katherine McCabe. 2024. "Context Matters: Understanding Student Usage, Skills, and Attitudes Toward AI to Inform Classroom Policies." *PS: Political Science & Politics* 57 (4): 1–5. DOI:10.1017/S1049096524001234.

Caluwaerts, Didier, Kamil Bernaerts, Rebekka Kesberg, Lien Smets, and Bram Spruyt. 2023. "Deliberation and Polarization: A Multi-Disciplinary Review." *Frontiers in Political Science* 5:1127372. DOI:10.3389/fpos.2023.1127372.

Delli Carpini, Michael X., Fay Lomax Cook, and Lawrence R. Jacobs. 2004. "Public Deliberation, Discursive Participation, and Citizen Engagement: A Review of the Empirical Literature." *Annual Review of Political Science* 7:315–44. DOI:10.1146/annurev.polisci.7.121003.091630.

Esterling, Kevin M., Michael A. Neblo, and David M. J. Lazer. 2011. "Means, Motive, and Opportunity in Becoming Informed About Politics: A Deliberative Field Experiment with Members of Congress and Their Constituents." *Public Opinion Quarterly* 75 (3): 483–503. DOI:10.1093/poq/nfr017.

Eveland, William P., Jr. 2004. "The Effect of Political Discussion in Producing Informed Citizens: The Roles of Information, Motivation, and Elaboration." *Political Communication* 21 (2): 177–93. DOI:10.1080/10584600490443877.

Freelon, Deen. 2010. "Analyzing Online Political Discussion Using Three Models of Democratic Communication." *New Media & Society* 12 (7): 1172–90. DOI:10.1177/1461444809357927.

Goldstein, Josh, and Girish Sastry. 2023. "The Coming Age of AI-Powered Propaganda." *Foreign Affairs*, April 7. www.foreignaffairs.com/united-states/coming-age-ai-powered-propaganda.

Goodman, Sylvia. 2022. "The Elusive Civil Classroom." *The Chronicle of Higher Education*, August 29. www.chronicle.com/article/the-elusive-civil-classroom.

Kelley, Katie, Jakob Miller, and Peter Staritz. 2026. "Replication Data for 'AI's Role in Deliberative Discussion.'" *PS: Political Science & Politics*. DOI:10.7910/DVN/7DBZYD.

Minozzi, William, Ryan Kennedy, Michael A. Neblo, Kevin M. Esterling, and Ryan Jewell. 2023. "Testing the Benefits of Public Deliberation." *American Journal of Political Science* 68 (1): 319–34. DOI:10.1111/ajps.12747.

Mutz, Diana C. 2006. *Hearing the Other Side: Deliberative Versus Participatory Democracy*. Cambridge, UK: Cambridge University Press. DOI:10.1017/CBO9780511617201.

Pew Research Center. 2018. "More Now Say It's 'Stressful' to Discuss Politics with People They Disagree With." www.pewresearch.org/politics/2018/11/05/more-now-say-its-stressful-to-discuss-politics-with-people-they-disagree-with.

Price, Vincent, Joseph N. Cappella, and Lilach Nir. 2002. "Does Disagreement Contribute to More Deliberative Opinion?" *Political Communication* 19 (1): 95–112. DOI:10.1080/105846002317246506.

Rossiter, Erin, and Taylor Carlson. 2024. "Cross-Partisan Conversation Reduced Affective Polarization for Republicans and Democrats Even After the Contentious 2020 Election." *Journal of Politics* 86 (4): 1608–12. DOI:10.1086/729931.

Sanders, Lynn M. 1997. "Against Deliberation." *Political Theory* 25 (3): 347–76. DOI:10.1177/0090591797025003002.

Settle, Jaime. 2025. "The Promise and Peril of Interpersonal Political Communication." *Advances in Political Psychology* 46 (1): 167–212. DOI:10.1111/pops.12990.

Skjuve, Marita, Asbjørn Følstad, Knut Inge Fostervold, and Petter Bae Brandtzaeg. 2021. "My Chatbot Companion: A Study of Human–Chatbot Relationships." *International Journal of Human-Computer Studies* 149:102601. DOI:10.1016/j.ijhcs.2021.102601.

https://doi.org/10.1017/S1049096526102108 Published online by Cambridge University Press