Research Design

Grace Flandreau

Table of contents

1 Introduction

Artificial Intelligence is a field that has exploded in the last few years. Software like chat bots have become normalized in workplace and academic environments because of their ease of use and quick answers. With these new developments, research has started to emerge about the impacts of AI on the general population. One topic that has been notable is the privacy concerns associated with AI and its applications. Since many AI algorithms, such as ChatGPT and Gemini, utilize information from the internet to inform users, as well as context from past conversations, there is growing concern about the amount of personal information AI is collecting. The main question I want to answer is: How do American cultural values and political ideology

shape the AI governance framework, particularly with respect to the prioritization of commercial interests over individual privacy protections? This question is imperative to the current policy landscape because of the substantial amount of information that is being collected by these AI sources and potential security risks that may emerge at the personal or national level. Compared to other nations, the United States places a higher priority on free enterprise and technical advancement. Understanding the cultural influences behind AI perception can allow policymakers to create legislation that reflects US attitudes.

The recency of AI technology is arguably one of most important reasons why research on this topic is necessary. At this point, there is still not a strong public understanding of how AI systems actually work or what their long-term implications may be. As people become more familiar with AI and its capabilities, I suspect there will be increasing pressure for stronger regulations surrounding data collection and distribution. This will likely become even more urgent as AI continues to integrate into everyday life. In the past few years, AI has been embedded into common technologies such as email platforms, online shopping, job search tools, and writing assistants. These systems rely on collecting and analyzing large amounts of personal data in order to function effectively. When companies collect this amount of information, concerns arise about corporate misuse and potential data breaches. While issues such as these are not new, the scale and sophistication of AI driven data collection may amplify them. As a result, it is imperative that policymakers and companies strike a balance between technological innovation and privacy protections.

However, public opinion about AI-related privacy concerns is unlikely to be uniform across the United States. Cultural values and political identities often shape how individuals interpret emerging technologies and assess risk. In a highly polarized political environment, attitudes toward regulation, corporate power, and government oversight may influence how

concerned people feel about AI's use of personal data. For example, individuals who prioritize individual liberty may view data collection as a greater threat, while others who emphasize technological progress may focus more on innovation than risk. By examining how US culture and politics shape public opinion on AI privacy concerns, this research can help policymakers better understand the social and political forces that influence support for regulation and data

2 Previous Literature

Since Artificial Intelligence is a relatively new concept, breaking out into mainstream success and recognition in early 2023, there are limits to the research that has been done on the topic. There are AI policies in the US created as early as 2019, with the American AI Initiative and the creation of the National Artificial Intelligence Initiative Office. However, the majority of policies were passed between 2024 and 2025. To get around the issue of a limited amount of policies due to recency of the topic, understanding the American view of internet privacy sheds light on some of the cultural stances that may have an impact on AI policy, since privacy is a common concern for many user input LLM systems. One way to view this is through the privacy paradox. This concept focuses on the gap between the information that a person says they are willing to share compared to what information they actually share (Adjerid, Peer, and Acquisti 2018, 466). An example of this is how opting into personalized advertisements would promote products and services someone is interested in, but would also give companies access to private information. A possible explanation for the willingness to give personal information when concerned about privacy is reliant on the possible benefits one would receive. In a study conducted by Adjerid, Peer, and Acquisti, they test participants to see if the level of privacy protection influences the amount of personal information participants are willing to disclose. The

two surveys supported the idea that increases in perceived privacy protections lead to more information being shared (Adjerid, Peer, and Acquisti 2018, 477–81). This finding elevates what is known about US privacy culture because it supports the idea that Americans prefer to have their information protected, even with the benefits of a potential technology.

Another way social media can be used to examine how people value privacy is through a study aimed to see how clearer messaging affects how people will interact with a social platform, with Instagram used as the platform of interest (Franz and Benlian 2022). The goal of the study was to see if more explicit messaging when asking for access to the participant's address book would make them less likely to share their contacts. In the updated privacy version of the survey, participants were asked if they had the consent of their contacts' to share information, which is something that is assumed in the regular notification (Franz and Benlian 2022, 2298–99). To the expectation of the researchers, having the added insert led to a 62% decrease in participants' disclosure of others' personal information (Franz and Benlian 2022, 2304). This finding suggests that people tend to value their privacy, but are not always aware of the ways that platforms and companies extract data. It is important to note that this study's participants are EU citizens, which may have different opinions on privacy compared to Americans. Despite this, the study demonstrates why clear privacy terms can influence how people view the privacy of their personal information.

Walker's paper on transparency and technology expands on the idea of the privacy paradox by introducing a theory that individuals share their information digitally without fully understanding the ramifications of that decision (Walker 2016, 145). She argues that the "information overload" that the general public experiences in the digital age leads to consumers not having the time or attention span to learn about how to protect their personal data. Because of this, people are more likely to be taken advantage of by private companies looking to use their

data (Walker 2016, 146). Furthermore, she equates the lack of defenses put up by government entities and individuals to the surrender to technology. The significance of this paper is that the average person does not understand the underlying issues that a lack of data privacy brings. While the average person may believe that they value the privacy of their data, when tested on this, people are often willing to share personal information. Transitioning from general privacy concerns to AI specific issues, Walker's 2024 paper with Milne suggests that the introduction of AI has led to greater vulnerability for personal information. The paper explains that AI has become much more commonly used in social media platforms within the last few years. AI can process information much faster than people, which allows platforms to run more efficiently (Walker and Milne 2024). The main issue the authors are concerned about is that data collection is not transparent when LLMs are brought into the picture. There is a worry that marketers and social media executives will value the ability to more accurately target potential consumers over data privacy. This shows a different perspective from Adjerid, Peer, and Acquisti, who observed that Americans can identify strong and weak data protection policies, and choose how to move forward accordingly, often favoring more protection. On the profit centered side of the argument, AI is a tool that allows for a clearer flow between company and consumer, whilst retaining accuracy and lowering costs. Based on this perspective, a laissez faire system for AI integration would help to boost the economy in the United States.

The opinions of AI experts also influence how policies are created surrounding AI usage. Akakpo et al. seek to examine how AI adoption will shape the future of information privacy and to develop a conceptual framework for understanding that relationship, drawing on the perspectives of 42 AI experts to identify six core themes: human agency, data use and abuse, AI transparency and opacity, weaponization of information, cyberbullying, and privacy enforcement (Akakpo et al. 2025, 12). These goals bear directly on the analysis of how US culture and policy

are adapting to AI's expanding role in daily life, as the themes the authors identify map closely onto live debates in American political and regulatory discourse around surveillance, data ownership, and the limits of government oversight. Particularly relevant is the study's forecast that advancing AI will erode meaningful privacy ownership over time, and its observation that the United States has thus far responded with a sector-specific approach as opposed to a unified federal framework (Akakpo et al. 2025, 10–11). This pattern reflects a broader willingness among its users, institutions, and policymakers to accept data abuse as a normal cost of technological progress. The framework offers a valuable analytical vocabulary for understanding not just the technical dimensions of AI and privacy, but the cultural and political conditions that shape how a society chooses to govern them.

When looking at the cultural specificities that shape the US's AI policy, it is relevant to compare with what other governments are doing in this space. Two notable governments to compare are the European Union and China. Both of these regions have measurable cultural differences compared to the United States. Radanliev presents a study that compares the policy frameworks across different countries and regions around the world, focusing on these three major powers. According to this study, all three governments have created policies surrounding ethical standards and safety, while still having unique focuses. As discussed, the US favors policies that advance innovation and highlight scientific integrity and collaboration. The EU is stricter with data collection and has policies that enforce human oversight and non-discrimination. China values more collective frameworks, with a focus on global cooperation and national security (Radanliev 2025, figs. 4, 5, and 6). There are also key differences in how the governments regulate the individual steps of the AI process. The EU is most concerned with data collection, the US with model development, and China with deployment. The way these metrics were constructed was through a system of rating AI ethics policy on eight main dimensions: "transparency,

accountability, fairness, privacy, human oversight, ethical standards, innovation encouragement, and national security" (Radanliev 2025, 20–21). Each of these dimensions was rated on a 1-10 scale using an analysis of public documents, policy white papers, and academic literature. The value of this study as a whole is that it allows a numeric system that depicts the differences in policy. This can help when examining how the US stands out.

The cultural differences between the US and other countries can be seen in the paper by Xing, He, and Zhang, where they conducted a study examining the differences in privacy opinions between US and Chinese people. The research design of the study investigated the opinions of Twitter and Weibo users on AI privacy by using a Python model. Twitter was used to assess US opinions and Weibo, a very similar Chinese platform, assessed the other group. This model was set to the key words "artificial intelligence" and "user privacy," which allowed the model to collect all Tweets, retweets, comments, and user information that pertained to these topics. A similar process was done on Weibo with the equivalent data (Xing et al. 2023, 494–95). The results of the study suggested that Twitter users were more likely to express concerns about security and protection policies, while Weibo users care more about new technologies and applications. Evidence also implied that Weibo users tended to speak more positively about AI, with 56.50% of comments showing support. Twitter users were more pessimistic with 35.60% of users having negative opinions and 40.81% being neutral (Xing et al. 2023, 498–99). The findings from this study take a slightly different direction from Radanliev, whose study suggested that Americans were more likely to support AI innovation when looking at policies that have been created. However, it is important to note that the results from Xing, He, and Zhang are specifically looking at users of a popular online platform, and do not necessarily represent the larger demographics of the respective countries. Despite this, this paper is a valuable one to observe to compare public opinions, as it gives a public perspective, while Radanliev's research leans towards a law maker

perspective. Both sides are important to look at when attempting to see the bigger picture.

Overall, prior research suggests that the average person has concerns about their private information being shared. However, government officials and technology industries value the growth that AI adaptation provides through the ability to increase accuracy while lowering costs. In the testing section of this paper, my hypothesis is: American cultural values rooted in individualism, free-market ideology, and limited government intervention predispose US policymakers to prioritize industry growth and commercial innovation in AI governance, resulting in a regulatory framework that systematically subordinates individual privacy protections to the economic interests of technology firms and the broader capitalist imperative of data-driven profit.

3 Data and Design

In the analysis section of this paper, I will employ a two-strand quantitative design to evaluate the central hypothesis that American cultural values rooted in individualism and free-market ideology predispose US policymakers to prioritize industry growth and commercial innovation in AI governance. This framework suggests that the United States systematically subordinates individual privacy protections to the economic interests of technology firms. Strand 1 will test this cross-nationally, investigating whether a country's Hofstedian cultural profile predicts the rights-orientation of its AI governance framework and whether the United States occupies a culturally explicable position within that distribution (Berry 2023). Strand 2 will examine domestic propositions, asking whether technology industry lobbying expenditure in the US correlates with weaker privacy legislative outcomes over time. This would serve as evidence that market-oriented cultural values create a permissive political environment for industry

influence. Both strands will be implemented entirely in R, proceeding through a sequence of data import, variable construction, exploratory analysis, and inferential modeling.

Data preparation for Strand 1 will begin by importing country-level CSV files using the readr package and merging them into a single analytic dataframe with countrycode as the common key. The cross-national dataset draws on Hofstede's dimension scores as explained by Berry 2023 (Berry 2023). To measure the dependent variable, I will utilize a composite index of AI policy rights-orientation, potentially drawing from Radanliev's numbered score index or Rismani et al.'s Responsible AI measures dataset (Radanliev 2025; Rismani et al. 2025). The rights-orientation index will be treated as approximately continuous, with raw scores drawn from the selected dataset and standardized prior to modeling. The primary independent variables will be drawn from Hofstede's six cultural dimensions: individualism, power distance, uncertainty avoidance, long-term orientation, masculinity, and indulgence. These scores will be standardized using the scale() function. The analytic sample of fifteen to twenty countries is determined by the intersection of countries for which Hofstede scores, AI governance index scores, GDP per capita estimates, and democratic governance ratings are simultaneously available. Countries will be selected to maximize regional and developmental diversity within that constraint, avoiding a sample that disproportionally represents Western liberal democracies, which would artificially compress variation on both the cultural and regulatory dimensions of interest. Countries missing scores on more than two of the six Hofstede dimensions will be excluded from the analytic sample, while remaining missing values will be addressed through mean imputation with sensitivity analyses confirming that results are not driven by imputed values.

The United States will be retained in the analytic sample to determine if its regulatory permissiveness is consistent with its cultural profile or represents a more extreme deviation. Exploratory analysis will include summary statistics via skimr::skim() and a correlation matrix to

screen for multicollinearity. The primary inferential model will be an ordinary least squares regression estimated via the lm() function, controlling for GDP per capita and democratic governance ratings. Although hierarchical modeling might prove advantageous for data nested within specific institutional clusters, the limited scope of the cross-national dataset prevents the complex parameter estimation with such specifications demand. Consequently, an ordinary least squares approach is utilized as a more robust and accessible alternative, employing bootstrapped confidence intervals to address the constraints of a smaller sample. While an ordered logistic model was evaluated in light of the ordinal characteristics of the rights-orientation metric, OLS was ultimately selected to maintain alignment with prevalent governance literature and to treat the index as a nearly continuous variable. Results will be visualized using ggplot2, with a scatter plot of individualism against policy rights-orientation to anchor the findings in the governing research question.

To assess the robustness of Strand 1 findings, the core OLS specification will be re-estimated using Schwartz's cultural value orientations as an alternative to Hofstede's dimensions, drawing on World Values Survey data where country coverage allows. This substitution tests whether the findings are an artifact of the particular cultural measurement framework employed or reflect a more general relationship between individualist cultural orientations and regulatory permissiveness. A secondary specification will also exclude the United States from the estimation sample entirely, allowing the model coefficients to be generated from the remaining cross-national variation before the US is reintroduced as a predicted case.

Strand 2 will shift to a domestic longitudinal analysis to see if the cultural predispositions identified in Strand 1 manifest within the political process. Specifically, I will examine whether technology industry lobbying expenditure systematically suppresses the passage and stringency of privacy-oriented AI legislation in the United States. The federal lobbying

disclosure data is drawn from records maintained under the Lobbying Disclosure Act of 1995, which requires registered lobbyists to file semiannual reports detailing expenditure, employing organization, and the specific legislative issues targeted. These filings are publicly accessible through the Senate Office of Public Records and have been aggregated into machine-readable annual datasets by organizations including OpenSecrets. Expenditure figures will be filtered to include only organizations whose disclosed issue areas include technology, data privacy, artificial intelligence, or telecommunications, ensuring that the lobbying measure captures industry activity relevant to the legislative outcomes under examination rather than total corporate political spending.

Congressional data covering AI and privacy-related legislation from 2016 to the present will be retrieved and filtered via dplyr. Two dependent variable indicators will be constructed: the annual passage rate of relevant legislation and a coded regulatory stringency score for enacted laws. The regulatory stringency score will be constructed through systematic coding of enacted legislation along several dimensions including, whether the law establishes enforceable individual rights such as access, deletion, or consent requirements; whether it creates an independent oversight or enforcement body; whether it imposes meaningful financial penalties for non-compliance, and whether it applies universally or is limited to specific sectors or firm sizes. Each enacted law will be scored on these dimensions and summed into a composite stringency index, with intercoder reliability assessed through a secondary coder on a random subset of legislation.

Lobbying expenditure will be lagged by one year to reflect the temporal gap between political spending and legislative outcomes. This lag reflects both theoretical and practical reasoning: legislative outcomes are unlikely to respond instantaneously to lobbying pressure, as the policy process involves committee review, floor scheduling, and negotiation that typically

unfolds across a full congressional session, and a one-year lag is the most common specification in the interest group literature, ensuring the temporal ordering required to make even a correlational argument about direction of influence. Exploratory analysis will commence with time-series visualization in ggplot2 to identify co-movement between expenditure and passage rates. Pearson correlation coefficients and inferential modeling via lm() will follow, expanding to include GDP and political party control as covariates.

To further assess the stability of Strand 2 findings, the primary model will be re-estimated using alternative lag structures of two years and a distributed lag specification to confirm that the one-year lag assumption does not drive the results. The regulatory stringency score will also be disaggregated into its component dimensions and each modeled separately against lobbying expenditure, testing whether industry spending is more strongly associated with suppression of enforcement mechanisms than with the formal absence of rights provisions, or vice versa. Taken together, these robustness checks across both strands are designed to distinguish findings that are stable across reasonable analytic decisions from those that are sensitive to particular specifications, grounding the paper's conclusions in the most defensible possible evidentiary base. The results of Strand 2 will be interpreted alongside Strand 1 to understand the political mechanism through which cultural values are translated into concrete policy outcomes that subordinate privacy protections to commercial interests.

4 Description of Statistical Results

If the statistical results support the central hypothesis, the findings across both strands would present a coherent picture linking cultural individualism to weakened AI privacy governance and domestic lobbying to the political mechanism through which that cultural

predisposition is realized.

In Strand 1, the scatter plot of Hofstede individualism scores against Radanliev AI ethics scores would not show a simple linear negative relationship across the full sample, but would instead reveal a more nuanced pattern that is nonetheless consistent with the hypothesis. The EU would score highest on the rights-orientation index despite only moderate individualism, while China and India would score low on both dimensions, and the United States would sit at the high end of individualism with among the lowest rights-protective scores in the sample. This pattern would suggest that individualism is a meaningful predictor of weak AI governance specifically among wealthy liberal democracies, where variation in cultural orientation is not confounded by developmental or institutional factors. The observed ethics score of the United States sitting below what its cultural profile would predict, this would suggest that there are other factors that are affecting the regulatory permissiveness of American AI Policy.

In Strand 2, the dual time-series would show lobbying expenditure rising steadily from approximately $52 million in 2016 to over $150 million by 2024, while the annual passage rate of privacy-relevant legislation declines over the same window from around 34 percent to 13 percent. The two series would move in opposite directions across nearly every year in the sample, producing a Pearson correlation of approximately -0.78. The OLS model would confirm this relationship after controlling for party control of Congress and GDP growth, with the one-year lagged lobbying coefficient returning a negative and statistically significant estimate. The Durbin-Watson statistic would fall within the acceptable range, and results would hold under HC1 sandwich standard errors, ruling out autocorrelation and heteroskedasticity as alternative explanations. The regulatory stringency model would produce a consistent but slightly attenuated coefficient, indicating that lobbying suppresses not only the volume of privacy legislation passed but also the strength of what does pass.

These two strands would support the argument that the United States sits at the permissive end of the cross-national distribution in a way that its cultural profile helps explain, while domestic lobbying shapes the cultural orientation that produces legislative outcomes. Neither strand alone would be sufficient, and the findings do not establish causation. When combined, they are consistent with the hypothesis that individualist cultural values create a political environment that industry lobbying then exploits, ultimately subordinating privacy protections to commercial interests.

5 Limitations

Several limitations of this design warrant acknowledgment. In Strand 1, the most significant constraint is sample size, as the cross-national dataset will encompass only approximately fifteen to twenty countries, restricting statistical power and generalizability. This will be partially addressed through bootstrapped confidence intervals via the boot package and by reporting effect sizes alongside conventional p-values. A related concern is that Hofstede's cultural dimension scores are static measures derived from survey data collected decades ago and may not accurately reflect contemporary cultural values. Where possible, supplementary indicators from the World Values Survey will be consulted to assess robustness. Additionally, treating the United States as a cultural monolith presents a meaningful risk, as American society is internally heterogeneous across lines of region, race, class, and political affiliation, and a single national score necessarily flattens that variation. This limitation will be acknowledged in the interpretation of results, and state-level or demographic disaggregation will be noted as a direction for future research. In Strand 2, the primary concern is omitted variable bias, as factors such as electoral cycles and broader public opinion on technology may influence both lobbying intensity and legislative

outcomes independently of the included covariates. The inclusion of party control of Congress and the use of sandwich estimators will partially mitigate this risk. Finally, both strands establish correlation rather than causation, and findings will be framed accordingly as consistent with or in tension with the hypothesis rather than as definitive causal evidence.

6 Conclusion

The cross-national analysis positions the United States as an outlier among comparable liberal democracies, scoring far lower on rights-protective AI governance than its institutional profile would suggest, while the domestic longitudinal analysis shows that rising technology industry lobbying expenditure correlates with declining rates of privacy legislation passage over time. Together, these findings are consistent with the argument that individualist cultural values and free-market ideology create a permissive political environment that industry lobbying then reinforces. Neither strand establishes causation, and the cross-national sample is small, but the analysis offers a structured account of why the United States governs AI the way it does, rooted in the deeper cultural and political conditions that shape what kinds of regulation are considered legitimate in the first place. As AI continues to integrate into everyday life and pressure for stronger data protections grows, understanding those conditions will be essential for anyone seeking to change them.

References

Adjerid, Idris, Eyal Peer, and Alessandro Acquisti. 2018. "Beyond The Privacy Paradox: Objective Versus Relative Risk in Privacy Decision Making1." *MIS Quarterly* 42(2): 465–88. https://misq.umn.edu/misq/article/42/2/465/1704/Beyond-The-Privacy-Paradox-Objective-Versus (April 9, 2026).

Akakpo, Alfred, Evans Akwasi Gyasi, Bentil Oduro, and Sunny Akpabot. 2025. "Artificial Intelligence (AI) and the Future of Information Privacy: Expert Viewpoints." *Journal of Global Information Management* 33(1). https://www.proquest.com/scholarly-journals/artificial-intelligence-ai-future-information/docview/3225540468/se-2?accountid=11752.

Berry, Alison. 2023. *Hofstede's Cultural Dimensions*. 2455 Teller Road, Thousand Oaks California 91320 United States: SAGE Publications, Inc. https://sk.sagepub.com/foundations/hofstedes-cultural-dimensions (May 7, 2026).

Franz, Anjuli, and Alexander Benlian. 2022. "Exploring Interdependent Privacy – Empirical Insights into Users' Protection of Others' Privacy on Online Platforms." *Electron Markets*. https://rdcu.be/e5bjA.

Radanliev, Petar. 2025. "AI Ethics: Integrating Transparency, Fairness, and Privacy in AI Development." *Applied Artificial Intelligence* 39(1): 2463722. https://www.tandfonline.com/doi/full/10.1080/08839514.2025.2463722 (April 9, 2026).

Rismani, Shalaleh et al. 2025. "Responsible AI Measures Dataset for Ethics Evaluation of AI Systems." *Scientific Data* 12(1): 1980. https://www.nature.com/articles/s41597-025-06021-5 (April 9, 2026).

Walker, Kristen L. 2016. "Surrendering Information Through the Looking Glass: Transparency, Trust, and Protection." *Journal of Public Policy & Marketing* 35(1): 144–58. https://journals.sagepub.com/doi/10.1509/jppm.15.020 (April 9, 2026).

Walker, Kristen L., and George R. Milne. 2024. "AI-Driven Technology and Privacy: The Value of Social Media Responsibility." *Journal of Research in Interactive Marketing* 18(5): 815–35. http://www.emerald.com/jrim/article/18/5/815-835/1237824 (April 9, 2026).

Xing, Yunfei, Wu He, Justin Zuopeng Zhang, and Gaohui Cao. 2023. "AI Privacy Opinions Between US and Chinese People." *Journal of Computer Information Systems* 63(3): 492–506. https://www.tandfonline.com/doi/full/10.1080/08874417.2022.2079107 (April 9, 2026).