

Understanding support for AI regulation: A Bayesian network perspective

International Journal of Engineering
Business Management
Volume 17: 1–18
© The Author(s) 2025
Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/118479790251383310
journals.sagepub.com/home/enb



Andrea Cremaschi¹, Dae-Jin Lee¹ and Manuele Leonelli¹ 

Abstract

As artificial intelligence (AI) becomes increasingly embedded in public and private life, understanding how citizens perceive its risks, benefits, and regulatory needs is essential. To inform ongoing regulatory efforts such as the European Union's proposed AI Act, this study models public attitudes using Bayesian networks learned from the nationally representative 2023 German survey Current Questions on AI. The survey includes variables on AI interest, exposure, perceived threats and opportunities, awareness of EU regulation, and support for legal restrictions, along with key demographic and political indicators. We estimate probabilistic models that reveal how personal engagement and techno-optimism shape public perceptions, and how political orientation and age influence regulatory attitudes. Sobol indices and conditional inference identify belief patterns and scenario-specific responses across population profiles. We show that awareness of regulation is driven by information-seeking behavior, while support for legal requirements depends strongly on perceived policy adequacy and political alignment. Our approach offers a transparent, data-driven framework for identifying which public segments are most responsive to AI policy initiatives, providing insights to inform risk communication and governance strategies. We illustrate this through a focused analysis of support for AI regulation, quantifying the influence of political ideology, perceived risks, and regulatory awareness under different scenarios.

Keywords

artificial intelligence governance, Bayesian networks, probabilistic modeling, public attitudes, survey analysis

Received: 8 July 2025; accepted: 5 September 2025

Introduction

Artificial intelligence (AI) is no longer a speculative technology: it is actively shaping how decisions are made, services are delivered, and risks are managed across nearly every domain of life. As AI systems become increasingly embedded in public infrastructure and private platforms, questions about how society should govern these technologies have gained new urgency. At the heart of this debate lies the public: how individuals perceive the risks and benefits of AI, how much they trust regulatory institutions, and whether they support legal constraints. These attitudes shape not only the democratic legitimacy of AI deployment, but also the political feasibility of governance frameworks.^{1,2} Across Europe and beyond, policymakers are seeking to balance innovation with accountability, often amid incomplete or contested public consensus. The

question of how people engage with AI has thus become central to both empirical research and policy design.^{3,4}

The European Union's proposed AI Act provides a timely institutional backdrop. Structured around a tiered, risk-based approach to regulation,^{5,6} the Act seeks to prohibit certain applications outright while tightly overseeing others deemed "high-risk." Although widely praised for its ambition, the Act also raises questions about public understanding and perceived legitimacy.⁷ Early evidence suggests that awareness of the regulation is uneven and that

¹School of Science and Technology, IE University, Madrid, Spain

Corresponding author:

Manuele Leonelli, School of Science and Technology, IE University, Paseo de la Castellana 259, Madrid 28029, Spain.
Email: manuele.leonelli@ie.edu



support varies with individuals' broader beliefs about AI's societal impact, their exposure to information, and political orientation.^{8,9} Yet we still lack a coherent framework for understanding how these factors interact to shape public attitudes.

Existing research has identified many relevant predictors of AI-related beliefs, including age, education, ideology, and trust in institutions.^{10–12} However, most studies use linear or additive models that estimate average effects, often overlooking the conditional relationships among variables and the underlying structure of belief systems.^{13,14} As a result, we know relatively little about how people assemble their views: how, for example, exposure to information interacts with techno-optimism or how policy evaluations mediate support for regulation. Nor do most studies offer a transparent way to simulate how attitudes might shift under different scenarios or in different segments of the population. This lack of structural insight is particularly problematic in the context of AI governance, where beliefs about regulation are often interdependent and shaped by both political cues and perceptions of institutional adequacy.³ Without accounting for these conditional relationships, policy communication strategies risk targeting the wrong segments or overlooking key mediators, potentially leading to ineffective or even counterproductive outcomes.

This paper addresses gaps by modeling public attitudes toward AI using Bayesian networks (BNs)¹⁵ learned from a nationally representative German survey conducted in 2023. Our approach captures the joint structure of beliefs across political, psychological, and informational dimensions, enabling us to explore how attitudes toward AI regulation are shaped by deeper patterns of perception and engagement. We find that support for legal requirements depends strongly on whether people view existing regulations as adequate, while awareness of those regulations is primarily driven by interest and information-seeking behavior. We also show that risk-oriented and opportunity-oriented respondents exhibit distinct logics of belief formation: one more ideologically structured, the other more technocratic. By making these relationships explicit, our model provides a data-driven foundation for designing more targeted communication strategies and responsive governance frameworks.

Literature review

The growing use of data-driven systems in public governance has expanded the role of computational methods in informing decision-making, monitoring risks, and delivering citizen services. From fraud detection to automated decision-making and citizen assistance, advanced analytics have become central to managing both operational complexity and public expectations.^{16,17} As AI is increasingly deployed in sensitive domains, understanding how citizens perceive these technologies, and how such perceptions vary

across population groups, has become a strategic concern for effective policy design and regulation.

Public attitudes toward AI have been studied from diverse disciplinary perspectives, highlighting the roles of sociodemographics, ideology, trust, and regulatory awareness. While these factors are interdependent, existing research models them in isolation, using approaches that overlook the conditional structure of belief systems. The next sections review key predictors of AI-related attitudes and identify methodological gaps our study aims to address.

Sociodemographic drivers of attitudes

A broad literature investigates the factors shaping public perceptions of AI and automation. Demographic characteristics such as age, gender, and education are consistently associated with distinct attitudinal profiles: younger individuals and those with higher education levels tend to express greater optimism toward AI, while older and less educated respondents are more skeptical or fearful.^{1,4,10} These associations have been observed across national contexts and survey instruments. Gender effects also emerge, with men typically expressing more confidence in AI and women showing greater concern about potential risks or inequalities.

However, demographics alone do not explain the complexity of AI perceptions. Psychological constructs such as technological self-efficacy, anxiety, and institutional trust have been shown to mediate the effects of background characteristics.^{3,14} In particular, AI-specific anxiety (encompassing concerns about bias, transparency, and control) has emerged as a distinct dimension that influences both expectations and policy preferences. Trust in government, science, and the media further modulates attitudes, with high-trust individuals more likely to accept the integration of AI in public life.

Political ideology also plays a central role. Research has shown that support for AI and automation is often polarized along left–right lines, especially when regulation or labor concerns are salient.^{9,18} Importantly, trust and ideology interact: political identity influences perceptions of institutional legitimacy, which in turn shapes how information about AI is received and evaluated. Cross-national comparisons reveal that these patterns are contingent on cultural and economic context. For instance, studies from Southern Europe and Latin America have found that fears about AI-driven job loss are particularly acute in low-trust environments, where technological change is perceived as externally imposed or unregulated.^{11,12,19}

Regulatory preferences and legal frameworks

Public support for AI regulation is generally high but conditional. Studies show that support tends to increase when individuals are made aware of specific risks, such as

misuse in surveillance, discrimination in automated decisions, or the erosion of human agency.^{2,4} Regulatory support is also stronger when the perceived domain of application involves public goods, such as health care, education, or democratic participation. Yet this support is not uniform: ideological priors, institutional trust, and personal experience with technology all influence whether individuals see regulation as protective, burdensome, or symbolic.

Awareness of regulatory efforts, such as the EU AI Act, also varies across population segments. Recent studies suggest that public knowledge of such initiatives is limited, and that support depends not only on factual awareness but on how the regulation is framed and evaluated.⁸ Some individuals support legal restrictions in principle but view current institutional efforts as either inadequate or excessive; others may reject specific regulations while expressing trust in broader governance processes. These findings suggest that regulatory preferences are best understood as embedded within multidimensional belief systems, rather than as direct functions of exposure or ideology alone.

Modeling challenges and interpretability

Most empirical studies of public attitudes toward AI use linear regression, logistic models, or structural equation modeling to test directional hypotheses.^{13,14} While effective for identifying average effects, these methods can obscure conditional dependencies, mediated paths, or interaction patterns that may be central to belief formation. Moreover, they typically rely on predefined structures, limiting the ability to discover emergent associations from data. This can be especially limiting when attempting to uncover causal mechanisms or identify actionable targets for communication, as linear models often fail to detect mediated, non-additive, or context-dependent pathways.^{20,21}

Recent work on interpretable machine learning emphasizes the value of transparent, flexible models in high-stakes domains.²² Probabilistic graphical models, including BNs, provide a structured way to capture the joint distribution of beliefs while retaining interpretability. BNs support both exploratory and confirmatory analysis, allowing researchers to model how information flows through a system of interconnected variables.

In public policy and social science research, BNs have been applied to domains such as environmental risk perception,²³ stakeholder engagement,²⁴ public health communication,²⁵ and EU climate policy modeling.²⁶ These studies demonstrate that BNs are especially valuable when variables are interdependent and when inference under uncertainty is required. BNs are well suited to the analysis of complex public opinion systems²⁷ like those surrounding AI and its governance. Their ability to reveal hidden belief structures, simulate counterfactuals, and perform subgroup inference makes them particularly promising for this emerging area.

We contribute to this literature by applying BNs to the study of AI-related public attitudes using nationally representative European survey data. Methodologically, our approach captures how variables interact within a structured system of conditional dependencies, uncovering indirect pathways and belief hierarchies often missed by linear models. Substantively, we show how regulatory support is not simply the result of demographic traits or isolated attitudes, but emerges from a dynamic interplay of trust, exposure, and techno-perception. By making these relationships explicit, our model supports transparent scenario-based inference and offers practical tools for tailoring AI governance strategies to diverse public segments.

Materials and methods

Data and feature engineering

We use data from the survey *Current Questions on AI (June 2023)*, conducted by FORSA on behalf of the Press and Information Office of the German Federal Government and archived by GESIS.²⁸ The dataset consists of interviews with a probability-based sample of 1506 individuals aged 14 and above, representative of the German-speaking population. Data were collected over 3 days (June 26–28, 2023) using computer-assisted telephone interviewing. Respondents were selected using a dual-frame sampling strategy combining landline and mobile numbers, in accordance with standard practice in Germany.²⁸ The survey provides non-aggregated, individual-level responses to a wide array of questions concerning AI. Given the scarcity of recent public datasets with this level of granularity, it offers a valuable snapshot of current public perceptions.

The questionnaire covers interest in and exposure to AI (e.g., media usage, information-seeking, interpersonal conversations); general perceptions of AI as a risk or opportunity; attitudes toward AI regulation, including views on the EU AI Act; beliefs about AI's social impacts (e.g., in healthcare or job automation); and sociodemographic and political background variables, such as age, sex, education, income, and voting intention. Respondents who identified AI as primarily a risk or an opportunity were asked to elaborate in open-ended follow-ups. These answers were then coded by the survey team into discrete variables capturing key thematic concerns, and are included in the released dataset.

We harmonized and recoded the original survey variables for consistency. Ordinal variables (e.g., interest in AI, perceived information level) were recoded into labeled four-point scales, while binary exposure items (media use, discussions, information search) were standardized as “Yes” or “No.” Rare responses such as “Don’t know” were treated as missing (threshold: <50 cases). Voting intention was recoded into a three-level political preference variable (Left, Right, Other), grouping non-voters and unclassifiable

answers under “Other.” Observations with missing values in the selected variables were excluded. The set of variables used is listed in Table 1.

Two additional datasets were created based on whether respondents viewed AI as a risk, an opportunity, or both. Those selecting “opportunity” or “both” provided open-ended responses on perceived benefits, coded into 24 binary indicators by the survey team. These were grouped into five themes: work and automation; health and care; everyday life and education; technological innovation; and general positivity. Similarly, responses from those identifying AI as a “risk” or “both” were coded into 18 binary indicators and grouped into five themes: job displacement and dehumanization; loss of autonomy and control; misuse and regulatory failure; data security and misinformation; and social uncertainty. These grouped thematic variables are also summarized in Table 1.

Bayesian Network modeling

Basic principles of Bayesian networks. Multivariate data often involve complex dependencies that make direct modeling of the joint distribution intractable, especially as the number of variables increases. A common strategy is to simplify the joint distribution by leveraging assumptions about how variables relate to each other. One standard approach is to represent these relationships using a graph structure that encodes conditional dependencies. In Bayesian networks (BNs), this structure is a directed acyclic graph (DAG), where each node represents a random variable, and directed edges indicate direct probabilistic influence.^{15,29} For each variable, the set of *parents* consists of those variables with edges pointing into it. This graphical structure then determines how the full joint distribution factorizes into simpler conditional distributions.

Formally, let X_1, \dots, X_n be n discrete random variables and let \mathcal{G} be a DAG over these variables. The joint distribution encoded by the BN factorizes as:

$$P(X_1, X_2, \dots, X_n) = \prod_{i=1}^n P(X_i | \text{Pa}_{\mathcal{G}}(X_i)),$$

where $\text{Pa}_{\mathcal{G}}(X_i) \subseteq X_1, \dots, X_n \setminus X_i$ denotes the set of parent variables of X_i in the graph \mathcal{G} . This decomposition allows for a compact and modular representation of the joint distribution by exploiting conditional independence assumptions.

In our application, all variables are treated as discrete. Ordinal variables are modeled as categorical, and conditional probability tables are estimated directly for each configuration of parent variables.

Learning Bayesian networks. Both the network structure and the conditional probability tables were learned directly from the data using a fully data-driven approach. Structure learning was carried out using a score-based method that

combines the Tabu search algorithm³⁰ with the Akaike Information Criterion (AIC), a well-established metric that balances model fit and complexity. Given a graph structure \mathcal{G} and dataset \mathcal{D} , the AIC score is defined as:

$$\text{Score}_{\text{AIC}}(\mathcal{G} | \mathcal{D}) = \log P(\mathcal{D} | \mathcal{G}, \hat{\theta}) - d,$$

where $P(\mathcal{D} | \mathcal{G}, \hat{\theta})$ is the likelihood of the data under the maximum likelihood parameters $\hat{\theta}$, and d is the number of free parameters in the model. For discrete variables, the log-likelihood can be written as:

$$\log P(\mathcal{D} | \mathcal{G}, \hat{\theta}) = \sum_{i=1}^n \sum_{j=1}^{q_i} \sum_{k=1}^{r_i} N_{ijk} \log \frac{N_{ijk}}{N_{ij}},$$

where X_i is a node with r_i possible states and q_i parent configurations, and N_{ijk} is the number of observations where $X_i = x_k$ and its parents are in configuration x_j . The AIC is particularly suitable for our setting, offering a more permissive penalty than BIC, which is appropriate given the moderate sample size and number of variables.³¹

To improve robustness, we applied a non-parametric bootstrap procedure,³² generating 2000 resampled datasets and learning a directed acyclic graph (DAG) for each. The final consensus network was obtained using the *averaged network* method in the *bnlearn* package,³¹ which optimizes the inclusion threshold for each edge based on expected predictive accuracy.

To guide learning and enhance interpretability, we enforced a tiered structure using blacklists. Demographic and political variables were set as roots and were not allowed to receive incoming edges from downstream concepts such as exposure, perception, or regulation. AI-related perceptions and regulatory attitudes were allowed to depend on exposure but constrained from influencing earlier tiers.

Once the structure was learned, the conditional probability tables were estimated using Bayesian parameter learning with a uniform Dirichlet prior.³³ This regularization approach, equivalent to Laplace smoothing commonly used in machine learning,³⁴ prevents zero-probability estimates in sparse configurations by assigning a small positive probability to all possible outcomes, including those not observed in the sample. Specifically, given observed counts N_{ijk} for variable X_i taking value x_k with parent configuration x_j , the posterior mean estimate is:

$$\hat{P}(X_i = x_k | \text{Pa}(X_i) = x_j) = (N_{ijk} + \alpha) / (N_{ij} + r_i \alpha),$$

where $N_{ij} = \sum_k N_{ijk}$, r_i is the number of states of X_i , and $\alpha = 1$ corresponds to a non-informative uniform prior. This estimator shrinks all probabilities away from the extremes and guards against overfitting in small samples. Importantly, it does not imply that all unobserved events are equally likely, nor that they are plausible in a substantive sense, only that they are not impossible under the model.

Table 1. List of variables used in the BN analysis, including structured survey responses and grouped thematic codes derived from open-ended justifications about AI. Variables are grouped by conceptual domain: demographics, exposure to AI, perceptions, political orientation, regulatory opinions, and emergent themes reflecting perceived opportunities or risks. For selected demographic variables, sample proportions are shown in parentheses.

Group	Variable	Description (question or theme)	Levels
Demographics	Sex	Respondent's gender	Female (47.4%), male (52.6%)
	Age	Respondent's age group	14–29 (8.2%), 30–44 (14.7%), 45–59 (28.2%), 60+ (48.4%), Other (0.5%)
	Education	Highest level of completed education	Low (11.6%), medium (26.2%), high (59.7%), Other (2.5%)
	Income	Self-reported income category	Low (18.3%), medium (29.7%), high (34.9%), Other (17.1%)
	Municipality	Size of municipality of residence	<5k (13.9%), 5k–19k (26.7%), 20k–99k (28.1%), 100k–499k (13.8%), 500k+ (17.5%)
Politics	VoteIntent	Political orientation based on voting intention	Left (29.8%), right (34.7%), Other (35.5%)
Exposure	InterestAI	How interested are you in information about AI?	Not at all, less strongly, strongly, very strongly
	InformedAI	How well informed do you feel about AI?	Very poor, rather poor, rather good, very good
	MediaAI	Have you recently seen/heard anything about AI in the media?	Yes, No
	FriendsAI	Have you discussed AI with friends/acquaintances?	Yes, No
	SearchAI	Have you actively searched for information about AI?	Yes, No
Regulation	AIRegulations	Should there be legal requirements for AI?	Yes, No
	HeardEURegulation	Have you heard of the EU AI regulation?	Yes, No
	EUAppropriateRegulation	What is your view of the EU AI regulation?	Appropriate, too strict, not strict enough, Don't know
Perceptions	AIReduceShortageWorkers	AI will reduce the impact of labor shortages	Strongly disagree to strongly agree
	AI Easier Life	AI will make everyday life easier	Strongly disagree to strongly agree
	AI Healthcare Benefit	AI will bring major benefits in healthcare	Strongly disagree to strongly agree, Don't know
	AI Field Benefit	I Can imagine which job areas benefit from AI	Strongly disagree to strongly agree
	AI vs Human	AI will produce content indistinguishable from human work	Strongly disagree to strongly agree
	AI False Info	AI will lead to greater spread of misinformation	Strongly disagree to strongly agree
	AI Uncontrollable Develop AI	I Fear AI will become uncontrollable Do you see AI developments as a risk or opportunity?	Strongly disagree to strongly agree Risk, Opportunity, both
Opportunity themes	PosWork	AI as beneficial for the labor market, automation, and job creation	Mentioned, not mentioned
	PosHealth	AI as a tool to advance healthcare, medicine, and care services	Mentioned, not mentioned
	PosLife	AI improving daily life, efficiency, and educational contexts	Mentioned, not mentioned
	PosTech	AI associated with innovation, future technology, and research	Mentioned, not mentioned
	PosGeneral	General positive outlook, openness to AI, or hope for risk mitigation	Mentioned, not mentioned

(continued)

Table 1. (continued)

Group	Variable	Description (question or theme)	Levels
Risk themes	RiskJobs	Fears of job loss, redundancy, or dehumanization in the workplace	Mentioned, not mentioned
	RiskLossOfControl	Concerns about losing control, autonomy, or freedom due to AI	Mentioned, not mentioned
	RiskMisuseRegulation	Ethical concerns, potential misuse, and inadequate regulation	Mentioned, not mentioned
	RiskData	Worries about data privacy, surveillance, or misinformation	Mentioned, not mentioned
	RiskSociety	Uncertainty about AI's societal impact, especially in education	Mentioned, not mentioned

The fitted models are publicly available via the bnRep R package.³⁵

Analyzing the Bayesian network model. Given the learned BN, we analyze its structure using probabilistic inference and sensitivity methods. These tools help identify key relationships, quantify the influence of individual variables, and assess the robustness of conclusions.

We perform evidence propagation by conditioning on specific variable states and examining the resulting changes in marginal distributions, highlighting directional effects and dependencies. Scenario-based inference is used to simulate joint configurations of interest and analyze downstream consequences on target variables.

To assess global variable importance, we compute Sobol variance indices,³⁶ which decompose the variance of a target variable into additive contributions from other nodes. These indices measure how much of the uncertainty in a target's distribution is attributable to variation in a given input. For a target variable Y , the first-order Sobol index for input X_i is:

$$S_i = \text{Var}_{X_i}(\mathbb{E}(Y|X_i)) / \text{Var}(Y). \quad (1)$$

This quantifies the proportion of output variance explained by X_i alone, marginalizing over all other variables. The resulting scores serve a role analogous to coefficients in regression, but capture both direct and indirect (mediated) influences within the BN structure.

Finally, we conduct local sensitivity analysis by perturbing entries of the conditional probability tables and quantifying their effect on selected target distributions.^{37,38} This reveals which parameters most strongly influence predictions and supports robustness checks of specific pathways. Together, these analyses offer a layered understanding of the BN's logic and predictive behavior, enabling interpretable, data-driven insights.

Results

We begin by presenting the structure and implications of the BN estimated over all respondents. We then proceed to

analyze two submodels learned separately for individuals who primarily perceive AI as a risk and those who see it as an opportunity.

Bayesian network over all respondents

Network structure and dependencies. In the network over all respondents (Figure 1), *InterestAI* emerges as a central hub, directly influencing both beliefs and information behaviors. It connects to a wide array of downstream nodes, including *DevelopAI* (the core framing of AI as risk/opportunity), as well as *AIEasierLife*, *AIFieldBenefit*, *InformedAI*, *MediaAI*, *FriendsAI*, *SearchAI*, and *HeardEURegulation*. This suggests that personal engagement with AI serves as a primary lens through which both informational exposure and attitudinal formation are structured.

DevelopAI itself feeds directly into *AIEasierLife*, which acts as a secondary hub: it shapes beliefs around AI's utility (*AIFieldBenefit*, *AIReduceShortageWorkers*), societal consequences (*AIUncontrollable*), and indirectly, through further downstream effects, beliefs about AI's role in healthcare (*AIHealthcareBenefit*) and misinformation (*AIFalseInfo*, via *AIUncontrollable*). Notably, *AIUncontrollable* links to both *AIFalseInfo* and *AIvsHuman*, and feeds into concerns about regulatory sufficiency through *EUAppropriateRegulation*, which in turn influences *AIRegulations*.

The regulation-related structure features two converging paths. One starts from *InterestAI*, proceeds through *HeardEURegulation*, and terminates in *AIRegulations* via *EUAppropriateRegulation*. A parallel path stems from *VoteIntent*, which exerts direct influence on *MediaAI*, *AIRegulations*, and *HeardEURegulation*. The convergence of these pathways at *AIRegulations* highlights the interplay between perceived institutional exposure and political predispositions in shaping regulatory support.

Demographic variables lie upstream in the structure. *Sex* is a root node influencing *InterestAI*, *Age*, *Education*, and *Income*. *Education* is a pivotal factor, shaping not only *InterestAI*, but also *FriendsAI*, *SearchAI*, *Age*, *Income*, *Municipality*, and *VoteIntent*. *Age* further informs *MediaAI*, *AIEasierLife*, *Income*, and *VoteIntent*. Finally, *Municipality*

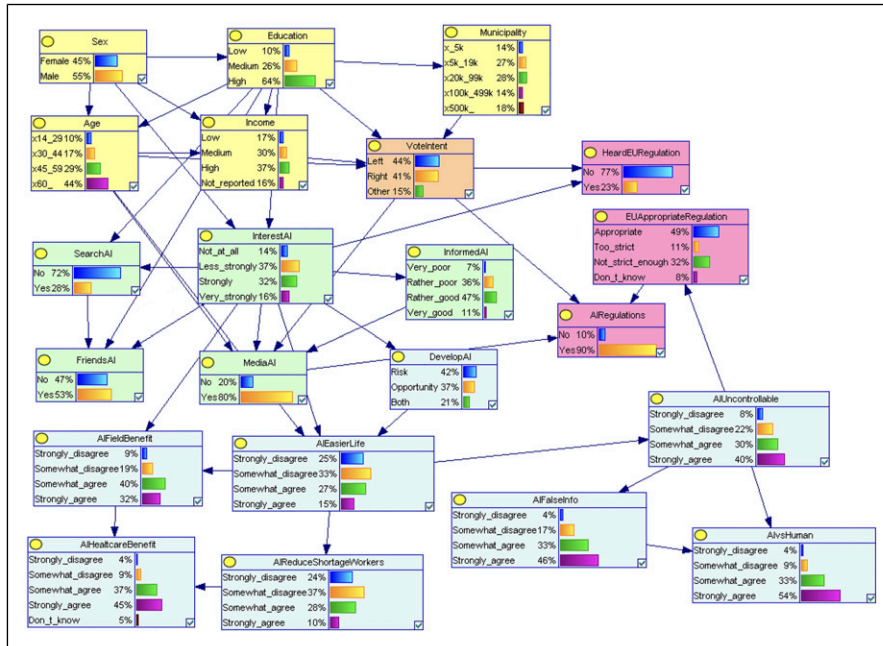


Figure 1. BN learned from all respondents. Thematic “Risk” and “Opportunity” variables are not included, as they were only collected for specific subgroups. Nodes are color-coded by variable group: demographics (yellow), exposure (green), regulation (pink), political orientation (orange), and AI perceptions (blue). Visualization produced in GeNIe.

and *Income* also feed into *VoteIntent*, which acts as a gateway from sociodemographics into the political and attitudinal core of the network.

Variance decomposition via Sobol indices. Table 2 reports variance-based Sobol indices, which quantify the proportion of output variance attributable to each input variable, for five representative outcomes capturing legislative attitudes (*AIRegulations*, *EUAppropriateRegulation*), AI-related beliefs (*DevelopAI*, *AIUncontrollable*), and policy awareness (*HeardEURegulation*).

Perceptions of AI are primarily shaped by expectations about its usefulness in daily life. The belief that AI will make everyday life easier is the strongest single contributor to how respondents frame AI (49.6% of variance in *DevelopAI*), followed by self-reported interest in AI (37.3%). This suggests that optimistic, personally relevant benefits play a dominant role in shaping overall attitudes. Notably, fear that AI may become uncontrollable (*AIUncontrollable*) is itself driven by the same belief about AI’s benefits (41.0%), but also by concerns over misinformation (*AIFalseInfo*, 38.2%) and the perception that existing regulation is too weak (*EUAppropriateRegulation*, 28.6%). This combination of optimism and anxiety, mutually reinforcing or counterbalancing, highlights the layered nature of public sentiment toward AI technologies.

Legislative attitudes are influenced through different pathways. Awareness of the EU AI regulation (*HeardEURegulation*) is shaped almost entirely by engagement and information-seeking behaviors, especially interest in AI

(70.0%) and online search activity (15.8%), with additional influence from political orientation (26.6%). In contrast, support for legal regulation (*AIRegulations*) depends more on perceived policy adequacy: views on whether the EU AI Act is appropriate explain 32.9% of variance, followed by political alignment and media exposure. Finally, opinions on whether EU regulation is too strict or too lenient (*EUAppropriateRegulation*) are overwhelmingly driven by techno-anxiety, particularly fear of losing control over AI systems (87.7%). These results reveal a clear division: perceptual engagement and information-seeking shape awareness, while emotional responses to risk, especially anxiety over control, are central to policy evaluation.

Conditional probabilities and belief dynamics. We next examine how input variables with an explanatory contribution above 10% influence the distribution of key target variables when entered as evidence. Tables 3–7 display conditional probability tables for the five main outcomes of interest, showing how their distributions shift across relevant predictor categories.

As shown in Table 3, personal attitudes and perceived utility are key drivers of how respondents interpret AI developments. Those who believe that AI will make life easier are substantially more likely to frame it as an opportunity (e.g., 67.9% for “Strongly agree”) than as a risk (14.6%). A similar gradient appears for interest in AI: respondents reporting very strong interest choose “Opportunity” (51.0%) or “Both” (23.7%) far more often than “Risk” (25.3%). These shifts reinforce the earlier variance

Table 2. Variance-based Sobol indices (percentage of output variance explained) for each input variable and five target outcomes in the BN model. Bold values indicate variables that explain more than 10% of the variance for the corresponding output. Dashes (–) indicate the variable was the target and thus excluded from its own analysis.

Input variable	DevelopAI	HeardEURegulation	AIRegulations	EUAppropriateRegulation	AIUncontrollable
AI Easier Life	49.6%	9.5%	0.0%	10.6%	41.0%
Interest AI	37.3%	70.0%	0.5%	1.3%	5.6%
AI Field Benefit	11.3%	8.5%	0.0%	1.2%	4.8%
Search AI	7.1%	15.8%	0.0%	0.3%	1.2%
Friends AI	6.2%	11.6%	0.1%	0.2%	1.0%
AI Uncontrollable	6.1%	1.1%	0.4%	87.7%	–
Informed AI	5.6%	11.7%	0.5%	0.2%	0.9%
AI Reduce Shortage Workers	4.5%	0.8%	0.0%	0.9%	3.4%
Education	2.4%	6.2%	0.3%	0.1%	0.6%
AI Healthcare Benefit	1.9%	1.1%	0.0%	0.3%	1.1%
Heard EU Regulation	1.8%	–	0.3%	0.1%	0.3%
Media AI	1.8%	4.7%	12.3%	0.0%	0.1%
Sex	0.9%	1.7%	0.0%	0.0%	0.2%
AI False Info	0.7%	0.1%	0.0%	8.4%	38.2%
EU Appropriate Regulation	0.6%	0.1%	32.9%	–	28.6%
AI vs Human	0.3%	0.1%	0.0%	5.0%	13.3%
Income	0.2%	1.8%	0.6%	0.0%	0.1%
Age	0.2%	0.3%	0.5%	0.6%	2.3%
Vote Intent	0.1%	26.6%	13.4%	0.0%	0.0%
Municipality	0.0%	0.7%	0.2%	0.0%	0.0%
Develop AI	–	3.9%	0.0%	1.7%	7.2%

decomposition, which identified *AI Easier Life* and *Interest AI* as the strongest contributors to variance in *Develop AI*. At the opposite end, respondents who strongly disagree with the idea that AI will ease daily life are more than 65% likely to view it as a risk, confirming that techno-optimism plays a protective role against risk-based framing.

As shown in [Table 4](#), the strongest contrasts in regulatory awareness (*Heard EU Regulation*) emerge between individuals with differing levels of AI interest and information engagement. Those with no interest in AI have a 93.2% probability of being unaware of EU regulations, compared to just 59.8% among those very interested. Awareness also rises with self-reported knowledge (*Informed AI*) and discussion with peers (*Friends AI*).

[Table 5](#) shows that support for AI-specific legal regulation (*AI Regulations*) is primarily shaped by policy evaluations: respondents who view current regulation as “Not strict enough” support legal requirements at a rate of 93.4%, while those who believe it is “Too strict” show only 72.9% support. This suggests that perceptions of regulatory adequacy serve as a key driver of willingness to endorse legal constraints.

[Table 6](#) illustrates how views on the appropriateness of EU AI regulation are shaped by both perceived benefits and risks. Individuals who strongly believe that AI will make life easier are more likely to consider existing regulation “Appropriate” (53.1%) or even “Too strict” (13.2%),

suggesting that techno-optimism dampens regulatory concern. In contrast, those who fear AI may become uncontrollable express higher rates of dissatisfaction: among respondents who “Strongly agree” with this concern, nearly 48% believe regulation is “Not strict enough,” compared to just 16.5% of those who “Strongly disagree.” These contrasts highlight the role of emotional responses in shaping not just support for regulation, but evaluations of institutional adequacy itself.

Concerns about AI becoming uncontrollable (*AI Uncontrollable*) reflect a layered interplay of optimism, anxiety, and policy evaluation. As shown in [Table 7](#), individuals with strong techno-optimism (*AI Easier Life: Strongly agree*) are less likely to express high concern (23.3% for “Strongly agree”) than those who reject AI’s usefulness (63.9%). However, the most influential factors stem from techno-anxiety: respondents who strongly believe AI spreads misinformation (*AI False Info: Strongly agree*) have a 55.8% probability of fearing AI is uncontrollable. Similarly, those who perceive current regulation as “Not strict enough” are 59.8% likely to express strong concern, compared to just 28.7% among those who see it as “Appropriate”.

Joint scenario reasoning. BNs support joint reasoning under multiple sources of evidence. Instead of conditioning on a single variable, we define realistic *scenarios* by fixing combinations of age, education, interest in AI, political

Table 3. Conditional probabilities for DevelopAI, grouped by evidence. Baseline shown in the first row.

Evidence variable	Evidence value	Both	Opportunity	Risk
Baseline		0.2143	0.3666	0.4191
AIEasierLife	Strongly disagree	0.1819	0.1534	0.6648
	Somewhat disagree	0.2420	0.2713	0.4867
	Somewhat agree	0.2309	0.5133	0.2558
	Strongly agree	0.1753	0.6785	0.1463
AIFieldBenefit	Strongly disagree	0.1782	0.2314	0.5904
	Somewhat disagree	0.2035	0.2925	0.5040
	Somewhat agree	0.2223	0.3727	0.4050
	Strongly agree	0.2209	0.4421	0.3370
InterestAI	Not at all	0.1340	0.2095	0.6565
	Less strongly	0.1823	0.3600	0.4577
	Strongly	0.2761	0.3724	0.3516
	Very strongly	0.2372	0.5101	0.2527

orientation, and information exposure. Table 8 presents ten illustrative profiles, ranging from the *Young Informed Left* to the *Disengaged Elder*, to compare attitudes across diverse population segments.

Figure 2 shows the predicted probabilities for three central variables. The *DevelopAI* plot reveals how optimistic beliefs (viewing AI as an opportunity) are prevalent in highly informed and interested groups, while apathy and disengagement drive risk-dominant perceptions. Support for legal regulation of AI (*AIRegulations*) is remarkably high overall but varies modestly with political and engagement variables, with *Young Apathetic* and *Disengaged Elder* profiles showing more skepticism. In contrast, concern over loss of control (*AIUncontrollable*) shows wide variation: techno-anxiety peaks in disengaged and older respondents and subsides among those with balanced or moderate exposure.

Table 4. Conditional probabilities for HeardEURegulation, grouped by evidence. Baseline shown in the first row.

Evidence variable	Evidence value	No	Yes
Baseline		0.7741	0.2259
FriendsAI	No	0.8167	0.1833
	Yes	0.7360	0.2640
InformedAI	Very poor	0.8593	0.1407
	Rather poor	0.7977	0.2023
	Rather good	0.7640	0.2360
	Very good	0.6839	0.3161
InterestAI	Not at all	0.9317	0.0683
	Less strongly	0.8204	0.1796
	Strongly	0.7382	0.2618
	Very strongly	0.5980	0.4020
SearchAI	No	0.8031	0.1969
	Yes	0.6981	0.3019
VoteIntent	Left	0.7091	0.2909
	Right	0.8072	0.1928
	Other	0.8717	0.1283

Parametric sensitivity analysis. To identify which upstream nodes most strongly influence each target, we conducted a parametric sensitivity analysis alongside the Sobol variance decomposition. For conciseness, sample results are shown in Appendix Figures 5–6, including DAGs colored by local sensitivity values and a tornado plot for *HeardEURegulation*. This analysis quantifies how small changes in the conditional probabilities of ancestral nodes affect downstream distributions. For example, *DevelopAI* is most sensitive to *InterestAI*, with additional effects from *Sex* and *Education*; *HeardEURegulation* is highly influenced by *InterestAI* and *VoteIntent*; and *AIRegulations* is primarily driven by *EUAppropriateRegulation*.

Subpopulation networks: Risk and opportunity

As before, both subpopulation models are estimated on the same backbone of demographic, exposure, and policy-opinion variables, but now the detailed AI-perception block includes the *grouped thematic variables* derived from open-ended responses.

Table 5. Conditional probabilities for AIRegulations, grouped by evidence. Baseline shown in the first row.

Evidence variable	Evidence value	No	Yes
Baseline		0.0996	0.9004
EUAppropriateRegulation	Not strict enough	0.0658	0.9342
	Appropriate	0.0776	0.9224
	Too strict	0.2712	0.7288
	Don't know	0.1242	0.8758
MediaAI	No	0.1773	0.8227
	Yes	0.0803	0.9197
VoteIntent	Left	0.0559	0.9441
	Right	0.1236	0.8764
	Other	0.1606	0.8394

Table 6. Conditional probabilities for EUAppropriateRegulation, grouped by evidence. Baseline shown in the first row.

Evidence variable	Evidence value	Not strict enough	Appropriate	Too strict	Don't know
Baseline		0.3180	0.4892	0.1143	0.0785
AI EasierLife	Strongly disagree	0.3857	0.4307	0.1037	0.0799
	Somewhat disagree	0.3340	0.4804	0.1081	0.0775
	Somewhat agree	0.2687	0.5316	0.1222	0.0776
	Strongly agree	0.2570	0.5313	0.1318	0.0799
AI Uncontrollable	Strongly disagree	0.1652	0.5456	0.1960	0.0932
	Somewhat disagree	0.1397	0.6449	0.1397	0.0756
	Somewhat agree	0.2785	0.5402	0.1087	0.0725
	Strongly agree	0.4780	0.3528	0.0878	0.0815

Network structure and dependencies. In the *risk-oriented* network (Figure 3), vote intention (*VoteIntent*) directly influences both regulatory awareness (*HeardEURegulation*) and support for legal requirements (*AIRegulations*), as well as media exposure (*MediaAI*), confirming its central role in shaping both institutional opinions and information behavior. The link from perceived adequacy of the EU regulation (*EUAppropriateRegulation*) to *AIRegulations* remains, while all exposure-related arcs into the regulation block, such as *MediaAI* → *AIRegulations* in the model over all respondents, disappear. Within the thematic risk layer, all five grouped variables form a tightly connected subgraph, with *AI_Risk_Jobs* acting as the main source node influencing the others either directly or via short chains. Demographic and exposure variables have limited direct access to the perception layer, with only *Age* and *FriendsAI* entering this block through *AI_Risk_Jobs* and *AI_Risk_Data*, respectively.

In the *opportunity-oriented* network (Figure 4), vote intention still affects *MediaAI*, but no longer connects directly to regulatory support. Instead, support for AI

regulation (*AIRegulations*) depends exclusively on *EUAppropriateRegulation*, suggesting that among optimistic respondents, support for legal constraints is shaped more by evaluations of policy content than by political alignment. The grouped benefit themes form a dense, layered structure with *AI_Pos_Health* and *AI_Pos_Life* acting as upstream nodes. These flow into *AI_Pos_Work* and *AI_Pos_Tech*, which in turn feed into the most general category (*AI_Pos_General*). In contrast to the risk network, the thematic benefit nodes receive several direct inputs from demographic variables: *Sex* enters *AI_Pos_Health*, and *Age* influences *AI_Pos_Work*. Exposure and engagement variables such as *InterestAI*, *InformedAI*, *FriendsAI*, and *SearchAI* are again structured hierarchically but do not directly connect to the perception nodes. The paths into regulation mirror a technocratic progression: *InterestAI* drives awareness (*HeardEURegulation*) via standard exposure channels, while *EUAppropriateRegulation* alone governs support.

Together, these two networks reveal how overall framing (risk vs. opportunity) restructures the dependencies among background factors, exposure, and regulation, as well as the

Table 7. Conditional probabilities for AIUncontrollable, grouped by evidence. Baseline shown in the first row.

Evidence variable	Evidence value	Strongly disagree	Somewhat disagree	Somewhat agree	Strongly agree
Baseline		0.0804	0.2214	0.3003	0.3979
AI EasierLife	Strongly disagree	0.0624	0.0952	0.2033	0.6390
	Somewhat disagree	0.0278	0.1890	0.3576	0.4256
	Somewhat agree	0.0945	0.3238	0.3521	0.2296
	Strongly agree	0.2056	0.3221	0.2389	0.2334
AI FalseInfo	Strongly disagree	0.3517	0.2927	0.1680	0.1876
	Somewhat disagree	0.1447	0.3405	0.2923	0.2226
	Somewhat agree	0.0672	0.2486	0.3896	0.2946
	Strongly agree	0.0419	0.1504	0.2498	0.5579
AI vsHuman	Strongly disagree	0.1551	0.2010	0.2013	0.4425
	Somewhat disagree	0.1344	0.3340	0.2712	0.2604
	Somewhat agree	0.0661	0.2818	0.3898	0.2624
	Strongly agree	0.0746	0.1675	0.2584	0.4995
EUAppropriateRegulation	Appropriate	0.0897	0.2919	0.3316	0.2869
	Too strict	0.1379	0.2708	0.2858	0.3056
	Not strict enough	0.0418	0.0973	0.2630	0.5979
	Don't know	0.0955	0.2135	0.2777	0.4133

Table 8. Definition of the ten scenario profiles used in multi-variable evidence analysis.

Scenario name	Definition (key variables)
Baseline (no evidence)	No evidence provided; full marginal uncertainty retained
Young informed left	Age: 14–29; education: High; vote: Left; InterestAI: Very strongly; InformedAI: Very good; all exposure variables: Yes
Older uninformed right	Age: 60+; education: Low; vote: Right; InterestAI: Less strongly; InformedAI: Very poor; all exposure variables: No.
Middle educated moderate	Age: 45–59; education: Medium; vote: Other; InterestAI: Less strongly; InformedAI: Rather poor; MediaAI: Yes; FriendsAI, SearchAI: No.
Highly interested low info	Age: 30–44; education: Medium; vote: Left; InterestAI: Very strongly; InformedAI: Very poor; all exposure variables: Yes
High info low interest	Age: 30–44; education: High; vote: Right; InterestAI: Not at all; InformedAI: Very good; MediaAI: Yes; FriendsAI, SearchAI: No.
Urban conservative	Age: 30–44; education: Medium; vote: Right; InterestAI: Strongly; InformedAI: Rather good; all exposure variables: Yes
Rural progressive	Age: 45–59; education: High; vote: Left; InterestAI: Strongly; InformedAI: Rather good; MediaAI, FriendsAI: Yes; SearchAI: No.
Young apathetic	Age: 14–29; education: Low; vote: Other; InterestAI: Not at all; InformedAI: Very poor; all exposure variables: No.
Engaged moderate	Age: 30–44; education: Medium; vote: Other; InterestAI: Strongly; InformedAI: Rather good; all exposure variables: Yes
Disengaged elder	Age: 60+; education: Medium; vote: Other; InterestAI: Not at all; InformedAI: Rather poor; all exposure variables: No.

internal logic of belief formation. While the risk model is politically activated and thematically entangled, the opportunity one is demographically modulated and structurally layered.(Figures 5 and 6)

Variance decomposition via Sobol indices. Table 9 reports the first-order Sobol indices within the risk and opportunity subnetworks, focusing again on the three regulation-related outcomes and one output from each group of specific constructs: *AI_Risk_Jobs* for the risk subnetwork and *AI_Pos_Health* for the opportunity subnetwork. Outputs that were disconnected from their corresponding subnetwork, such as *AIRegulations* and *EUAppropriateRegulation* in the opportunity case, were not included since the Sobol indices are exactly zero for all inputs.

For *HeardEURegulation*, we observe a clear contrast between the two subnetworks. In the risk subnetwork, the

dominant contributors are *InterestAI* (65.7%), *SearchAI* (14.6%), and *VoteIntent* (26.0%), closely aligned with the results from the all respondents network where *InterestAI* and *SearchAI* also rank among the most influential. In the opportunity subnetwork, the variance is explained by exposure-related variables: *InterestAI*, *SearchAI*, *FriendsAI*, and *InformedAI*, with no meaningful contribution from demographics or *MediaAI*.

For *AIRegulations*, the most important factor in the risk subnetwork is *EUAppropriateRegulation* (57.1%), followed by *VoteIntent* (21.4%). This replicates the overall network structure, where *EUAppropriateRegulation* explains 32.9% of variance in *AIRegulations*, and *VoteIntent* accounts for 13.4%. The larger contribution observed here may result from isolating the pathway without interference from other, unrelated nodes. Notably, variables such as *InformedAI* and *FriendsAI* play a minimal role, suggesting that opinions on

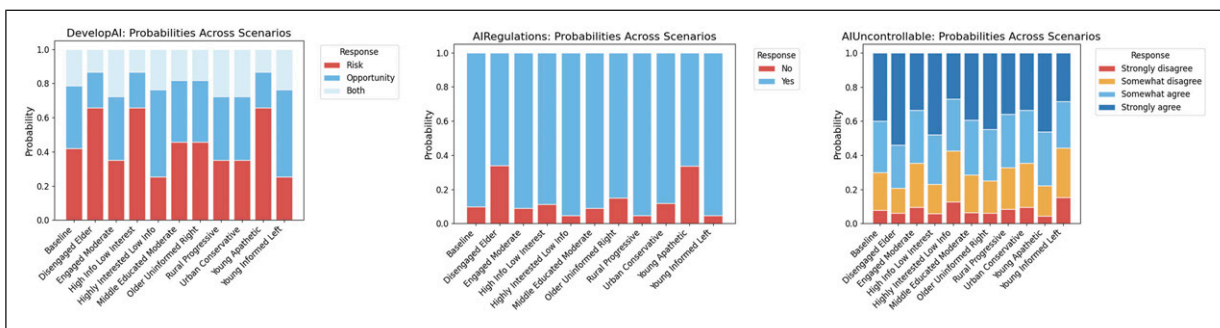


Figure 2. Posterior distributions across the ten scenarios for three key target variables: *DevelopAI*, *AIRegulations*, and *AIUncontrollable*.

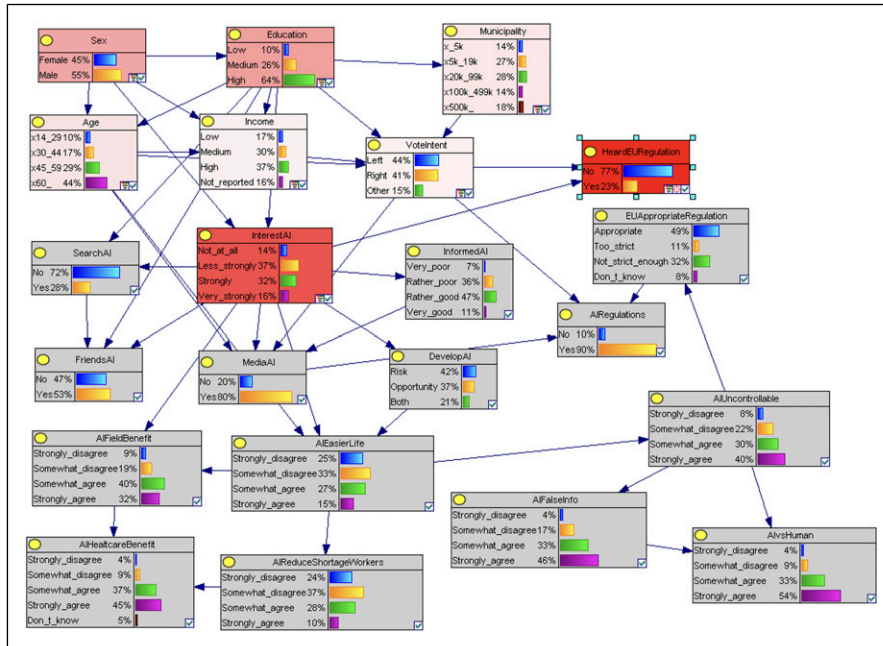


Figure 5. Bayesian network over all respondents with nodes shaded by their parametric sensitivity to the variable *HeardEURegulation*. Gray nodes have no measurable effect, while nodes shaded in red indicate increasing levels of influence, with darker shades corresponding to stronger effects.

AI regulation are more strongly tied to institutional and political cues than to general AI familiarity or sentiment.

For *EUAppropriateRegulation*, *Age* dominates the variance in the risk subnetwork (85.1%), a finding not visible in the model over the whole population where *AIEasierLife* and *AIUncontrollable* drive most of the variation. This difference highlights how the removal of highly influential nodes outside the regulation construct can expose underlying demographic patterns that were previously masked. The influence of *InformedAI* and *AIRegulations* remains modest, consistent with their lower Sobol scores in the previous analysis.

Turning to the specific outputs, no variable explains a large share of the variance. The highest contribution is from *AI_Risk_Misuse_Regulation* (10.6%) to *AI_Risk_Jobs*, indicating that no single input dominates the variance in these thematic outcomes.

Conditional probabilities and belief dynamics. We replicate the conditional probability analysis for the three regulation-related outputs in the subnetworks, focusing exclusively on demographic predictors.

For *HeardEURegulation*, the results from the subnetwork confirm a significant political gradient: respondents who identify as left-leaning are more likely to have heard about EU regulation (26.8%) compared to the baseline (19.6%), while right-wing respondents fall slightly below the baseline (15.7%), and those selecting “Other” display the lowest awareness (11.6%). These patterns align with those reported in Table 4, although the magnitude of the contrast is slightly more pronounced in the subnetwork.

Notably, this confirms the stable role of political orientation in shaping regulatory awareness, even when isolating regulation-specific constructs.

For *AIRegulations*, we again observe that left-leaning respondents express stronger support for AI regulation (95.4% compared to a baseline of 90.5%), whereas right-leaning and “Other” respondents express slightly reduced support (87.2% and 86.7%, respectively). The direction and magnitude of these effects are highly consistent with the previous results, reinforcing the interpretation that political stance is a reliable predictor of regulatory preferences.

In the case of *EUAppropriateRegulation*, age stratification reveals substantial heterogeneity. Respondents aged 14–29 are significantly more likely to consider current regulation “Appropriate” (70.7% versus a baseline of 40.9%), with a corresponding drop in the “Not strict enough” category. Middle-aged groups (30–59) align more closely with the baseline, while older respondents (60+) show a slightly lower probability of selecting “Appropriate” and a marginal increase in “Don’t know”. These age-related shifts were largely absent over the whole population, where attitudes were primarily shaped by optimism and concern variables.

Discussion

Patterns in public attitudes toward AI

Our findings reveal that how individuals frame artificial intelligence (as primarily a risk, an opportunity, or both) is




Figure 6. Tornado diagram for the variable *HeardEURegulation = No*, showing the effect of changing parameters of upstream nodes. Each bar represents a conditional configuration (e.g., joint levels of *InterestAI* and *VoteIntent*), and its length corresponds to the shift in probability under that scenario. Colors indicate the direction of parameter change: green for increases in the input variable, and red for decreases.

strongly shaped by perceived usefulness and personal engagement. This is broadly consistent with earlier work showing that younger, more educated, and higher-trust individuals tend to express greater techno-optimism and openness to AI applications.^{3,4,10} The belief that AI will make life easier emerged as the most influential driver of opportunity framing, reinforcing previous research that links techno-optimism to age and education.^{4,10} At the same time, fear that AI may become uncontrollable is associated with heightened concern about misinformation and insufficient regulation, echoing work on techno-anxiety and institutional mistrust.^{1,11} Importantly, these attitudes are not mutually exclusive: many respondents endorse both hopeful and fearful views, supporting recent claims that AI perceptions are multi-dimensional rather than polarized.^{3,14} This pattern is especially salient in our German sample, where institutional trust and policy awareness are relatively high. While the specific role of political orientation may vary across

national contexts, depending on partisan cues and regulatory discourse, the central influence of perceived usefulness and information engagement on AI framing is likely to generalize across countries, as these mechanisms are consistently linked to technology acceptance in the literature.¹⁰

Rather than ideology exerting a direct effect on regulatory support, as assumed in some accounts of politically motivated reasoning, our findings align with recent research showing that trust in AI governance is shaped by how individuals evaluate the adequacy, transparency, and reliability of institutional frameworks.³⁹ The variable EU Appropriate Regulation consistently ranked among the most influential predictors of support for legal requirements, suggesting that citizens do not simply support or oppose regulation based on partisan cues, but critically assess whether current efforts go far enough. This finding refines the picture presented by earlier work on regulatory preferences^{2,7}, indicating that perceived institutional

Table 9. Variance-based Sobol indices (percentage of output variance explained) for each input variable across target variables in the Risk and Opportunity subnetworks. Bold values indicate variables that explain more than 10% of output variance. Double dashes (–) indicate that the variable was the target output in that column. NA indicates that the variable was not included in the corresponding BN submodel.

Input variable	Risk subnetwork					Opportunity subnetwork				
	EUAppropriateRegulation					AI_Regulations				
	EUAppropriateRegulation	AI_Regulations	HeardEURRegulation	AI_Risk_Jobs	AI_Pos_Health	EUAppropriateRegulation	AI_Regulations	HeardEURRegulation	AI_Risk_Jobs	AI_Pos_Health
Age	85.1%	0.2%	0.2%	7.6%	0.1%	85.1%	0.2%	0.2%	7.6%	0.1%
Education	2.6%	0.0%	4.9%	0.2%	0.0%	2.6%	4.9%	0.2%	0.2%	0.0%
Income	0.0%	0.4%	0.7%	0.0%	0.1%	0.0%	0.7%	0.3%	0.0%	0.1%
Municipality	1.1%	0.6%	1.2%	0.1%	0.0%	1.1%	1.2%	0.1%	0.1%	0.0%
Sex	0.5%	0.0%	0.0%	0.1%	7.0%	0.5%	0.0%	0.0%	0.1%	7.0%
VotIntent	0.0%	21.4%	26.0%	0.0%	0.1%	0.0%	26.0%	0.0%	0.0%	0.1%
FriendsAI	2.3%	0.0%	9.5%	0.2%	0.0%	2.3%	9.5%	13.4%	0.2%	0.0%
InformedAI	1.7%	0.0%	8.9%	0.0%	0.0%	1.7%	8.9%	16.9%	0.0%	0.2%
InterestAI	0.2%	0.0%	65.7%	0.0%	0.0%	0.2%	65.7%	100.0%	0.0%	0.0%
MediaAI	0.0%	0.4%	5.4%	0.0%	0.0%	0.0%	5.4%	0.4%	0.0%	0.0%
SearchAI	0.2%	0.0%	14.6%	0.0%	0.0%	0.2%	14.6%	21.7%	0.0%	0.0%
AI_Regulations	1.0%	–	0.5%	0.0%	0.0%	1.0%	0.5%	0.0%	0.0%	0.0%
EUAppropriateRegulation	–	57.1%	0.0%	0.2%	0.0%	–	0.0%	0.0%	0.2%	0.0%
HeardEURRegulation	0.0%	0.5%	–	0.0%	0.0%	0.0%	–	–	0.0%	0.0%
AI_Risk_Jobs	1.9%	0.0%	0.0%	–	0.0%	1.9%	0.0%	–	–	NA
AI_Risk_Misuse_Regulation	0.1%	0.0%	0.0%	10.6%	NA	0.1%	0.0%	NA	10.6%	NA
AI_Risk_LossOfControl	0.1%	0.0%	0.0%	9.4%	NA	0.1%	0.0%	NA	9.4%	NA
AI_Risk_Society	0.1%	0.0%	0.0%	6.0%	NA	0.1%	0.0%	NA	6.0%	NA
AI_Risk_Data	0.0%	0.0%	0.1%	3.2%	NA	0.0%	0.1%	NA	3.2%	NA
AI_Pos_Work	NA	NA	NA	NA	NA	NA	NA	0.0%	NA	2.8%
AI_Pos_Tech	NA	NA	NA	NA	NA	NA	NA	0.0%	NA	7.0%
AI_Pos_General	NA	NA	NA	NA	NA	NA	NA	0.0%	NA	5.5%
AI_Pos_Life	NA	NA	NA	NA	NA	NA	NA	0.0%	NA	1.9%
AI_Pos_Health	NA	NA	NA	NA	NA	NA	NA	0.0%	NA	1.9%

performance, rather than exposure alone, drives willingness to endorse governance interventions. While prior studies acknowledge the role of trust and policy content, our BN approach makes the mediating structure explicit, revealing how evaluations of regulatory adequacy channel the effect of political orientation.

Engagement with information sources, particularly self-reported interest in AI and active search behavior, played a central role in shaping awareness of EU regulation. Respondents who reported high interest were significantly more likely to have heard of regulatory initiatives, while levels of informedness and interpersonal discussion also contributed. However, these factors did not directly translate into policy support. This confirms findings from prior studies^{3,14} that awareness is driven by cognitive engagement and information exposure, while support for regulation rests on deeper evaluative and emotional layers, such as trust, perceived risk, and institutional credibility.

The comparison between subpopulation networks, those viewing AI as a risk vs. an opportunity, offers further insight into belief structures. Among risk-oriented respondents, regulatory attitudes are more politically activated, and thematic concerns are tightly interconnected. In this group, political orientation plays a prominent role, consistent with findings from polarized and lower-trust contexts such as Southern Europe and Latin America.^{9,12} In contrast, opportunity-oriented respondents display a more demographically driven pattern, where support for regulation flows through assessments of policy content rather than ideology. This suggests that logics of belief formation differ across framing groups: risk attitudes appear more ideologically and emotionally embedded, while opportunity attitudes align more closely with technocratic reasoning.

Engaging the public in governance

These findings have important implications for the design of AI governance and public engagement strategies. Scenario-based inference revealed strong attitudinal contrasts between population segments, such as the *Young Informed Left* and the *Disengaged Elder*. Such contrasts suggest that a one-size-fits-all communication strategy is unlikely to succeed. Instead, engagement efforts should be tailored: emphasizing practical benefits and accessibility for disengaged groups, while focusing on procedural adequacy and institutional transparency for politically skeptical individuals. The observed alignment between interest, exposure, and awareness also implies that public education campaigns may be most effective when targeting interest-driven audiences who already engage with technology-related content.

The analysis further suggests that concerns about uncontrollability and policy insufficiency are central to regulatory dissatisfaction, aligning with studies that link fear of AI to weakened trust in governance.^{1,2} This highlights the need for risk communication strategies that go beyond

factual reassurances and address perceived control, fairness, and legitimacy. Policymakers should recognize that public unease often reflects not just fear of the technology itself, but a broader anxiety about who governs its use and how.

In light of the EU AI Act's phased implementation, our results underscore the conditional nature of public support. Even individuals who express general agreement with regulation adjust their attitudes based on how adequate current efforts appear. This supports critiques of the Act's technocratic framing,^{5,7} which may fail to resonate with public expectations around fairness, autonomy, and accountability. Making regulatory content accessible and interpretable to the public may be just as important as enforcing compliance or technical standards.

Finally, the modeling framework introduced here provides a practical tool for policymakers seeking to monitor and respond to public opinion. By identifying which population segments are most responsive to particular aspects of regulation (such as awareness campaigns, policy framing, or risk communication) our BN approach enables a more targeted and adaptive form of governance. Future extensions could integrate this model into participatory policymaking platforms or media monitoring systems to support real-time adjustment of engagement strategies.

Contributions, limitations, and paths forward

Methodologically, this study demonstrates the value of probabilistic graphical models for analyzing public attitudes toward emerging technologies. BNs offer a transparent and flexible approach for modeling belief systems, capturing not only direct associations but also conditional dependencies and mediated effects. Unlike traditional regression or structural equation models, BNs accommodate complex interrelations among variables without assuming linearity or unidirectional causality. The use of bootstrapped structure learning, Sobol variance decomposition, and scenario-based inference provides both robustness and interpretability, aligning with recent calls for more transparent and exploratory modeling in high-stakes policy domains.²²

Nonetheless, the study has several limitations. First, the analysis is based on cross-sectional data, which limits our ability to draw causal conclusions. While the tiered structure and blacklists reflect plausible causal ordering, they cannot establish temporal precedence. Second, although the dataset is rich, it pertains to a single national context and results may not generalize to other countries without adjustment for institutional and cultural factors. For instance, political alignment effects are likely to be more context-specific, shaped by national party systems and trust in institutions, whereas the influence of interest in AI, perceived usefulness, and information-seeking behavior may reflect more universal patterns in technology acceptance and belief formation. Third, as the network grows in size and complexity, interpretability may diminish. Future studies could explore

the use of regularization techniques, sparse priors, or expert elicitation to constrain structure learning and improve scalability.

Several avenues for future research emerge. Longitudinal or repeated cross-sectional data could be used to study belief dynamics over time, particularly in response to policy changes or media events. Comparative studies across different EU countries, or between European and non-European contexts, could reveal how institutional design and public trust interact to shape attitudes. Finally, integrating BNs with qualitative data, such as open-ended survey responses or focus group transcripts, could deepen understanding of the cultural narratives and personal experiences that underpin AI-related beliefs. Such extensions would help further bridge the gap between interpretability, policy relevance, and methodological rigor in the study of technology and society.

Conclusions

This study presents the first application of a BN to analyze public attitudes toward AI regulation using nationally representative European survey data. By modeling political, psychological, and informational factors within a unified BN framework, we uncover how beliefs co-vary and condition one another, revealing distinct pathways that shape awareness, trust, and support for legal oversight. Our findings emphasize that regulatory attitudes are shaped not only by ideology or exposure but by perceived adequacy and emotional responses to risk. The model's interpretability enables targeted insights for policymakers, offering a transparent tool for aligning AI governance with public sentiment. Future work can extend this approach across countries and over time to monitor evolving attitudes and inform adaptive regulation in an increasingly AI-integrated society.

ORCID iD

Manuele Leonelli  <https://orcid.org/0000-0002-2562-5192>

Funding

The authors received no financial support for the research, authorship, and/or publication of this article.

Declaration of conflicting interests

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

References

- Li J and Huang JS. Dimensions of artificial intelligence anxiety based on the integrated fear acquisition theory. *Technol Soc* 2020; 63: 101410.
- Chiarini A, Grando A, Venturini S, et al. Do automation and AI impact on job reduction? A study on perceived risk of losing job among white-collars in the Italian manufacturing companies. *Prod Plann Control* 2024; 35(16): 2198–2211.
- Montag C, Kraus J, Baumann M, et al. The propensity to trust in (automated) technology mediates the links between technology self-efficacy and fear and acceptance of artificial intelligence. *Computers in Human Behavior Reports* 2023; 11: 100315.
- Department for Science, Innovation and Technology. Public attitudes to data and AI. *Tracker survey* 2024; wave 4.
- Kop M. EU Artificial intelligence act: the European approach to AI, 2021.
- Edwards L. Expert explainer: the EU AI act. *Technical report, European Parliamentary Research Service (EPRS)* 2022, PE 698792.
- Laux J, Wachter S and Mittelstadt B. Trustworthy artificial intelligence and the European Union AI act: on the conflation of trustworthiness and acceptability of risk. *Regul Gov* 2024; 18(1): 3–32.
- Magistro B, Borwein S, Alvarez RM, et al. Attitudes toward artificial intelligence (AI) and globalization: common microfoundations and political implications. *Am J Polit Sci* 2025.
- Gur T, Hameiri B and Maaravi Y. Political ideology shapes support for the use of AI in policy-making. *Front Artif Intell* 2024; 7: 1447171.
- Bergdahl J, Latikka R, Celuch M, et al. Self-determination and attitudes toward artificial intelligence: cross-national and longitudinal perspectives. *Telematics Inf* 2023; 82: 102013.
- Dong M, Conway JR, Bonnefon JF, et al. Fears about artificial intelligence across 20 countries and six domains of application. *Am Psychol* 2024.
- Yarovenko H, Kuzior A, Norek T, et al. The future of artificial intelligence: fear, hope or indifference? *Human Technology* 2024; 20(3): 611–639.
- Cremaschi A, Lee DJ and Leonelli M. Will AI take my job? Evolving perceptions of automation and labor risk in Latin America. *arXiv preprint arXiv:250508841* 2025.
- Schiavo G, Businaro S and Zancanaro M. Comprehension, apprehension, and acceptance: understanding the influence of literacy and anxiety on acceptance of artificial intelligence. *Technol Soc* 2024; 77: 102537.
- Pearl J. *Probabilistic reasoning in intelligent systems; network of plausible inference*. Morgan Kaufmann, 1988.
- Berisha B, Mëziu E and Shabani I. Big data analytics in cloud computing: an overview. *J Cloud Comput* 2022; 11(1): 24.
- Ho GTS, Tsang YP, Wu Q, et al. ck-FARM: an R package to discover big data associations for business intelligence. *SoftwareX* 2023; 22: 101341.
- Yang S, Krause NM, Bao L, et al. In AI we trust: the interplay of media use, political ideology, and trust in shaping emerging AI attitudes. *Journal Mass Commun Q* 2025; 102(2): 382–406.
- Wilde N. Fear of artificial intelligence or fear of looking in the mirror? Revisiting the Western machine-takeover imaginary. *AI Soc* 2025: 1–11.

20. Imai K, Keele L and Tingley D. A general approach to causal mediation analysis. *Psychol Methods* 2010; 15(4): 309–334.
21. Glynn AN. The product and difference fallacies for indirect effects. *Am J Polit Sci* 2012; 56(1): 257–269.
22. Rudin C. Stop explaining Black box machine learning models for high stakes decisions and use interpretable models instead. *Nat Mach Intell* 2019; 1(5): 206–215.
23. Kaikkonen L, Parviainen T, Rahikainen M, et al. Bayesian networks in environmental risk assessment: a review. *Integrated Environ Assess Manag* 2020; 17(1): 62–78.
24. Carrick J, Fitzsimmons C and Gray T. Using Bayesian belief networks and participatory action research to improve stakeholder engagement. *Transdisciplinary Marine Research*. Routledge, 2022, pp. 85–106.
25. Pérez S, German-Labaume C, Mathiot S, et al. Using Bayesian networks for environmental health risk assessment. *Environ Res* 2022; 204: 112059.
26. Maciejowski J and Leonelli M. Uncovering drivers of EU carbon futures with Bayesian networks. *arXiv preprint arXiv: 250510384* 2025.
27. Cugnata F, Kenett RS and Salini S. Bayesian networks in survey data: robustness and sensitivity issues. *J Qual Technol* 2016; 48(3): 253–264.
28. Presse- und Informationsamt der Bundesregierung. Current questions on AI 2023. <https://doi.org/10.4232/1.14223,2024.ZA8719Data>
29. Koller D and Friedman N. *Probabilistic graphical models: principles and techniques*. MIT Press, 2009.
30. Tsamardinos I, Brown LE and Aliferis CF. The max-min hill-climbing Bayesian network structure learning algorithm. *Mach Learn* 2006; 65: 31–78.
31. Scutari M. Learning Bayesian networks with the bnlearn R package. *J Stat Software* 2010; 35: 1–22.
32. Scutari M and Nagarajan R. Identifying significant edges in graphical models of molecular networks. *Artif Intell Med* 2013; 57(3): 207–217.
33. Heckerman D, Geiger D and Chickering DM. Learning Bayesian networks: the combination of knowledge and statistical data. *Mach Learn* 1995; 20: 197–243.
34. Bishop CM. *Pattern recognition and machine learning*. ISBN: Springer, 2006.
35. Leonelli M. bnRep: a repository of Bayesian networks from the academic literature. *Neurocomputing* 2025; 624: 129502.
36. Ballester-Ripoll R and Leonelli M. Computing Sobol indices in probabilistic graphical models. *Reliab Eng Syst Saf* 2022; 225: 108573.
37. Ballester-Ripoll R and Leonelli M. You only derive once (YODO): automatic differentiation for efficient sensitivity analysis in Bayesian networks. *International Conference on Probabilistic Graphical Models*. PMLR, pp. 169–180.
38. Leonelli M, Ramanathan R and Wilkerson RL. Sensitivity and robustness analysis in Bayesian networks with the bnmonitor R package. *Knowl Base Syst* 2023; 278: 110882.
39. Tamò-Larrieux A, Guitton C, Mayer S, et al. Regulating for trust: can law establish trust in artificial intelligence? *Regulation & Governance* 2024; 18(3): 780–801.