

Do Citizens Agree with the EU AI Act? Public Perspectives on Risk and Regulation of AI Systems

Gabriel Lima

Max Planck Institute for Security and Privacy
Bochum, Germany
gabriel.lima@mpi-sp.org

Gustavo Gil Gasiola

Department of Informatics
Karlsruhe Institute of Technology
Karlsruhe, Germany
gustavo.gasiola@kit.edu

Frederike Zufall

Department of Informatics
Karlsruhe Institute of Technology
Karlsruhe, Germany
Waseda Institute for Advanced Study
Waseda University
Tokyo, Japan
zufall@kit.edu

Yixin Zou

Max Planck Institute for Security and Privacy
Bochum, Germany
yixin.zou@mpi-sp.org

Abstract

The European Union (EU) has spearheaded the regulation of artificial intelligence (AI) with the AI Act, which regulates AI systems based on the risks they pose to fundamental rights and other protected values. AI systems that pose unacceptable risks are prohibited, high-risk AI systems must comply with mandatory requirements, and minimal risk AI systems are encouraged—but not required—to adopt voluntary standards. Motivated by concerns that the AI Act may not reflect the public’s opinions, we investigate how laypeople ($N=1,421$) assess 48 different AI systems concerning their risk and regulation. We find that people believe all 48 AI systems pose moderate levels of risk and should be regulated (albeit without outright prohibitions). Our findings challenge the AI Act’s tiered approach, showing that people might support horizontal regulation requiring minimal standards for AI systems, and provide implications for developers seeking to develop AI aligned with public expectations.

CCS Concepts

• **Applied computing** → Law; Psychology; • **Social and professional topics** → Governmental regulations; • **Human-centered computing** → Empirical studies in HCI.

Keywords

Artificial Intelligence, AI System, Regulation, European Union, EU, AI Act, Risk, Regulation, Public, Laypeople

ACM Reference Format:

Gabriel Lima, Gustavo Gil Gasiola, Frederike Zufall, and Yixin Zou. 2026. Do Citizens Agree with the EU AI Act? Public Perspectives on Risk and Regulation of AI Systems. In *Proceedings of the 2026 CHI Conference on*

Human Factors in Computing Systems (CHI '26), April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 24 pages. <https://doi.org/10.1145/3772318.3790535>

1 Introduction

Artificial intelligence (AI) systems have become embedded in numerous fields, from medicine and education to employment and law enforcement. Their widespread deployment in consequential domains has prompted governments from around the world—such as the United States (US) [41], China [72], Brazil [4], and South Korea [69]—to examine whether and how to regulate AI systems in ways that balance their potential risks and benefits.

The European Union (EU) has become a global leader in AI regulation with the introduction of the first comprehensive regulatory framework for AI: the AI Act. By recognizing that AI systems can pose risks to fundamental rights, health, safety, democracy, the rule of law, the environment, and other protected values (Art. 1(1) AI Act), the AI Act seeks to strike a balance between AI innovation and the protection of fundamental rights by imposing requirements on AI systems based on the risk they pose.

The AI Act categorizes AI systems according to a three-tier risk framework. AI systems deemed to pose *unacceptable risk* are prohibited by the AI Act as their excessive risks cannot be reduced to an acceptable level by mitigation strategies. *High-risk* AI systems pose significant risks that can be reduced to an acceptable level through compliance with mandatory requirements prescribed by the AI Act. AI systems that do not pose a significant risk to fundamental rights and other protected values fall into the residual category of *minimal risk* and are thus merely encouraged to voluntarily comply with the AI Act’s risk mitigation measures.

Although the AI Act has been applauded for employing public consultations during its drafting process, there remains the question of whether the AI Act—as it has come into effect—is aligned with the opinions of those who will be subjected to the risks posed by AI systems: the general public [18]. In fact, scholars have argued that the final version of the AI Act is more aligned with the interests of corporate actors than the views of the general public [86].



This work is licensed under a Creative Commons Attribution 4.0 International License. *CHI '26, Barcelona, Spain*

© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2278-3/26/04
<https://doi.org/10.1145/3772318.3790535>

For instance, the AI Act introduces mandatory requirements only for high-risk AI systems; according to the consultation report that informed the regulation, 55% of industry and business responses agreed with this approach, whereas only 39% of surveyed citizens endorsed this proposal [25]. This misalignment is particularly worrying given that the classification of AI systems into the high-risk level relies on self-assessments by the companies to determine the existence of significant risks (Art. 6(3) AI Act).

Our research exists side-by-side with recent efforts in the EU to operationalize the AI Act. The EU Commission has recently proposed a regulation seeking to simplify the implementation of the AI Act (Digital Omnibus on AI)¹ that leaves the general risk-based approach unchanged, but underlines the fact that the debate and regulation of AI is an ongoing process. At the same time, we note that the requirements put forward by the AI Act are currently being concretized through implementing regulations, guidelines, or codes of conduct, which could benefit from insights from the general public [11, 84, 86].

1.1 Findings and Contributions

In this paper, we take a step towards collecting the public's opinion of AI systems in light of the AI Act. Instead of exploring people's opinions about AI regulation in a general sense, we design a study to investigate whether the general public's expectations of AI regulation are aligned with how the AI Act operates. Here, we present the results of this study, which collects laypeople's ($N=1,421$) perceptions of 48 different AI systems in relation to their perceived risk and potential regulation. We then contrast these lay opinions with how the AI Act classifies and regulates the same AI systems. By recruiting participants from three EU countries (Germany, Spain, and France) and the United States (US), our study investigates the extent to which their opinions about AI regulation are aligned—or misaligned—with the AI Act's risk-based approach. More specifically, we focus on three central themes of the AI Act:

- (1) We examine the extent to which laypeople believe AI systems pose risks to fundamental rights and protected values. By doing so, we aim to inform potential amendments to the list of high-risk AI systems (Art. 7 AI Act) and shed light on potential reasons behind user acceptance (or rejection) of AI systems.
- (2) We investigate which AI systems people believe should be regulated and prohibited, providing implications for future attempts to regulate AI beyond the EU and informing the development of future AI systems. Our findings have concrete ramifications for potential future legislative and regulatory amendments and for policymakers seeking to develop policies that are aligned with the public's perception of AI systems [5, 83].
- (3) Finally, we explore the perceived importance of several requirements put forward by the AI Act. Our findings provide implications for AI developers aiming to develop systems that are perceived to be safe and legal [27, 52, 54] by identifying what laypeople expect from AI systems in terms of both voluntary and mandatory requirements.

Our results indicate that participants believe all 48 AI systems—no matter their risk level according to the AI Act—pose moderate levels of risk and should thus be regulated. Yet, participants were largely contrary to outright prohibitions, instead calling for broader regulations that prescribe the AI Act's regulatory requirements to all AI systems. We also show that participants across all four countries largely agreed that all AI systems are risky and should be subjected to strict requirements.

Our findings have implications for HCI researchers, policymakers, and AI developers. Laypeople's assessments that all AI systems pose similar levels of risk suggest that they may not differentiate between the varying levels of risk posed by distinct AI systems, demonstrating the importance of HCI research seeking to advance AI literacy and design AI in ways that end users are able to better assess potential risks [39, 45, 58, 90]. Our findings also challenge the risk-based tiered framework adopted by the AI Act, instead providing support for horizontal rules that set minimal standards for *all* AI systems [7, 31, 79, 84]. Our results also demonstrate that users may expect AI developers to comply with the AI Act's requirements even for minimal risk AI, for which they are not mandatory. Finally, our study provides initial public support for internationally applicable and legally binding rules for AI systems.

2 Background

2.1 The EU AI Act

The AI Act² is a pioneering comprehensive legal framework for AI systems and general-purpose AI (GPAI) models [84]. The regulation is based on the EU competence under Art. 114 TFEU,³ which allows for harmonized legislation to establish free movement of goods and services across the internal EU market, and based on the competence for rules on the protection of personal data (Art. 16(2) TFEU). However, on its substantive level, the subject matter and purpose of the AI Act are not limited to the improvement of the functioning of the internal market. The regulation also aims to promote the uptake of “human-centered and trustworthy artificial intelligence” (Art. 1(1) AI Act) [18] by building upon a study prepared by the High-Level Expert Group, which identified key elements for “trustworthy AI” aligned with ethical principles [67]. In this sense, the Act repeatedly refers to the protection of EU fundamental rights and values, explicitly mentioning democracy, the rule of law, and the environment (Art. 1(1) and Recitals 1, 6, 8, 28, 58, 59, 60 and 176 AI Act). Consequently, the AI Act introduces prohibitions, mandatory requirements, and other legislative measures to guarantee that AI systems are developed and deployed in line with fundamental rights and other important protected values. Considering its broad (extra-)territorial scope (Art. 2(1) AI Act), it also creates obligations for providers located outside the EU and could also serve as a role model for upcoming regulatory regimes beyond the EU—the so-called “Brussels effect” [2, 9, 33, 78]. That is not to say that the AI Act is immune to critiques. For instance, scholars have criticized the regulation for its overreliance on self-certification [86] and the difficulty in assessing the risks raised by

¹COM(2025) 836 final, Brussels, 19.11.2025.

²Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonized rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139, and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797, and (EU) 2020/1828.

³Treaty on the Functioning of the European Union, OJ C 326, 26.10.2012, pp. 47–390.

specific use cases of AI [13, 18, 23, 79], raising questions about its effectiveness and international influence [68].

2.2 The Risk-Based Approach

Finding the right balance between AI innovation and its free movement on the EU market on one side, and the protection of fundamental rights and protected values on the other, is the central regulatory task of the AI Act's legal framework. Under the principle of legislative proportionality [77], any limitation to the freedom to conduct business (Art. 16 of the Charter of Fundamental Rights of the European Union)—which includes the freedom to develop and implement AI systems—is only allowed if it is necessary to meet objectives of general interest or to protect the rights and freedoms of others (Art. 52(1) EU Charter). As the risk posed by AI systems varies widely according to their capabilities and context of application [82], the AI Act relies on a “clearly defined risk-based approach” (Recital 26 AI Act) to ensure the proportionality between legislative measures and potential threats to fundamental rights and protected values [16, 19, 28, 55].

The AI Act's risk-based approach comprises three levels of risk: unacceptable, high, and minimal. **Unacceptable AI systems** pose excessive risks that cannot be reduced to an acceptable level by risk mitigation measures and are thus prohibited. This excessive [10, 55] or unbearable [56] risk to fundamental rights and protected values justifies the prohibition of these AI practices (Art. 5(1) and Recital 26 AI Act) [13]. **High-risk AI systems** pose significant risks (Recitals 7 and 46 AI Act) that can be reduced to an acceptable level [19, 28, 74, 75]. Accordingly, the AI Act sets out mandatory requirements for such systems to ensure that their risks are properly addressed. AI systems that do not pose a significant risk to fundamental rights or protected values fall into the residual category of **minimal risk** [18]. For these systems, the AI Act merely encourages the adoption of codes of conduct and voluntary compliance with other requirements (Art. 95 AI Act).

Besides these three risk levels for AI systems, the AI Act also introduces a regulatory regime for GPAI models with specific rules tailored to potential systemic risks associated with them (see Art. 51-56 AI Act). In our study, we focus on the basic risk-based regime for AI systems that pose unacceptable, high, or minimal risk.

2.3 Public Participation in Drafting and Implementing the AI Act

Direct public participation in the EU's legislative process is an important step towards increasing democratic legitimacy [57] and ensuring that political decisions are more coherent, transparent (Art. 11(3) TEU⁴), credible, and trustworthy [15]. The right to participate in the democratic process (Article 10(3) TEU) is exercised through representative democracy (Article 10(1) TEU) and, complementarily, through the direct participation of citizens in decision-making processes (Article 11 TEU) [46]. Article 11(1) TEU also establishes a general obligation for EU institutions to ensure citizens can share their views on key policy decisions [57]. Specifically, Article 11(3) TEU requires the European Commission to carry out broad consultations with stakeholders.

As such, public consultations informed the AI Act's drafting procedure. Alongside workshops, technical consultations, and the use of a multi-stakeholder online platform [16], the Commission carried out a broad public consultation [25] from February to June 2020. This process comprised an online survey that received 406 responses from citizens (around 34 percent of all respondents) about the regulatory options for AI, including the classification of high-risk AI systems and mandatory requirements. From July to September 2020, the Commission also requested feedback on the Inception Impact Assessment, a more concrete description and justification of the Commission's decision, which received only seven contributions from citizens (around 5 percent of all respondents) [16].

After its adoption and entering into force, the practical implementation of the AI Act also requires extensive work by the Commission in developing guidelines (Art. 96 AI Act) and adopting delegated and implementing acts (Art. 41 and 97 AI Act), as well as by the European standardization organizations in drawing up standards (Art. 40 AI Act). While direct public participation is not explicitly mentioned in this phase, the Commission is conducting public consultations, such as those regarding the guidelines on prohibited practices [17] and the Code of Practices for GPAI models [29].

But to what extent does the risk-based regulatory framework adopted by the AI Act actually align with citizens' perceptions? As discussed above, the legislative process behind the AI Act has been praised by its numerous public consultations [18]. Yet, these public consultations suffer from significant limitations that could have limited the public's ability to propose different regulatory approaches. For instance, the survey questions employed by the previous public consultation [25] portray the proposed risk-based approach as the only alternative, priming respondents to answer positively (e.g., “Do you agree with the approach to determine ‘high-risk’ AI applications proposed in Section 5.B of the White Paper?”). The results of this consultation also illustrate misalignments between the general public, corporate actors, and the final version of the AI Act: the public was much less likely to agree that mandatory requirements should be limited to high-risk AI systems [25], as implemented by the AI Act.

Even though the public may have had some influence in the drafting process of the AI Act, it is unclear whether citizens will continue to have a voice in how the regulation will be applied. The requirements set by the AI Act are now being implemented through a set of implementing regulations, guidelines, and codes of conduct that are bound to suffer from corporate influence [84, 86]. It is unclear whether civil society will have an adequate representation in the development of these standards [11]. With these limitations in mind, this paper investigates the general public's perceptions of AI systems in relation to the AI Act's risk-based approach.

2.4 Public Opinion About AI Regulation

Research has examined people's attitudes towards AI systems through several lenses, such as fairness [35, 36, 48, 50, 53, 80, 81], trust [26, 26, 47, 50], and risk [3, 73, 76, 88]. More relevant to the current research, prior studies have collected the public's attitudes towards AI regulation through surveys [20] and interactive platforms [61]. For instance, a survey conducted in 17 countries focusing on AI

⁴Treaty on the European Union, OJ C 202 7.6.2016, p. 13.

applications in healthcare, security, human resources, and recommender systems identified that a majority of the public believes AI regulation is necessary, calling for external, independent oversight [32]. Another survey covering 11 countries also identified a strong consensus that “AI requires careful management” [21]. A study conducted in Germany and Spain shows a similar trend, with most participants supporting a stricter regulatory framework for AI [85].

Studies conducted by the Pew Research Center also found that people around the world are more concerned than excited about AI and believe higher safety standards are needed [70, 71]. In contrast, a study conducted by Google identified that the public is now more excited than concerned about AI [42]. Focusing on local United States (US) policymakers, Hatz et al. [37] found that they also support government oversight of AI, with increased attention to specific issues, such as privacy, unemployment, and fairness. As summarized in a review paper of English-language studies examining the public opinion about AI [20], support for external, independent AI governance largely outweighs opposition.

In this paper, we focus on the public’s opinions about AI systems and their regulation from the perspective of the EU AI Act. Instead of exploring whether people call for regulation in a broader sense, we rigorously investigate whether the public perception of AI systems are aligned with how the AI Act views the same AI systems. As such, our study focuses on core concepts of the AI Act: whether and to what extent AI systems pose risks to EU fundamental rights and other protected values and how these perceived risks translate into regulatory requirements.

3 Methods

We conducted a multi-country large-scale study ($N=1,421$) to capture laypeople’s perceptions of the risk posed by 48 different AI systems, as well as to collect their opinions about how each AI system should be regulated. Furthermore, we also examined the perceived importance of some of the requirements prescribed by the AI Act for high-risk AI systems. Our study was designed to strictly focus on the AI Act and the EU context, examining how this particular regulation is (mis)aligned with public opinion, rather than collecting the public’s attitudes towards AI regulation in a more general sense. The study was approved by an Ethical Review Board (ERB), and all of the data and scripts used for analysis are available online: <https://tinyurl.com/PublicPerspective-AIAct>.

3.1 List of AI Systems

Considering the broad scope of the AI Act (Art. 2(1) AI Act), we sought to collect the perceived risk of a wide range of AI systems. To do so, two authors with an HCI background first drafted a list of 158 applications that could potentially fall under the concept of AI systems under Art. 3(1) AI Act. The drafted list was inspired by AI systems discussed in the Ethics Guidelines for Trustworthy AI [67], the public consultation on the White Paper on AI [25], the Impact Assessment on the AI Act [16], as well as the AI applications mentioned on the Recitals, Art. 5, and Annex III of the AI Act. For AI systems that could potentially fall into the minimal risk category, the two authors also came up with applications based on their

personal and professional experiences (e.g., AI systems designed to help programmers code).

Two authors with a legal background then determined whether each AI application previously identified indeed fell under the definition of AI system (Art. 3(1) AI Act) and then classified each AI system following the approach proposed by Gasiola [30]. They first determined whether the AI system was included in the list of prohibited AI practices set out in Art. 5 AI Act. Prohibited AI systems were classified as *unacceptable risk*. If the AI system was not prohibited, they considered two consecutive criteria to further classify it. The first criterion evaluated whether the AI system was a safety component of a product, or whether the AI system itself was a product covered by EU product safety harmonization legislation that is subject to third-party conformity assessment (Art. 6(1) AI Act). The second criterion referred to whether the AI system was listed in the critical areas described in Annex III of the AI Act (Art. 6(2) AI Act). If at least one of these criteria was met, the AI system was classified as *high-risk*; otherwise, it was considered *minimal risk*. AI systems were removed from the list if it was not possible to clearly determine whether the classification criteria had been met.

All of the authors then met to collectively and iteratively refine the list and the descriptions, ensuring that the AI systems would clearly fall into one of the three risk levels, while at the same verifying that their description was as clear and concise as possible. We sought to have at least one AI system for each paragraph in Annex III AI Act and Art. 5(1) AI Act, also including some that would fall into the residual category of minimal risk.

The final list includes 48 AI systems, covering AI applications explicitly prohibited in Art. 5(1) AI Act and systems considered high-risk listed in Annex III AI Act [33], including also some AI systems that would fall into the minimal risk category. We present the full list and the mapping between each AI system and their risk level in Appendix A. Our list comprises a wide range of domains, from applications in law enforcement and the government to AI systems deployed at the workplace and educational institutions. We present some example AI systems in Table 1. Most AI systems ($n=26$) are categorized as high-risk by the AI Act (**high-risk AI systems**). We also included 10 AI systems that would be prohibited by the AI Act (**unacceptable risk AI systems**). Finally, we incorporated 12 AI systems that would neither be classified as unacceptable- nor high-risk (**minimal risk AI systems**). For consistency, the study introduced all AI systems to participants using the language employed in Annex III AI Act, which introduces high-risk AI systems as “AI systems intended to be used for [application in a particular critical area].”⁵

3.2 Participants

Given the AI Act’s objective to harmonize AI regulation across the EU, we aimed to recruit study participants from several EU countries to examine whether risk perceptions vary within the EU.

⁵The AI Act is sometimes explicit about the user of a particular AI system. For instance, Annex III 5(a) AI Act explicitly describes “AI systems intended to be used by *public authorities or on behalf of public authorities* to evaluate the eligibility of natural persons for essential public assistance.” In such cases, our description of the AI system was also explicit about the deployer (refer to Appendix A for the exact phrasing of all AI systems).

	AI system description	Risk level
#2	An AI system intended to be used to engage in conversation to help programmers code	Minimal risk
#5	An AI system intended to be used to reply to emails without human intervention	Minimal risk
#9	An AI system intended to be used to summarize books	Minimal risk
#18	An AI system intended to be used to grade students' written exams	High-risk (Annex III (3b))
#21	An AI system intended to be used to decide which employees receive promotions	High-risk (Annex III (4b))
#32	An AI system intended to be used to estimate the risk of a person becoming a victim of a criminal offense for the police	High-risk (Annex III (6a))
#42	An AI system intended to be used to estimate the risk of a person committing a criminal offense based on their personality for the police	Unacceptable risk (Art. 5(1)d)
#44	An AI system intended to be used to infer students' emotions during class	Unacceptable risk (Art. 5(1)f)

Table 1: Example AI systems covered by our study. Refer to Appendix A for the complete list of 48 AI systems. Articles and Annexes refer to the AI Act.

Taking into account the AI Act's potential—albeit limited [78]—extraterritorial scope of application to third countries [2], we also sought to recruit participants from the United States (US), considering that most large AI providers are US companies. Hence, we recruited study participants from the US, Germany, France, and Spain—the latter ones being major countries in the EU economically and in terms of population.

We recruited participants through Qualtrics, a survey panel aggregation company. Qualtrics recruits participants from several panel providers to mitigate potential panel recruitment biases and employs several automated fraud detection techniques for data quality. It also identifies and discards invalid responses from participants based on their behavior while completing the study (e.g., those speeding through the survey and not paying enough attention). Qualtrics compensates study participants according to the policies of the panels they aggregate; for instance, panels could compensate participants with airline miles, gift cards, and points in bonus programs.

Before recruiting participants, we conducted a simulation-based power analysis based on the mixed-effects models we report in the paper (see Section 3.5) to determine our sample size [34]. Using the mean perceived risk of a subset of 12 minimal, high-, and unacceptable risk AI systems gathered through a pilot study, we identified that a sample of 200 participants would be enough to detect similar effect sizes with 0.80 power. To account for inattentive participants and a larger number of scenarios, we sought to recruit 300 participants from each country with quotas matching its gender and age distribution.

In total, we analyzed 1413 responses after removing participants who failed any of the attention and comprehension checks (see Section 3.3). Our final sample comprised 724 (51.23%) participants who self-identified as men, 681 (48.20%) as women, and eight as non-binary or who preferred to self-describe or not disclose their gender. Participants' mean age was 50.7 years old ($SD = 15.9$), with the youngest participants being 19 and the oldest being 89 years old. Most participants did not have any background in disciplines related to computer science ($n = 1,127$, 79.76%) or law ($n = 1,222$, 86.48%). Participants were distributed similarly across the four countries ($n_{Germany} = 330$, $n_{France} = 348$, $n_{Spain} = 386$, $n_{US} = 349$). We

provide a detailed breakdown of participants' demographics by country in Table 5 in Appendix B.

3.3 Study Design

Figure 1 presents an overview of our methodology. After agreeing to the research terms, participants read the study's instructions and the AI Act's definitions of the terms "AI system" and "risk" (Art. 3 (1),(2) AI Act). Participants were then asked two comprehension questions to ensure that they understood these concepts as defined by the regulation. They had to indicate that "risk" refers to the "combination of the probability of a harm occurring and the severity of such harm" and that an "AI system" can be described as "a system that operates with varying levels of autonomy and adaptiveness and generates outputs" (see Appendix C for details). We discarded responses from participants who failed these comprehension checks.

Participants were then shown four random AI systems out of the 48 as described in Section 3.1, one at a time and in random order. For each AI system, participants answered questions concerning its potential risks and indicated how they believed it should be regulated. For the last AI system (of the four they read about), participants additionally indicated the perceived importance of a list of requirements prescribed by the AI Act. The study concluded with an instructed response attention check question, which was used to filter inattentive participants out, and a series of exploratory and demographic questions. We provide detailed study materials in Appendix C.

3.3.1 Study Translations: All study materials, including AI systems' descriptions, were initially drafted in English and then translated into German, French, and Spanish. The translations were also done by Qualtrics through an iterative process involving three professional translators. Each translation was then checked by a researcher from the authors' institutions who spoke the respective language natively. Study participants could choose their preferred language and change it at any moment while completing the study.

3.4 Measures

3.4.1 Perceived Risk of AI Systems: Rather than collecting the perceived risk of AI systems in an abstract manner, we decided to

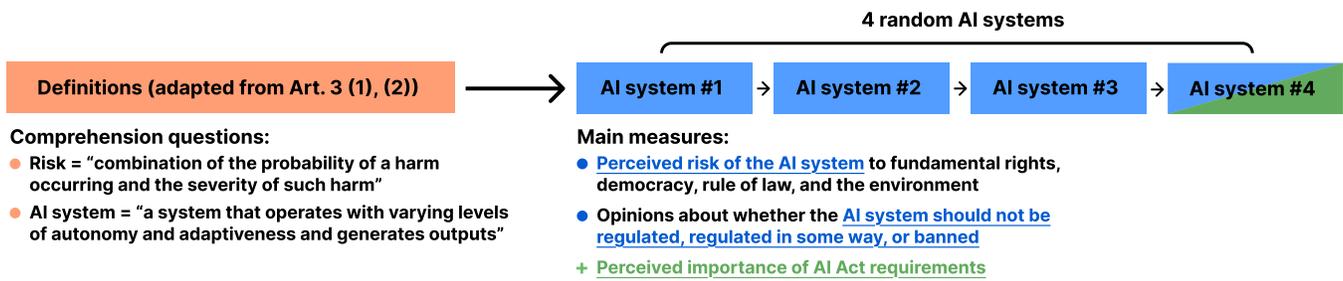


Figure 1: Overview of our study design.

ground participants’ risk assessments on fundamental rights and protected values that the AI Act strives to protect and promote. Given the legal definition of risk as the “the combination of the probability of an occurrence of harm and the severity of that harm” (under Art. 3(2) AI Act), we first sought to identify the object of harm to which this definition refers.

As stated in Art. 1(1), the AI Act aims to ensure a “high level of protection of health, safety, fundamental rights [...], including democracy, rule of law, and environmental protection.” We thus identified an extensive (albeit not exhaustive) list of rights and protected values that could be at risk by the AI systems covered by our research. Given our objective to study the AI Act in depth, we once more turned to EU law to determine these rights and protected values. The list was crafted based on the fundamental rights and values enshrined in the EU Charter, the TEU, and described in the AI Act (e.g., health and safety). We prioritized fundamental rights that were relevant to the AI systems we examined in the study, striving to find a balance between extensive coverage of EU fundamental rights and a manageable number of items that could be asked in a survey without overwhelming participants.

Before conducting the final study, we recruited participants to take part in a pilot study where we asked them to indicate to what extent they understood a number of fundamental rights and protected values using a closed-ended scale and an open-ended text box (see Appendix C.1 for additional information). Through this pilot study, we learned that laypeople did not fully understand some of these concepts. For instance, participants’ open-ended explanations of “rule of law” described it as a set of legal rules, without mentioning that the rule of law prescribes that the law applies to everyone equally, such that even governments can only act within legal limits. We thus decided to introduce each fundamental right and protected value using a language simpler than that used by legal scholarship, explicating the main real-world implication of each legal concept.

Table 2 shows how each fundamental right and protected value was presented in the study, as well as its legal basis. Participants reported the perceived risk of each AI system by rating the extent to which it poses a risk to each fundamental right and protected value (e.g., democracy and the rule of law) using a 7-point scale (0 = No risk at all, 3 = Moderate risk, 6 = Extreme risk). The fundamental rights and protected values were presented in random order between participants.

3.4.2 Opinions about AI Regulation: As explained in Section 2.1, the AI Act first categorizes AI systems based on their risk to EU fundamental rights and values to then determine the subset of rules applicable to each AI system. AI systems deemed to pose unacceptable risks are prohibited. High-risk AI must satisfy a list of mandatory requirements, while those deemed to pose only minimal risk are mostly subjected to voluntary requirements. In the context of this study, we translate these three options as the AI system i) being banned, ii) being regulated, or ii) not being regulated.

After reporting the perceived risk of a particular AI system, participants indicated their view on regulation by selecting which of the following “best describes [their] view on what should happen with” it:

- “The AI system **should be banned**” ⇒ equivalent to being prohibited under the AI due to an unacceptable risk;
- “The AI system **should be regulated in some way**” ⇒ equivalent to being high-risk under the AI Act and thus subjected to mandatory requirements;
- “The AI system **should not be regulated**” ⇒ equivalent to posing only minimal risks under the AI Act and thus not subjected to mandatory requirements.

3.4.3 Perceived Importance of AI Act Requirements: The AI Act imposes several requirements that are mandatory for high-risk AI systems (Art. 8-27 AI Act) and voluntary for minimal risk AI systems (Art. 56 AI Act). Furthermore, it also prescribes certain transparency requirements for applicable AI systems (Art. 50 AI Act). Yet, the technical details of these requirements are still being drawn up through the form of harmonized standards (Art. 40(2) AI Act), guidelines (Art. 96 AI Act), and future delegated acts (Art. 97 AI Act). With the objective of informing future specifications and implementations, we also asked participants to rate the perceived importance of some of these requirements.

Due to the numerous obligations in the AI Act and our attempt to avoid overwhelming participants, we focused on the general requirements imposed on providers of high-risk AI systems. To this end, we excluded specific obligations for importers (Art. 23 AI Act), distributors (Art. 24 AI Act), deployers (Art. 26 AI Act), providers of public services (Art. 27 AI Act), and those related to the AI value chain (Art. 25 AI Act). Other obligations were not considered due to their direct relation to already listed requirements (e.g., the obligation to create technical documentation under Art. 11 AI Act and the obligation to keep documentation pursuant to Art. 18 AI Act) or because they are not directly related to AI systems (e.g., the

“How much risk does the AI system pose to...?”

Exact phrasing in the survey question	Concrete right/value	Legal basis
Fundamental rights		
Your health	Right to health	Art. 1(1) AI Act
Your safety	Right to safety	Art. 1(1) AI Act
The protection of your personal data	Right to data protection	Art. 7 and 8(1) EU Charter
Your equal treatment under the law regardless of your gender, race, nationality, or other personal characteristics	Right to equality and non-discrimination	Art. 20 and 21 EU Charter / Art. 2 TEU
Your freedom to express ideas and receive information without interference	Right to freedom of expression	Art. 11 EU Charter
Your access to a fair trial	Right to a fair trial	Art. 47 EU Charter
Your freedom to seek employment	Right to engage in work	Art. 15 EU Charter
Your access to working conditions that respect your health, safety, and dignity	Right to fair and just working conditions	Art. 31 EU Charter
The protection of your intellectual property	Right to intellectual property	Art. 17(2) EU Charter
Your freedom to act without undue physical restraint or coercion	Right to liberty	Art. 6 EU Charter
Protected values		
The protection of the electoral system in a democratic State	Democracy	Art. 2 TEU / Art. 1(1) AI Act
The guarantee that governments can only act within the limits set by the law	Rule of law	Art. 2 TEU / Art. 1(1) AI Act
The environment	Environment	Art. 37 EU Charter / Art. 1(1) AI Act

Table 2: List of fundamental rights and protected values examined by our study. The list includes rights and values enshrined in the Charter of Fundamental Rights of the European Union (EU Charter), the Treaty on the Functioning of the European Union (TEU), and the AI Act. For each AI system participants read about, they indicated the extent to which it poses a risk to each of these concepts using a 7-point scale (0 = No risk at all, 3 = Moderate risk, 6 = Extreme risk).

obligation to cooperate with public authorities under Art. 21 AI Act).

We thus focused on the requirements imposed on AI providers by Art. 8-17 AI Act, as well as the transparency rules from Art. 50 AI Act. We asked participants how important they considered each of the requirements shown in Table 3 using a 7-point scale (0 = Not important at all, 3 = Moderately important, 7 = Extremely important). The list of requirements was presented in random order between participants.

3.5 Analysis Plan

We employed mixed-effects regressions and ANOVA tests to analyze our data. We treated participants’ risk assessments and the perceived importance of AI Act requirements as continuous variables.⁶ Participants’ opinions concerning whether AI systems should be banned or regulated were mapped onto categorical variables. To account for repeated measures across participants and AI systems (i.e., participants evaluated several AI systems, and AI systems were evaluated by multiple participants), we included crossed random intercepts for both participants and AI systems in the regressions.⁷

⁶We present robustness tests using ordinal regressions in the study’s online repository; all results are consistent with those reported in the main text. We report the results of linear regressions in the main text for easier interpretation.

⁷When analyzing the perceived importance of the AI Act requirements, we only include random intercepts for AI systems since participants only answered this question for one AI system (instead of four, as in the other survey questions).

Whenever an ANOVA test identified a significant effect at the $\alpha < 0.05$ level, we tested for pairwise differences using contrasts and applied Bonferroni corrections to account for multiple comparisons.

First, we explored participants’ risk judgments concerning AI systems. We regressed the perceived risk of each AI system to a dummy variable encoding its category under the AI Act: unacceptable, high, or minimal risk. We analyzed the risk to each fundamental right or protected value separately. This analysis focuses on whether participants assess the risk posed by AI systems similarly to how the AI Act categorizes them.

Second, we investigated whether participants had different opinions about AI regulation depending on the AI system and how it is categorized under the AI Act. For our regression analysis, we mapped participants’ responses into two different outcomes. We consider that a participant wants to *prohibit* a particular AI system if they indicated that it should be banned; inversely, we deem that the participant does not want to prohibit an AI system if they indicated that the AI system should “be regulated in some way” or “not be regulated.” Moreover, we consider that a participant wants a particular AI system to be *regulated* if they indicated that it should either be “banned” *or* “regulated in some way.” In contrast, we interpret that a participant does not believe the AI system should be regulated if they said so explicitly in the survey. We regressed

Requirements	Legal basis
To establish a risk management system for the AI system’s entire lifecycle to monitor and mitigate risks	Art. 9
To ensure that training, validation, and testing data are relevant, representative, and error-free	Art. 10
To create technical documentation to demonstrate compliance and assist authorities in assessment	Art. 11
To design the AI system to automatically record events related to risks and system updates	Art. 12
To provide clear instructions to deployers to help them maintain compliance	Art. 13
To design the AI system to support human oversight	Art. 14
To ensure the AI system is accurate, robust, and safe	Art. 15
To implement a quality management system to maintain compliance with requirements	Art. 17
To inform end users that they are interacting with the AI system when it is not obvious	Art. 50
To disclose that predictions, content, recommendations, or decisions are AI-generated when applicable	Art. 50

Table 3: List of requirements imposed by the AI Act explored in our study, as well as its source in the regulation. Participants indicated the importance of each requirement using a 7-point scale (0 = Not important at all, 3 = Moderately important, 7 = Extremely important). Articles refer to the AI Act.

these mapped variables to dummy variables encoding the AI Act’s risk categorization.⁸

Third, we analyzed whether the perceived risk of AI systems is associated with participants’ opinions about their potential prohibition or regulation. We thus regressed reported support for banning or regulating each AI system to its perceived risk, aiming to examine whether risk is associated with the intention to regulate AI—as done by the AI Act.

Finally, we examined the perceived importance of AI Act requirements. We regressed participants’ responses to the dummy variable encoding the AI system’s AI Act risk level to explore whether people have different expectations of compliance depending on the AI system and its regulated risk level.

We first report the results of our analysis considering responses from all countries together. We then discuss between-country differences and report a detailed analysis in Appendix D.1. We do so for brevity since all takeaways are consistent across all four countries. All of our data is available for replication in the study’s online repository.

4 Results

4.1 Perceived Risk of AI Systems

Figure 2 presents the mean perceived risk of all AI systems, as well as the risk of AI systems grouped by their AI Act risk category (see Figure 6 in Appendix D for the distribution). On average, participants considered all AI systems moderately risky to all fundamental rights and values the AI Act strives to protect. The right to the protection of personal data was deemed under the most risk ($M=3.76$, $SE=0.02$), whereas the environment was rated as the value under the least risk ($M=2.57$, $SE=0.02$). All in all, participants indicated that AI systems pose non-negligible risks to all fundamental rights and protected values examined by our study.

When we explored differences in the perceived risk between AI systems based on how the AI Act categorizes them, we identified

⁸Although participants’ opinions about banning or regulating AI were modeled as binary variables, we report the results of linear regressions for consistency with the other models and due to their more intuitive interpretation. For a discussion on the suitability and benefits of using linear models for binary outcomes, see Hellevik [38]. Nonetheless, we note that we observe consistent results using logit regressions.

significant differences for some—although not all—protected values. As Table 4 shows, participants rated unacceptable and high-risk AI systems as riskier than their minimal risk counterparts with respect to the fundamental rights to (i) health, (ii) safety, (iii) equality and non-discrimination, (iv) access to a fair trial, (v) freely seek employment, and (vi) fair and just working conditions. We found the same trend concerning AI systems’ perceived risk to the rule of law and the environment. We did not identify statistically significant differences between the risk posed by unacceptable- and high-risk AI systems to these fundamental rights and protected values.

Looking at the risk to freedom of expression and to democracy, unacceptable-risk AI systems were deemed riskier than high- and minimal risk systems, which were by contrast judged similarly. Concerning the right to the protection of personal data, we only identified significant differences between unacceptable- and minimal risk AI systems, with the former posing higher risks. For the fundamental right to liberty, unacceptable AI systems were considered riskier than their high-risk counterparts, which were in turn deemed to pose more risk than minimal risk AI systems. We did not observe any differences in perceived risk based on the AI Act categories for the fundamental right to intellectual property.

From a broader perspective, participants differentiated between unacceptable- and minimal risk AI systems, indicating that the former pose higher risks (as proposed by the AI Act). However, the distinction between unacceptable- and high-risk AI systems was less pronounced considering that they were judged similarly concerning their risk to most fundamental rights, the environment, and the rule of law. All in all, the difference between the three AI Act risk categories is not as prominent as implemented in the regulation. Even between AI systems categorized as unacceptable- and minimal risk, we only observed differences of less than one point on our 7-point risk scale.

4.2 Opinions about Prohibiting and Regulating AI Systems

Figure 3 shows that participants overwhelmingly indicate that AI systems should be regulated—but not necessarily prohibited. Even

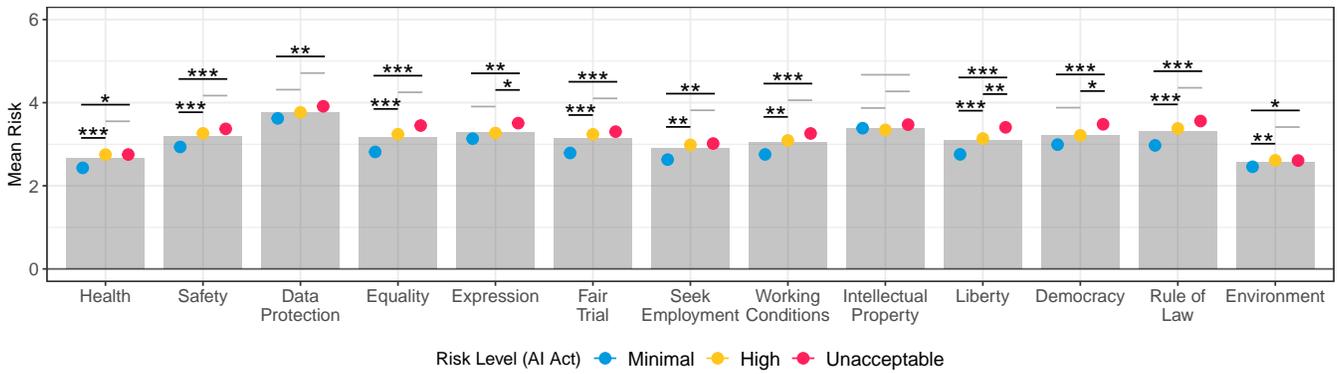


Figure 2: Perceived risk of AI systems to fundamental rights and protected values. Gray bars present the mean values across all AI systems, while circles represent mean perceived risk of AI systems grouped by their risk level according to the AI Act. The figure includes standard errors for the risk-level group means, though they are too small to be visible. Lines above the mean values denote pairwise comparisons, with symbols indicating statistically significant differences: * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

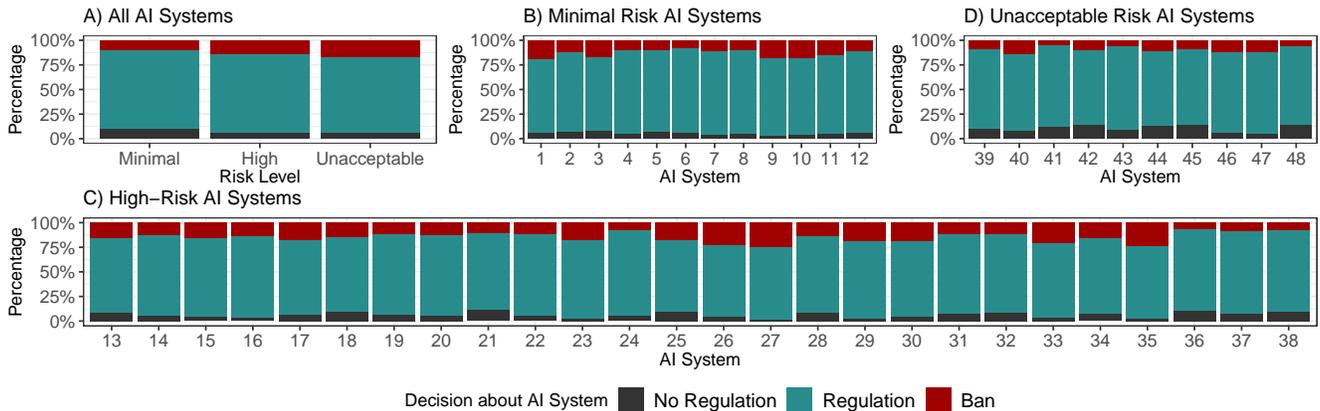


Figure 3: Distribution of participants’ views towards regulating or banning AI systems. Panel (A) groups AI systems based on their AI Act risk level. Panel (B)-(D) shows participants’ responses for each AI system examined in our study.

though participants were slightly more inclined to ban unacceptable-risk AI systems than high- and minimal risk ones (see Figure 3A and Table 6 in Appendix D), they largely suggest that all AI systems, no matter their risk level under the AI Act, should be regulated in some way other than through outright prohibitions.

A separate analysis of each scenario shows consistent patterns. Although one in four participants indicated that some unacceptable-risk AI systems (see #39 and #47 in Figure 3D) should be banned, most supported other forms of regulation. Similarly, Figure 3B shows that although a few participants believe minimal risk AI systems do not require regulatory intervention, most people showed support for regulation. We observed similar trends for high-risk AI systems, as shown in Figure 3C.

Examining the relationship between participants’ perceptions of the risks posed by an AI system and their views concerning its regulation, we identified that the riskier an AI system is to most fundamental rights and protected values, the more likely participants support its prohibition (see Table 6 in Appendix D).

Focusing on people’s opinions concerning any regulation (including a potential prohibition), AI systems deemed to pose higher risks to the fundamental rights to health, the protection of personal data, and equality, as well as to the rule of law, were more likely to be met with support for strict regulation. In contrast, we found the opposite trend for some values. Higher perceived risk to intellectual property was associated with lower support for prohibiting an AI system, while systems deemed to pose more risks to the environment were less likely to be considered for bans or regulation.

At a higher level, perceived risk was associated with reported intentions to prohibit and regulate AI systems, as proposed by the AI Act. Yet, the impact of these risk perceptions was only marginal because participants indicate that all AI systems—no matter their perceived risk to fundamental rights, democracy, the rule of law, and the environment, as well as its AI Act risk level—should be regulated in some way without outright prohibitions.

Comparison	Estimate	SE	df	<i>t</i>	<i>p</i> -value	Cohen's <i>d</i>	
Health			$F(2, 44.937) = 8.473, p < .001$				
Minimal - High	-0.3334	0.0835	44.58	-3.992	0.0007	0.31	
Minimal - Unacceptable	-0.3143	0.1028	45.08	-3.058	0.0112	0.29	
High - Unacceptable	0.0191	0.0893	45.11	0.214	1.0000	0.02	
Safety			$F(2, 45.052) = 12.117, p < .001$				
Minimal - High	-0.3685	0.0868	44.57	-4.247	0.0003	0.31	
Minimal - Unacceptable	-0.4731	0.1067	45.08	-4.432	0.0002	0.40	
High - Unacceptable	-0.1046	0.0928	45.12	-1.127	0.7972	0.10	
Data Protection			$F(2, 45.325) = 5.735, p = 0.006$				
Minimal - High	-0.1597	0.0844	44.59	-1.892	0.1949	0.14	
Minimal - Unacceptable	-0.3516	0.1038	45.08	-3.387	0.0044	0.30	
High - Unacceptable	-0.1919	0.0902	45.11	-2.126	0.1169	0.17	
Equality			$F(2, 44.785) = 22.673, p < .001$				
Minimal - High	-0.4482	0.0827	44.50	-5.418	<.0001	0.39	
Minimal - Unacceptable	-0.6453	0.1018	45.09	-6.337	<.0001	0.55	
High - Unacceptable	-0.1971	0.0885	45.14	-2.226	0.0931	0.17	
Expression			$F(2, 45.0008) = 6.079, p = 0.005$				
Minimal - High	-0.1364	0.0950	44.65	-1.436	0.4738	0.11	
Minimal - Unacceptable	-0.4016	0.1168	45.07	-3.439	0.0038	0.34	
High - Unacceptable	-0.2652	0.1015	45.10	-2.613	0.0365	0.22	
Fair Trial			$F(2, 44.642) = 11.570, p < .001$				
Minimal - High	-0.4481	0.1017	44.68	-4.405	0.0002	0.38	
Minimal - Unacceptable	-0.5076	0.1251	45.06	-4.059	0.0006	0.43	
High - Unacceptable	-0.0596	0.1087	45.09	-0.548	1.0000	0.05	
Seek Employment			$F(2, 44.981) = 7.520, p = 0.002$				
Minimal - High	-0.3690	0.1055	44.70	-3.498	0.0032	0.31	
Minimal - Unacceptable	-0.4331	0.1297	45.06	-3.339	0.0051	0.36	
High - Unacceptable	-0.0640	0.1127	45.09	-0.568	1.0000	0.05	
Work Conditions			$F(2, 45.121) = 11.931, p < .001$				
Minimal - High	-0.3459	0.0908	44.60	-3.811	0.0013	0.29	
Minimal - Unacceptable	-0.5203	0.1117	45.07	-4.660	0.0001	0.44	
High - Unacceptable	-0.1744	0.0971	45.11	-1.797	0.2372	0.15	
Intellectual Property			$F(2, 45.490) = 1.792, p = 0.178$				
Minimal - High	0.0321	0.0822	44.52	0.391	1.0000	0.03	
Minimal - Unacceptable	-0.1337	0.1012	45.09	-1.321	0.5795	0.11	
High - Unacceptable	-0.1658	0.0880	45.13	-1.885	0.1976	0.14	
Liberty			$F(2, 45.902) = 25.753, p < .001$				
Minimal - High	-0.4041	0.0779	44.45	-5.187	<.0001	0.35	
Minimal - Unacceptable	-0.6721	0.0959	45.10	-7.007	<.0001	0.58	
High - Unacceptable	-0.2680	0.0834	45.15	-3.214	0.0073	0.23	
Democracy			$F(2, 44.997) = 9.323, p < .001$				
Minimal - High	-0.2278	0.0933	44.64	-2.442	0.0559	0.19	
Minimal - Unacceptable	-0.4954	0.1147	45.07	-4.318	0.0003	0.41	
High - Unacceptable	-0.2676	0.0997	45.10	-2.683	0.0305	0.22	
Rule of Law			$F(2, 45.398) = 15.689, p < .001$				
Minimal - High	-0.4036	0.0891	44.61	-4.527	0.0001	0.34	
Minimal - Unacceptable	-0.5769	0.1097	45.07	-5.261	<.0001	0.48	
High - Unacceptable	-0.1733	0.0953	45.11	-1.818	0.2272	0.14	
Environment			$F(2, 44.757) = 7.280, p = 0.002$				
Minimal - High	-0.1631	0.0448	43.59	-3.643	0.0021	0.13	
Minimal - Unacceptable	-0.1643	0.0554	45.16	-2.965	0.0145	0.14	
High - Unacceptable	-0.0012	0.0482	45.25	-0.024	1.0000	0.00	

Table 4: Pairwise comparisons of the perceived risk posed by AI systems based on their AI Act risk level. We apply Bonferroni corrections to account for multiple comparisons.

4.3 Perceived Importance of AI Act Requirements

Participants assessed all AI Act requirements examined by our study to be important, as shown by Figure 4 (see also Figure 7 in Appendix D for the distribution of perceived importance). All requirements were rated close to four on our 7-point scale of perceived importance, with human oversight obligations being rated the least important ($M=4.05$, $SE=0.05$) and accuracy, robustness, and safety provisions as the most significant ($M=4.58$, $SE=0.04$).

We identified only marginal differences in the perceived importance of regulatory requirements for AI systems depending on how they are classified under the AI Act (see Figure 4). In line with our analysis of risk perceptions and regulatory intentions, some AI Act requirements were deemed more important when applied to high-risk AI systems (compared to minimal risk AI systems). Albeit statistically significant (see Table 7 in Appendix D), these differences were only marginal. Instead, we found that participants considered all requirements similarly important (approximately or over four in our 7-point scale) regardless of how the AI system was classified under the AI Act. All in all, the three-tier risk categories employed by the AI Act had little to no impact on people's opinions about AI Act requirements.

4.4 Between-Country Differences

Figure 5 shows that participants from Germany, Spain, France, and the US evaluate AI systems similarly with respect to risk. Across all risk levels proposed by the AI Act, participants from the four countries deemed all AI systems similarly risky, in line with our results reported above. Although participants from Germany and the US sometimes indicate that AI systems pose less risks, the differences are—albeit sometimes statistically significant—only marginal, as discussed in detail in Appendix D.1.

While participants from Germany and the US were marginally more likely to support the prohibition of AI systems classified as unacceptable in the AI Act, participants from the four countries largely suggest that all AI systems, regardless of their risk level under the AI Act, should be regulated without complete prohibitions (see Figure 8 in Appendix D.1). Similarly, participants from the four countries rated all regulatory requirements as similarly important (see Figure 9 in Appendix D.1). In summary, participants from the four countries had similar opinions concerning the risk that AI systems pose, as well as what should be done concerning their potential regulation and prohibition (see Appendix D.1).

5 Discussion

Below, we discuss the implications of our findings. Participants' views that all AI systems pose moderate levels of risk challenge the clearly defined risk levels proposed by the AI Act (§5.1). Furthermore, our findings reflect the general public's call for a more comprehensive and horizontal AI regulation that imposes requirements on all AI systems regardless of their level of risk (§5.2). The finding that the requirements imposed by the AI Act are perceived similarly important for all AI systems—irrespective of their regulatory risk level—demonstrates that end-users may expect AI providers to comply with voluntary obligations (§5.3). Our between-country

analysis also indicates potential for unified international governance for AI (§5.4). Finally, we reflect on our study's limitations (§5.5) and share some concluding remarks (§5.6).

5.1 All AI Systems Are Considered Risky—and to a Similar Extent

Participants deemed all AI systems—no matter their risk level according to the AI Act—moderately risky. Although we observed some small significant differences based on AI Act risk categories, we found that all AI systems examined in our study were perceived to pose largely similar levels of risk to fundamental rights, democracy, the rule of law, and the environment. In a similar vein, participants' assessments of risk did not vary much between different fundamental rights and protected values.

These results are surprising given that, through the lens of the AI Act, AI systems can pose different levels of risk to distinct fundamental rights and protected values. Had participants evaluated the risk posed by AI systems similarly to the AI Act, unacceptable AI systems would have been deemed extremely risky (e.g., around six on our 7-point scale); high-risk AI would have been considered moderate-to-highly risky (e.g., between three and five); while minimal risk AI systems should have been judged as posing only minimal risks (e.g., between zero and two).

Our results are largely aligned with prior work collecting risk assessments of AI systems in the media, health, psychology, transport, and justice domains [3, 73, 76]. In these studies, AI systems were rated from somewhat to moderately risky across all domains, with only marginal differences between scenarios due to opinions about the stakes involved in deploying AI in a particular domain [3]. Our findings extend these prior results to several other domains where AI can be deployed, such as education, employment, public services, and immigration. Considering the evidence that the perceived risk of AI impacts willingness to use it [76], our findings could also indicate that laypeople may be reluctant to use AI systems that pose minimal risk according to the AI Act. In contrast, the public may be relatively receptive to AI systems that the AI Act prohibits due to their excessive risks.

All in all, our findings suggest that laypeople recognize that AI can pose risks but may not differentiate between the varying levels of risk posed by distinct AI systems to their fundamental rights and protected values. In fact, our findings mirror the critique that the risk-based approach of the AI Act is hard to operationalize due to the inherent difficulty in estimating the risks of specific AI use cases [13, 18, 23, 68, 79].

It is possible that AI literacy interventions could play a role in assisting the general public and other relevant stakeholders (e.g., policymakers) in identifying the risks and benefits of different AI systems [89]. Research suggests that laypeople's understanding of AI is limited [8] and often detached from reality [12], which could impact their evaluation of potential risks [43]. In the context of the AI Act, AI literacy refers to the capacity to become aware about the potential risks and harms associated with AI (Art. 3(56) and Recital 20 AI Act). The AI Act calls for the promotion of AI literacy, assigning this task to the European Commission (Art. 66(f) and Recital 20 AI Act) and providers and deployers when it concerns their staff and persons operating the AI system (Art. 4 AI Act).

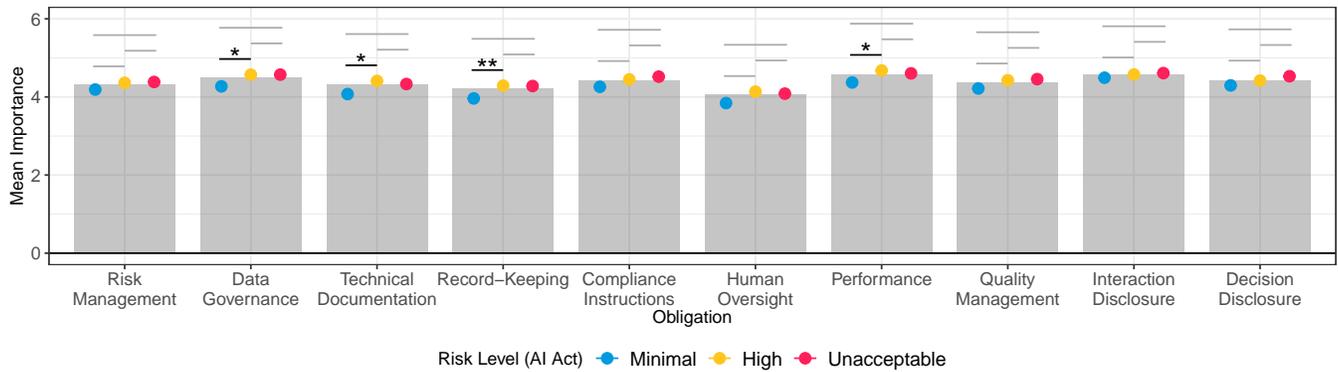


Figure 4: Perceived importance of AI Act requirements. Gray bars present the mean values across all AI systems, while circles represent mean perceived importance based on the AI Act risk levels. The figure includes standard errors for the risk-level group means, though they are too small to be visible. Lines above the mean values denote pairwise comparisons, with symbols indicating statistically significant differences: * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

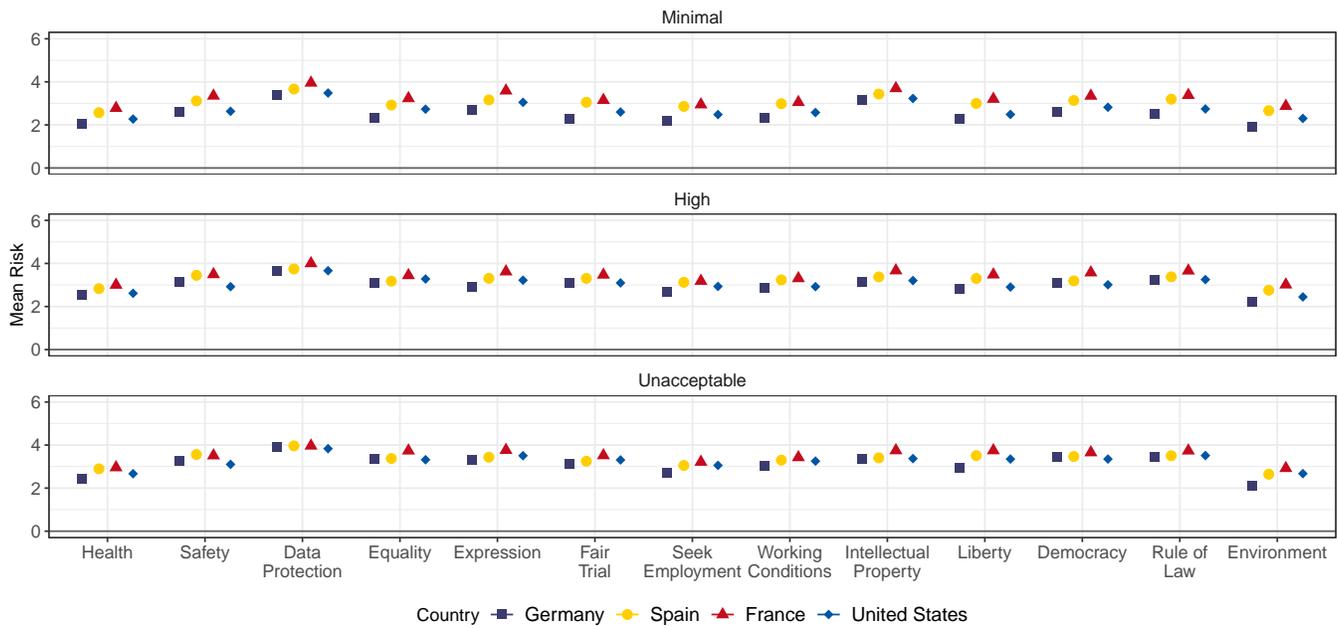


Figure 5: Perceived risk of AI systems to fundamental rights and protected values across Germany, Spain, France, and the US. Circles represent mean perceived risk of AI systems grouped by their risk level according to the AI Act and participants' country. The figure includes standard errors for the risk-level group means, though they are too small to be visible.

Accordingly, this study suggests that enhancing AI literacy is a key concern that requires the attention of the Commission, providers, and deployers.

From a development perspective, AI systems could also be designed in ways that promote appropriate reliance and accurate risk perceptions. Prior work suggests that designing explanations and reporting AI evaluation metrics impacts how much people rely on AI systems [45, 58, 90]. Similarly, a study illustrates that disclaimers about the limitations of AI could have an impact on people's perceptions of AI outputs [39]. It is imperative that AI systems are

developed in a way that fosters accurate perceptions of risk so that their end-users do not over- or underestimate real risks.

Without necessarily changing how AI systems are developed, future research could explore how to better introduce the risks posed by AI systems to the public. Those developing AI are incentivized to downplay potential risks and promote potential benefits of their systems [14, 49, 62], which could shape laypeople's perceptions of risk. For instance, the risks of AI could be better described in its model cards [60], which document important information about AI systems' training and evaluation. Yet, an audit of existing model

cards suggests that the potential risks of AI systems are one of the least reported information in model cards [51]. It is also unclear whether these model cards would be easily accessible to the general public, which could be a direction for future regulatory efforts on transparency.

From a legal perspective, our findings cast doubt on whether the AI Act's risk-based approach properly addresses the risks posed by AI systems. Legal scholars have pointed out that using lists of prohibited and high-risk practices to classify AI systems may not be well-suited to identify the risks posed by concrete applications [13, 18, 23, 79]. In particular, these classification lists are based on abstract risk assessments and may include AI systems whose concrete implementation could pose varying degrees of risk. Since the list of AI systems used in this study are based on the AI Act's classification criteria and phrased with a similar level of abstraction, our findings might be partially the result of participants picturing several different applications (and their corresponding risks) when introduced to one AI system. Without extensive details about how the AI system would be concretely implemented in the provided domain, participants might have taken a conservative approach when considering the risks.⁹

5.2 People Support Comprehensive AI Regulation, but Not Outright Prohibitions

Our study found that participants largely believe that all AI systems—irrespective of their risk level in the AI Act—should be regulated in some way. Furthermore, we identified that even if a particular AI system is deemed unacceptable and thus banned by the AI Act, people do not necessarily support its outright prohibition. In contrast to the three-level framework adopted by the AI Act, participants indicated that some form of regulation should be applied to AI systems of all risk levels.

Our findings are aligned with prior work collecting the public's attitudes towards AI regulation. As found in other surveys [20, 21, 32, 70, 85], our study participants agreed that AI systems should be subjected to stricter regulatory frameworks. Our study further demonstrates that this call for higher standards extend to a wide range of AI systems, including those that the AI Act deems minimal risk. Our careful focus on how the AI Act has been operationalized also illustrates that the public may be contrary to outright prohibitions, calling instead for wider oversight.

As proposed by the AI Act, we identified an association between perceived risk of an AI system and participants' opinions about its regulation. We identified stronger support for the regulation of an AI system if it was assessed to pose larger risks. Yet, the impact of perceived risk was only marginal since all AI systems in our study should be regulated in some form—no matter their perceived risk—in the eyes of the general public. Importantly, even if the

public may not clearly identify how different AI systems pose risks to their fundamental rights and protected values (as found in our study), they still recognize that AI systems pose risks and thus call for regulation that mitigates them.

A possible explanation for our results might be that laypeople support horizontal regulation that sets minimal standards for all AI systems. Many legal scholars have argued that the AI Act does not establish a minimum level of health, safety, or fundamental rights protection from AI systems [7, 31, 79, 84]. Our findings reflect these arguments: participants indicated that AI systems should respect a certain level of protection of fundamental rights and other values, which is at odds with the AI Act, as it does not prescribe any level of protection for AI systems that fall into the minimal risk category.

Finally, we note that the tendency to favor regulation over prohibition does not necessarily imply an aversion to the latter. If regulation is conceived as the establishment of a binding set of rules [6] that define the lawful use of AI systems, providers and deployers would consequently be prohibited from deploying AI systems that do not fulfill these standards. Our findings suggest, therefore, that while laypeople might not support general or unrestricted prohibitions, they might support the deployment of AI systems only if it fulfills certain minimum requirements.

5.3 AI Act Requirements Are Important for All AI Systems

All AI Act requirements we examined in our study were considered at least moderately important for all AI systems, regardless of their risk level under the AI Act. Our results are in line with the findings of a public consultation conducted during the drafting process of the AI Act [25], which found that requirements concerning training data, human oversight, performance, and record-keeping were considered important for high-risk AI systems. Yet, in our study, these requirements (and other obligations) were considered important not only for high-risk AI, but even for those categorized as minimal risk by the EU regulation.

Our findings suggest that people using or affected by AI systems may expect AI providers to comply with the AI Act requirements even if not mandated by the regulation, as in the case of AI systems deemed minimal risk by the regulation. Whereas the AI Act requires only high-risk systems to comply with regulatory requirements, our results suggest that the general public may call for providers of minimal risk AI systems to voluntarily comply with these requirements, potentially through the voluntary codes of conduct in Art. 95 AI Act. As suggested by prior work on public perceptions of AI ethical guidelines [44], complying with voluntary requirements can have a positive impact on people's acceptance of AI. Providers could thus consider complying with regulations even if not mandated by law—as in the case of minimal risk AI systems—to ensure that their AI systems are aligned with user expectations.

Considered alongside the findings described in Sections 5.1 and 5.2, our results cast doubt on the adequacy of the AI Act's risk-based approach for regulating AI systems from the perspective of citizens. In particular, our findings challenge applying the concept of risk to determine which AI systems are subject to mandatory requirements and which are not. We also note that there exist other risk-based regulatory approaches, some of which are already employed in

⁹The AI Act partially recognizes potential misalignments between its classification criteria and the risks posed by concrete applications of AI systems. Under Art. 6(3) AI Act, AI systems that would fall within the high-risk level according to Annex III but do not pose "significant risk" are not classified as high-risk. This exception thus acknowledges that the list in Annex III may include applications in critical areas, such as education or justice, that do not pose enough risk to be classified as high-risk. In contrast, there is no provision requiring the classification of AI systems that do not fall within the prohibited or high-risk lists but still pose significant risks to fundamental rights and other protected values—which is the case for some AI systems explored in our study.

the EU [19]. For example, under the EU General Data Protection Regulation (GDPR),¹⁰ all processing of personal data is subject to mandatory requirements (Art. 2(1) GDPR), and the measures necessary to ensure compliance are shaped by the risk posed by the specific processing [19, 31]. Our findings—that all AI systems were deemed moderately risky and that the set of requirements was considered relevant for all AI applications—align more closely with the GDPR model than with the tiered risk-based approach in the AI Act.

5.4 Alignment Across Jurisdictions—Unification in AI Regulation?

For this study, we recruited participants from three EU countries (Germany, France, and Spain) and one non-EU country (the US). Surprisingly, no substantial differences were found between countries. This was unexpected, particularly between EU countries and the US, given their distinct legal traditions and contrasting approaches to AI regulation [1] and prior work suggesting that individuals from these countries have different expectations about privacy and data protection [24, 40], which were fundamental rights covered by our study.

The observed agreement concerning the potential risks of AI, as well as measures to address them, might point to a convergence of regulatory approaches to AI. Ultimately, our findings provide initial support for the idea of internationally accepted minimum standards for AI systems to protect affected individuals and societies. While some significant efforts in this regard already exist [63–65], there are still no globally applicable and legally binding international rules for AI.¹¹

The extent to which the AI Act, with its extraterritorial scope of application—and the potential Brussels effect [2, 9, 33, 78]—will still continue to influence new legislative attempts to regulate AI systems remains to be seen. In fact, some scholars have questioned whether the EU approach will shape other international regulatory efforts [68]. Nevertheless, there is already a convergence in national efforts to regulate AI systems based on the notion of risk [23], or more specifically, on the classification of high-risk AI. This tendency is confirmed by the already adopted regulations in South Korea [59] and the state of Colorado in the US [66]. However, even when regulations adopt a risk-based approach, the notion of risk and the classification criteria for AI systems appear to differ.¹²

5.5 Limitations

Our findings are based on one-dimensional assessments of risk. We did so for simplicity as we sought to explore risks to several fundamental rights and protected values, as well as to examine multiple different AI systems. However, risk can be multidimensional [73, 76]. This multidimensionality is also acknowledged by

the AI Act, which defines risk as the combination of the likelihood and severity of a harm. Future work could also elicit risk assessments using different methods [22, 87].

Our study also did not cover all fundamental rights and values enshrined in EU treaties and other legal instruments as doing so would be unfeasible. Participants' assessments that all fundamental rights and protected values are similarly under risk was discussed as an indication that laypeople recognize that AI poses risks but may not differentiate between the risks posed by different AI systems. It is also possible that these results may have partially emerged due from our decision to ask about multiple fundamental rights and other protected values, which are multifaceted concepts that cannot be easily translated simple sentences. Participants could thus have struggled to interpret and differentiate between all the rights and values we examined, despite our best efforts in simplifying the language through iterative pilot testing. Future work could explore different ways of capturing perceptions of risk to different protected values; for instance, studies could focus on a specific fundamental right for each participant while varying the AI systems under consideration. Similarly, future studies could also explore different ways of referring to the fundamental rights and protected values.

The identified lack of support for outright prohibitions of AI systems could also have been the result of the difficulty in describing scenarios of prohibited practices according to Art. 5 AI Act. The AI practices described therein employ terms that are difficult to translate into concrete scenarios without influencing participants' responses. For example, Art. 5(1)(b) AI Act prohibits AI practices that exploit people's vulnerability to distort the behavior "in a manner that causes or is reasonable likely to cause that person or another person significant harm." While the use of sensible information (age, disability or social situation) can be neutrally described (see AI system #40 in Appendix A), describing the AI system as causing harm to people might nudge the response in favor of prohibitions.

Furthermore, we designed and chose the scenarios for AI systems under the assumption that they are placed on the market and would fall under the scope of the AI Act as prescribed by Art. 2. Thus, we did not consider private development of AI systems or use cases exempt from the AI Act, such as those developed and put into service solely for scientific research and development (Art. 2(6) AI Act). Our finding that laypeople support regulation of AI systems can thus only hold insofar these systems are placed on the market and do not fall under the exceptions.

Our research sought to provide an in-depth analysis of the AI Act and its European context in light of the public's perceptions of several AI systems. Instead of aiming for breadth by exploring several different approaches to AI regulation, we focused on depth, going deep into one particular regulatory strategy. This decision enabled us to rigorously examine whether the AI Act—as implemented by the EU—is aligned with the public's regulatory expectations and risk perceptions. However, our study design does not allow us to examine other approaches to AI regulation, such as those that are not risk-based or based on values other than those protected in the EU. Finally, we restricted our participant recruitment to three major EU countries and the US. However, AI regulation is under debate in many other countries around the world [4, 41, 69, 72]. Future work could compare expectations of AI regulation around the world.

¹⁰Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC.

¹¹Although [65] is legally binding and has been signed by the EU and other 16 countries, no country has yet ratified it.

¹²For example, the Colorado Consumer Protection for Artificial Intelligence Act [66] categorizes as high-risk "any artificial intelligence system that, when deployed, makes, or is a substantial factor in making, a consequential decision" (Part 17(9)a). Despite the inclusion of certain exceptions in Part 17(9)b, this classification criterion is more extensive and abstract compared to the criteria outlined in the AI Act.

5.6 Concluding Remarks

In this paper, we examined how the general public's perceptions of AI systems and their regulation are aligned (or misaligned) with the AI Act's risk-based approach to regulating AI. Our findings challenge the "clearly defined risk-based approach" (Recital 26 AI Act) that neatly categorizes AI systems into three distinct risk levels. Instead, we found that laypeople—both from the EU and the US—recognize that all AI systems can pose moderate levels of risks and should thus be subjected to mandatory requirements. We also showed how—compared to the AI Act—AI users may downplay the risks posed by prohibited AI systems, while at the same time overestimating the risks of AI classified as minimal risk. Our finding suggest that people could support AI regulation that focuses on setting minimal regulatory standards for AI systems irrespective of their associated risk level, establishing a minimal level of protection to fundamental rights and other protected values.

Acknowledgments

The research was partially funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC 2092 CASA - 390781972 and the Topic Engineering Secure Systems of the Helmholtz Association (HGF). It was also supported by KASTEL Security Research Labs in Karlsruhe, Germany. We thank Abraham Mhaidli and Jérémy Thibault for proofreading our study translations.

References

- [1] Sacha Alanoca, Shira Gur-Arieh, Tom Zick, and Kevin Klyman. 2025. Comparing Apples to Oranges: A Taxonomy for Navigating the Global Landscape of AI Regulation. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency (FAccT '25)*. Association for Computing Machinery, New York, NY, USA, 914–937. doi:10.1145/3715275.3732059
- [2] Marco Almada and Anca Radu. 2024. The Brussels side-effect: How the AI act can reduce the global reach of EU policy. *German Law Journal* 25, 4 (2024), 646–663.
- [3] Theo Araujo, Natali Helberger, Sanne Kruikemeier, and Claes H De Vreese. 2020. In AI we trust? Perceptions about automated decision-making by artificial intelligence. *AI & society* 35, 3 (2020), 611–623.
- [4] Alessandro Aveni. 2025. Overview of AI International and Brazil regulations. *Revista Jurídica da Presidência* 27, 142 (2025), 317–349.
- [5] Edmond Awad, Sohan Dsouza, Jean-François Bonnefon, Azim Shariff, and Iyad Rahwan. 2020. Crowdsourcing moral machines. *Commun. ACM* 63, 3 (2020), 48–55.
- [6] Robert Baldwin, Martin Cave, and Martin Lodge. 2012. *Understanding regulation: theory, strategy, and practice*. Oxford university press.
- [7] Lily Ballot Jones, Julia Thornton, and Daswin De Silva. 2025. Limitations of risk-based artificial intelligence regulation: a structuration theory approach. *Discover Artificial Intelligence* 5, 1 (2025), 14.
- [8] Arne Bewersdorff, Xiaoming Zhai, Jessica Roberts, and Claudia Nerdel. 2023. Myths, mis- and preconceptions of artificial intelligence: A review of the literature. *Computers and Education: Artificial Intelligence* 4 (2023), 100143.
- [9] Anu Bradford. 2020. *The Brussels effect: How the European Union rules the world*. Oxford University Press.
- [10] Peter Burgstaller. 2023. The Use of AI and Its Legal Boundaries in the EU. In *2023 1st International Conference on Optimization Techniques for Learning (ICOTL)*. IEEE, 1–5.
- [11] Marta Cantero Gamito and Christopher T Marsden. 2024. Artificial intelligence co-regulation? The role of standards in the EU AI Act. *International journal of law and information technology* 32, 1 (2024), eae011.
- [12] Stephen Cave and Kanta Dihal. 2019. Hopes and fears for intelligent machines in fiction and reality. *Nature Machine Intelligence* 1, 2 (2019), 74–78.
- [13] Johanna Chamberlain. 2023. The risk-based approach of the European Union's proposed artificial intelligence regulation: Some comments from a tort law perspective. *European Journal of Risk Regulation* 14, 1 (2023), 1–13.
- [14] Rachel Coldicutt. 2024. AI safety is a narrative problem. *Harvard Data Science Review Special Issue* 5 (2024).
- [15] European Commission. 2021. Better Regulation Guidelines.
- [16] European Commission. 2021. Impact Assessment Accompanying the Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts.
- [17] European Commission. 2025. Commission Guidelines on prohibited artificial intelligence practices established by Regulation (EU) 2024/1689 (AI Act).
- [18] Jerome De Cooman. 2022. Humpty dumpty and high-risk AI systems: the rationale materiae dimension of the proposal for an EU artificial intelligence act. *Mkt. & Competition L. Rev.* 6 (2022), 49.
- [19] Giovanni De Gregorio and Pietro Dunn. 2022. The European Risk-Based Approaches: Connecting Constitutional Dots in the Digital Age. *Common Market Law Review* 59 (2022), 473.
- [20] Noemi Dreksler, Harry Law, Chloe Ahn, Daniel Schiff, Kaylyn Jackson Schiff, and Zachary Peskowitz. 2025. What Does the Public Think About AI? An Overview of the Public's Attitudes Towards AI and a Resource for Future Research. (2025).
- [21] Noemi Dreksler, David McCaffary, Lauren Kahn, Kate Mays, Markus Anderljung, Allan Dafoe, M Horowitz, and Baobao Zhang. 2023. Preliminary Survey Results: US and European Publics Overwhelmingly and Increasingly Agree That AI Needs to Be Managed Carefully. *Centre for the Governance of AI* (2023).
- [22] Rithika Dulam, Christina Gore, and Jennifer Helgeson. 2024. Methods for elicitation of risk preferences and perceptions. (2024).
- [23] Martin Ebers. 2024. Truly risk-based regulation of artificial intelligence how to implement the EU's AI Act. *European Journal of Risk Regulation* (2024), 1–20.
- [24] Emma Engström, Kimmo Eriksson, Marie Björnstjerna, and Pontus Strimling. 2023. Global variations in online privacy concerns across 57 countries. *Computers in Human Behavior Reports* 9 (2023), 100268.
- [25] European Commission. 2020. White Paper on Artificial Intelligence: Public consultation towards a European approach for excellence and trust. <https://tinyurl.com/hjc66yzj>.
- [26] Tina Feldkamp, Markus Langer, Leo Wies, and Cornelius J König. 2023. Justice, trust, and moral judgements when personnel selection is supported by algorithms. *European Journal of Work and Organizational Psychology* (2023), 1–16.
- [27] Casey Fiesler. 2020. Lawful users: Copyright circumvention and legal constraints on technology use. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–11.
- [28] Henry Fraser and José-Miguel Bello y Villarino. 2024. Acceptable risks in Europe's proposed AI act: reasonableness and other principles for deciding how much risk management is enough. *European Journal of Risk Regulation* 15, 2 (2024), 431–446.
- [29] Gustavo Gil Gasiola. 2025. The GPAI Code of Practice. *Verfassungsblog* (2025).
- [30] Gustavo Gil Gasiola. 2025. Rebuilding the pyramid: The AI Act's risk-based approach using a binary decision diagram. *Computer Law & Security Review* 58 (2025), 106189.
- [31] Raphaël Gellert. 2021. The role of the risk-based approach in the General data protection Regulation and in the European Commission's proposed Artificial Intelligence Act: Business as usual? *Journal of Ethics and Legal Technologies* (2021).
- [32] Nicole Gillespie, Steven Lockey, Caitlin Curtis, Javad Pool, and Ali Akbari. 2023. Trust in artificial intelligence: A global study. *The University of Queensland and KPMG Australia* 10 (2023).
- [33] Delaram Golpayegani, Harshvardhan J Pandit, and Dave Lewis. 2023. To be high-risk, or not to be—semantic specifications and implications of the AI act's high-risk AI applications and harmonised standards. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. 905–915.
- [34] Peter Green and Catriona J MacLeod. 2016. SIMR: An R package for power analysis of generalized linear mixed models by simulation. *Methods in Ecology and Evolution* 7, 4 (2016), 493–498.
- [35] Nina Grgic-Hlaca, Elissa M Redmiles, Krishna P Gummadi, and Adrian Weller. 2018. Human perceptions of fairness in algorithmic decision making: A case study of criminal risk prediction. In *proc. of the Web conference*. 903–912.
- [36] Galen Harrison, Julia Hanson, Christine Jacinto, Julio Ramirez, and Blase Ur. 2020. An empirical study on the perceived fairness of realistic, imperfect machine learning models. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*. 392–402.
- [37] Sophia Hatz, Noemi Dreksler, Kevin Wei, and Baobao Zhang. 2025. Local US officials' views on the impacts and governance of AI: Evidence from 2022 and 2023 survey waves. *PLoS One* 20, 10 (2025), e0332919.
- [38] Ottar Hellevik. 2009. Linear versus logistic regression when the dependent variable is a dichotomy. *Quality & quantity* 43 (2009), 59–74.
- [39] Angelica Lermann Henestrosa and Joachim Kimmerle. 2025. "Always check important information!"—The role of disclaimers in the perception of AI-generated content. *Computers in Human Behavior: Artificial Humans* 4 (2025), 100142.
- [40] Franziska Herbert, Steffen Becker, Leonie Schaewitz, Jonas Hielscher, Marvin Kowalewski, Angela Sasse, Yasemin Acar, and Markus Dürmuth. 2023. A world full of privacy and security (mis) conceptions? Findings of a representative survey in 12 countries. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–23.

- [41] Emmie Hine and Luciano Floridi. 2023. The blueprint for an ai bill of rights: in search of enaction, at risk of inaction. *Minds and Machines* 33, 2 (2023), 285–292.
- [42] C Jackson. 2025. Google/Ipsos Multi-Country AI Survey 2025. *Ipsos, Washington, DC, Tech. Rep* (2025).
- [43] Ruogu Kang, Laura Dabbish, Nathaniel Fruchtler, and Sara Kiesler. 2015. “{My} data just goes {Everywhere:}” user mental models of the internet and implications for privacy and security. In *Eleventh symposium on usable privacy and security (SOUPS 2015)*. 39–52.
- [44] Kimon Kieslich, Birte Keller, and Christopher Starke. 2022. Artificial intelligence ethics by design. Evaluating public perception on the importance of ethical design principles of artificial intelligence. *Big Data & Society* 9, 1 (2022), 20539517221092956.
- [45] Sunnie SY Kim, Jennifer Wortman Vaughan, Q Vera Liao, Tania Lombrozo, and Olga Russakovsky. 2025. Fostering appropriate reliance on large language models: The role of explanations, sources, and inconsistencies. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–19.
- [46] Acar Kutay. 2015. Limits of participatory democracy in European governance. *European Law Journal* 21, 6 (2015), 803–818.
- [47] Markus Langer, Cornelius J König, Caroline Back, and Victoria Hemsing. 2023. Trust in Artificial Intelligence: Comparing trust processes between human and automated trustees in light of unfair bias. *Journal of Business and Psychology* 38, 3 (2023), 493–508.
- [48] Markus Langer, Cornelius J König, and Maria Papathanasiou. 2019. Highly automated job interviews: Acceptance under the influence of stakes. *International Journal of Selection and Assessment* 27, 3 (2019), 217–234.
- [49] Seth Lazar and Alondra Nelson. 2023. AI safety on whose terms? 138–138 pages.
- [50] Min Kyung Lee. 2018. Understanding perception of algorithmic decisions: Fairness, trust, and emotion in response to algorithmic management. *Big Data & Society* 5, 1 (2018), 2053951718756684.
- [51] Weixin Liang, Nazneen Rajani, Xinyu Yang, Ezinwanne Ozoani, Eric Wu, Yiqun Chen, Daniel Scott Smith, and James Zou. 2024. Systematic analysis of 32,111 AI model cards characterizes documentation practice in AI. *Nature Machine Intelligence* 6, 7 (2024), 744–753.
- [52] Gabriel Lima, Nina Grgic-Hlača, Jin Keun Jeong, and Meeyoung Cha. 2023. Who should pay when machines cause harm? Laypeople’s expectations of legal damages for machine-caused harm. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. 236–246.
- [53] Gabriel Lima, Nina Grgic-Hlača, Markus Langer, and Yixin Zou. 2025. Lay Perceptions of Algorithmic Discrimination in the Context of Systemic Injustice. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–30.
- [54] Gabriel Lima, Nina Grgic-Hlača, and Elissa M Redmiles. 2025. Public Opinions About Copyright for AI-Generated Art: The Role of Egocentricity, Competition, and Experience. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–32.
- [55] Tobias Mahler. 2021. Between risk management and proportionality: The risk-based approach in the EU’s Artificial Intelligence Act Proposal. *Nordic Yearbook of Law and Informatics* (2021).
- [56] Gianclaudio Malgieri and Frank Pasquale. 2024. Licensing high-risk artificial intelligence: Toward ex ante justification for a disruptive technology. *Computer Law & Security Review* 52 (2024), 105899.
- [57] Christian Marxsen. 2015. Open Stakeholder Consultations at the European Level—Voice of the Citizens? *European Law Journal* 21, 2 (2015), 257–280.
- [58] Siddharth Mehrotra, Chadha Degachi, Oleksandra Vereschak, Catholijn M Jonker, and Myrthe L Tielman. 2024. A systematic review on fostering appropriate trust in Human-AI interaction: Trends, opportunities and challenges. *ACM Journal on Responsible Computing* 1, 4 (2024), 1–45.
- [59] Ministry of Science and ICT of the Republic of Korea. 2024. Press Release: A New Chapter in the Age of AI: Basic Act on AI Passed at the National Assembly’s Plenary Session. <https://tinyurl.com/373jwdke>.
- [60] Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. 2019. Model cards for model reporting. In *Proceedings of the conference on fairness, accountability, and transparency*. 220–229.
- [61] Gustavo Moreira, Edyta Paulina Bogucka, Marios Constantinides, and Daniele Quercia. 2025. The Hall of AI Fears and Hopes: Comparing the Views of AI Influencers and those of Members of the US Public Through an Interactive Platform. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–27.
- [62] Arvind Narayanan and Sayash Kapoor. 2024. AI snake oil: What artificial intelligence can do, what it can’t, and how to tell the difference. In *AI Snake Oil*. Princeton University Press.
- [63] United Nations. 2024. Global Digital Compact.
- [64] OECD. 2019. Recommendation of the Council on Artificial Intelligence.
- [65] Council of Europe. 2024. Framework Convention on Artificial Intelligence.
- [66] General Assembly of the State of Colorado. 2024. Senate Bill 24-205 - Concerning Consumer Protections in Interactions with Artificial Intelligence Systems.
- [67] High-Level Expert Group on Artificial Intelligence. 2019. Ethics Guidelines for Trustworthy AI.
- [68] Ugo Pagallo. 2023. Why the AI Act won’t trigger a Brussels effect. *AI Approaches to the Complexity of Legal Systems (Springer 2024), Forthcoming* (2023).
- [69] Do Hyun Park, Eunjung Cho, and Yong Lim. 2024. A Tough Balancing Act—The Evolving AI Governance in Korea. *East Asian Science, Technology and Society: An International Journal* 18, 2 (2024), 135–154.
- [70] Jacob Poushter, Moira Fagan, and Manolo Corichi. 2025. *How People Around the World View AI: More are concerned than excited about its use, and more trust their own country and the EU to regulate it than trust the U.S. or China*. Pew Research Center.
- [71] Harrison Rainie. 2022. *AI and Human Enhancement: American’s Openness Is Tempered by a Range of Concerns: Public views are tied to how these technologies would be used, what constraints would be in place*. Pew Research Center.
- [72] Huw Roberts, Josh Cows, Jessica Morley, Mariarosaria Taddeo, Vincent Wang, and Luciano Floridi. 2021. The Chinese approach to artificial intelligence: an analysis of policy, ethics, and regulation. In *Ethics, governance, and policies in artificial intelligence*. Springer, 47–79.
- [73] Nadia Said, Andreea E Pottinteu, Irina Brich, Jürgen Buder, Hanna Schumm, and Markus Huff. 2023. An artificial intelligence perspective: How knowledge and confidence shape risk and benefit perception. *Computers in human behavior* 149 (2023), 107855.
- [74] Jonas Schuett. 2023. Defining the scope of AI regulations. *Law, Innovation and Technology* 15, 1 (2023), 60–82.
- [75] Jonas Schuett. 2024. Risk management in the artificial intelligence act. *European Journal of Risk Regulation* 15, 2 (2024), 367–385.
- [76] Rebekka Schwesig, Irina Brich, Jürgen Buder, Markus Huff, and Nadia Said. 2023. Using artificial intelligence (AI)? Risk and opportunity perception of AI predict people’s willingness to use AI. *Journal of Risk Research* 26, 10 (2023), 1053–1084.
- [77] Robert Schütze. 2015. EU Competences: Existence and Exercise. In *The Oxford Handbook of European Union Law*. Oxford University Press. doi:10.1093/oxfordhb/9780199672646.013.44
- [78] Charlotte Siegmund and Markus Anderljung. 2022. The Brussels effect and artificial intelligence: How EU regulation will impact the global AI market.
- [79] Nathalie A Smuha, Emma Ahmed-Rengers, Adam Harkens, Wenlong Li, James MacLaren, Riccardo Piselli, and Karen Yeung. 2021. How the EU can achieve legally trustworthy AI: a response to the European Commission’s proposal for an Artificial Intelligence Act. *SSRN* (2021).
- [80] Megha Srivastava, Hoda Heidari, and Andreas Krause. 2019. Mathematical notions vs. human perception of fairness: A descriptive approach to fairness for machine learning. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*.
- [81] Christopher Starke, Janine Baleis, Birte Keller, and Frank Marcinkowski. 2022. Fairness perceptions of algorithmic decision-making: A systematic review of the empirical literature. *Big Data & Society* 9, 2 (2022), 20539517221115189.
- [82] Kees Stuurman and Eric Lachaud. 2022. Regulating AI. A label to complete the proposed Act on Artificial Intelligence. *Computer Law & Security Review* 44 (2022), 105657.
- [83] Kevin Tobia. 2022. Experimental jurisprudence. *The University of Chicago Law Review* 89, 3 (2022), 735–802.
- [84] Michael Veale and Frederik Zuiderveen Borgesius. 2021. Demystifying the Draft EU Artificial Intelligence Act—Analysing the good, the bad, and the unclear elements of the proposed approach. *Computer Law Review International* 22, 4 (2021), 97–112.
- [85] Stepan Vesely and Byungdo Kim. 2024. Survey evidence on public support for AI safety oversight. *Scientific Reports* 14, 1 (2024), 31491.
- [86] Sandra Wachter. 2023. Limitations and loopholes in the EU AI Act and AI Liability Directives: what this means for the European Union, the United States, and beyond. *Yale JL & Tech.* 26 (2023), 671.
- [87] Robyn S Wilson, Adam Zwickle, and Hugh Walpole. 2019. Developing a broadly applicable measure of risk perception. *Risk Analysis* 39, 4 (2019), 777–791.
- [88] Benno G Wissing and Marc-André Reinhard. 2018. Individual differences in risk perception of artificial intelligence. *Swiss Journal of Psychology* 77, 4 (2018), 149.
- [89] Shixian Xie, John Zimmerman, and Motahhare Eslami. 2025. Exploring What People Need to Know to be AI Literate: Tailoring for a Diversity of AI Roles and Responsibilities. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–16.
- [90] Ming Yin, Jennifer Wortman Vaughan, and Hanna Wallach. 2019. Understanding the effect of accuracy on trust in machine learning models. In *Proceedings of the 2019 chi conference on human factors in computing systems*. 1–12.

A List of AI Systems

• Minimal-risk AI systems:

- (1) An AI system intended to be used to engage in small talk with users
- (2) An AI system intended to be used to engage in conversation to help programmers code
- (3) An AI system intended to be used to engage in conversation to give users travel advice
- (4) An AI system intended to be used to reply to emails without human intervention
- (5) An AI system intended to be used to assist users in writing emails
- (6) An AI system intended to be used to write novels without human intervention
- (7) An AI system intended to be used to summarize news stories about a particular topic
- (8) An AI system intended to be used to summarize books
- (9) An AI system intended to be used to infer people's preferences to recommend music
- (10) An AI system intended to be used to organize customer files in a company
- (11) An AI system intended to be used to retrieve user profiles from a social media platform based on a search query
- (12) An AI system intended to be used to recommend movies based on individual preferences

• High-risk AI systems:

- (13) An AI system intended to be used to assist judges in determining sentences - from (8a) Annex III
- (14) An AI system intended to be used to summarize legal documents for judges in court - from (8a) Annex III
- (15) An AI system intended to be used to infer people's emotions in a shopping mall - from (1c) Annex III
- (16) An AI system intended to be used to manage traffic based on the number of cars and pedestrians - from (2) Annex III
- (17) An AI system intended to be used to determine whether university applicants are admitted to university - from (3a) Annex III
- (18) An AI system intended to be used to grade students' written exams - from (3b) Annex III
- (19) An AI system intended to be used to determine to whom a job vacancy is advertised - from (4a) Annex III
- (20) An AI system intended to be used to evaluate job applicants - from (4a) Annex III
- (21) An AI system intended to be used to decide which employees receive promotions - from (4b) Annex III
- (22) An AI system intended to be used to decide which employees are terminated - from (4b) Annex III
- (23) An AI system intended to be used to monitor employees' behavior at work - from (4b) Annex III
- (24) An AI system intended to be used to monitor employees' performance at work - from (4b) Annex III
- (25) An AI system intended to be used to allocate tasks to employees based on personal traits - from (4b) Annex III
- (26) An AI system intended to be used to determine whether an applicant is granted housing loans by the government - from (5a) Annex III

- (27) An AI system intended to be used to determine whether an applicant has their housing loans revoked by the government - from (5a) Annex III
 - (28) An AI system intended to be used to assess a person's creditworthiness - from (5b) Annex III
 - (29) An AI system intended to be used to establish priority in the dispatch of ambulances - from (5d) Annex III
 - (30) An AI system intended to be used to assess personal risk factors to determine the pricing of private health insurance - from (5c) Annex III
 - (31) An AI system intended to be used to estimate the risk of a person committing a criminal offense based on their criminal records for the police - from (6d) Annex III
 - (32) An AI system intended to be used to estimate the risk of a person becoming a victim of a criminal offense for the police - from (6a) Annex III
 - (33) An AI system intended to be used to evaluate people's interests and behavior to assist the police in the detection of criminal offenses - from (6e) Annex III
 - (34) An AI system intended to be used to evaluate people's interests and behavior to assist the police in the investigation of criminal offenses - from (6e) Annex III
 - (35) An AI system intended to be used to assess the risk of irregular immigration posed by an immigrant for an immigration authority - from (7b) Annex III
 - (36) An AI system intended to be used to assess the security risk posed by an immigrant for an immigration authority - from (7b) Annex III
 - (37) An AI system intended to be used to examine visa applications for an immigration authority - from (7c) Annex III
 - (38) An AI system intended to be used to determine criminal sentences and jail time - from (8a) Annex III
- ### • Unacceptable-risk AI systems:
- (39) An AI system intended to be used to alter the face of a person in a political advertisement video to influence people's voting behavior - from Art. 5(1)a
 - (40) An AI system intended to be used to infer a person's income level to show targeted political advertisement - from Art. 5(1)b
 - (41) An AI system intended to be used to score people based on their social media profiles to determine whether they are provided financial assistance - from Art. 5(1)c
 - (42) An AI system intended to be used to estimate the risk of a person committing a criminal offense based on their personality for the police - from Art. 5(1)d
 - (43) An AI system intended to be used to store images of public transportation users for facial recognition - from Art. 5(1)e
 - (44) An AI system intended to be used to infer students' emotions during class - from Art. 5(1)f
 - (45) An AI system intended to be used to infer employees' emotions during work - from Art. 5(1)f
 - (46) An AI system intended to be used to identify a person's sexual orientation by their appearance and behavior for advertisement - from Art. 5(1)g
 - (47) An AI system intended to be used to identify a person's political orientation by their appearance and behavior for political advertisement - from Art. 5(1)g

- (48) An AI system intended to be used to monitor pedestrians to track them in real-time to prevent pickpocketing in a public park - from Art. 5(1)h

B Participants' Demographic Information

Table 5 presents the demographic distribution of our participants divided by country.

C Supplementary Methods

Section 3 presents an overview of our methodology. Here, we present a more detailed description of our study. After agreeing to the research terms, participants were shown the following definitions of AI system and risk from the AI Act (Art. 3 (1),(2)):

- An “artificial intelligence (AI) system” is a machine-based system that is designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment, and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments;
- “Risk” refers to the combination of the probability of an occurrence of harm and the severity of that harm.

Participants then answered two comprehension-check questions in random order—response options were also shown in random order:

- Q1) Which of the following best defines “risk” according to the definition we provided you?
- Combination of the probability of a harm occurring and the severity of such harm (correct answer);
 - Probability of experiencing delays in a project;
 - Severity of harm against animals.
- Q2) Which of the following best defines “artificial intelligence (AI) system” according to the definition we provided you?
- A system that operates with varying levels of autonomy and adaptiveness and generates outputs (correct answer);
 - A robot that demonstrates intelligence and operates in the physical world;
 - A type of software that can feel and demonstrate real emotions.

Participants were then asked to imagine four random AI systems, one at a time. For each AI system, participants were first asked “how much risk does the AI system pose to [right/value]?” Fundamental rights and protected values shown in Table 2 were formatted in a matrix table and presented in random order between participants. Participants rated the perceived risk to each right or concept using a 7-point scale anchored with 0 = No risk at all, 3 = Moderate risk, and 6 = Extreme risk. Participants were then asked “which of the following best describes your view on what should happen with the AI system” and could select that the AI systems “should be banned,” “should be regulated in some way,” or “should not be regulated.”

For the fourth (and final) AI system, participants were additionally asked to indicate “how important do you consider the inclusion of the following requirements in the regulation of the AI system?” Participants were shown all requirements presented in Table 3 formatted as a matrix table and in random order. Participants rated the

importance of each requirements using a 7-point scale anchored at 0 = Not important at all, 3 = Moderately important, 6 = Extremely important. Participants also indicated to what extent they agreed that the following should have a say in the regulation of the AI system using a 7-point scale (from strongly disagree to strongly agree): 1) business and industry, 2) individual citizens, 3) academics, 4) governments, and 5) non-governmental organizations (NGOs).

Participants were shown a instructed response attention check question (i.e., asked to select “Neither agree nor disagree”) and asked the extent to which they agreed with a several statements proposing different ways to regulate AI, the analysis of which we do not report in this paper. We concluded the study by asking several demographic questions, including gender, birth year, and education level, as reported in Table 5. Participants also indicated whether they had any training or work experience in professions “related to the legal system, including but not limited to law enforcement, legal services, legislation, etc” or “related to computer science, machine learning, artificial intelligence, or robotics.” We finally asked participants whether they paid attention to the study as a final attention check question.

C.1 Pilot Study

As explained in Section 3.4.1, we conducted a pilot ($N = 100$) to examine whether potential study participants understood the legal concepts we planned to investigate in our study. After making judgments about the perceived risk of AI systems, participants rated their understanding (“To what extent do you understand the concept of ...”) of the following concepts using a 7-point scale (0 = Not at all, 3 = Moderately well, 6 = Extremely well): 1) right to privacy and data protection; 2) right to equality and non-discrimination; 3) right to freedom of expression; 4) right to a legal defense and a fair trial; 5) right to engage in work and worker’s rights; 6) right to intellectual property; 7) right to liberty and security; 8) rule of law; 9) democracy; 10) AI system; and 11) risk. In a separate page, participants were also asked to explain each of these concepts using a text box (“Explain in your own words what you understand by ...”).

Although most of participants indicated that they understood these concepts moderately well (median reported understanding was four or higher for all concepts), participants’ open-ended explanations were sometimes incorrect. For instance, explanations of rule of law often described it as a “set of rules” rather than the idea that everyone—including institutions such as governments—is bound by the law. Another concepts with limited understanding was the “right to liberty and security,” which was double-barreled in this pilot; for the final study, we simplified it into the right to liberty. Our pilot results informed our decision to introduce fundamental rights and protected values using a language simpler than that used by legal scholarship, explaining the primary implications of each concept (see Table 2).

We also would like to acknowledge that participants’ explanations of “AI system” often mentioned large language models, likely because of their wide availability. Although this pilot also focused on specific AI applications, this finding further reinforced

Variable	United States	Germany	Spain	France
Sample Size	349 (24.70%)	330 (23.35%)	386 (27.32%)	348 (24.63%)
Gender				
Woman	196 (56.16%)	152 (46.06%)	164 (42.49%)	169 (48.56%)
Man	149 (42.69%)	177 (53.64%)	221 (57.25%)	177 (50.86%)
Non-Binary/Prefer not to respond	4 (1.15%)	1 (0.30%)	1 (0.26%)	2 (0.57%)
Age (Mean ± SD)	57.1 ± 17.5	52.6 ± 15.3	44.6 ± 13.0	49.7 ± 15.2
Education				
Less than a Bachelor’s degree	203 (58.17%)	214 (64.85%)	192 (49.74%)	180 (51.72%)
At least a Bachelor’s degree	141 (40.40%)	110 (33.33%)	192 (49.74%)	166 (47.70%)
Other/Prefer not to respond	5 (1.43%)	6 (1.82%)	2 (0.52%)	2 (0.57%)
Background in professions related to computer science				
Yes	59 (16.91%)	38 (11.52%)	122 (31.61%)	67 (19.25%)
No/I am not sure	290 (83.09%)	292 (88.48%)	264 (68.39%)	281 (80.75%)
Background in professions related to law				
Yes	38 (10.89%)	46 (13.94%)	66 (17.10%)	41 (11.78%)
No/I am not sure	311 (89.11%)	284 (86.06%)	320 (82.90%)	307 (88.22%)

Table 5: Demographic characteristics by country: counts with column-wise percentages for categorical variables; age is presented as mean and standard deviation (SD).

our decision to ground participants’ risk judgments in concrete AI applications (see Section 3.1).

D Supplementary Analysis

Tables 6 shows how participants’ opinions concerning AI regulation and prohibitions are correlated with the AI Act’s risk levels and perceived risk; see Section 4.2 for a detailed discussion of these results. Table 7 shows how participants rate the importance of obligations for AI systems according to their AI Act risk level, as discussed in Section 4.3.

Furthermore, Figure 6 presents the distribution of risk judgments of AI systems grouped by their AI Act risk level. Similarly, Figure 7 presents the distribution of importance judgments for the AI Act requirements explored in the study.

D.1 Differences Between Countries

We used mixed-effects ANOVA models to examine between-country differences in risk perceptions. We regressed participants’ risk judgments of a particular AI system to its AI Act risk level and participants’ country. We found that country was a significant factor in risk judgments for all EU fundamental rights and protected values. On average, participants from Germany and the US considered AI systems slightly less risky than participants from Spain and France (see Figure 5 in the main text). Due to the large number of pairwise comparisons, we present the detailed results in the study’s online repository: <https://tinyurl.com/PublicPerspective-AIAct>.

Figure 8 depicts the percentage of participants who believe AI systems should be prohibited or regulated based on their AI Act risk level and participants’ country. Table 8 presents how differences

between countries concerning the prohibition and regulation of AI systems. Participants from the United States were more likely to call for the prohibition of AI systems, whereas those from Germany were relatively less likely to suggest that AI systems should only be regulation, relying more on outright prohibitions.

Finally, Figure 9 shows how participants from different countries rate AI Act requirements similarly important. We also used mixed-effects ANOVA models to examine whether country is a significant factor in the perceived importance of AI Act requirement (while controlling for the AI system’s risk level). For brevity, we report the pairwise comparisons in the study’s online repository due to the large number of pairwise comparisons: <https://tinyurl.com/PublicPerspective-AIAct>. Although country is sometimes significant, the pairwise differences are again small.

	<i>Dependent Variable</i>			
	Prohibition		Any regulation (including ban)	
Risk Level/Perceived risk to...	(1)	(2)	(3)	(4)
Risk Level (AI Act) = High	0.051*** (0.010)		0.044*** (0.009)	
Risk Level (AI Act) = Unacceptable	0.091*** (0.013)		0.052*** (0.011)	
Health		0.008*** (0.003)		0.005* (0.003)
Safety		0.011*** (0.003)		0.002 (0.003)
Data Protection		0.002 (0.003)		0.011*** (0.003)
Equality		0.021*** (0.003)		0.011*** (0.003)
Expression		0.012*** (0.003)		0.001 (0.003)
Fair Trial		0.004 (0.003)		-0.002 (0.003)
Seek Employment		0.009*** (0.003)		-0.002 (0.002)
Work Conditions		-0.004 (0.003)		0.002 (0.003)
Intellectual Property		-0.008*** (0.003)		-0.001 (0.002)
Liberty		0.013*** (0.003)		0.003 (0.003)
Democracy		0.004 (0.003)		0.003 (0.003)
Rule of Law		0.013*** (0.003)		0.011*** (0.003)
Environment		-0.015*** (0.003)		-0.004* (0.003)
Constant	0.085*** (0.011)	-0.090*** (0.013)	0.894*** (0.009)	0.798*** (0.011)
Observations	5,652	5,652	5,652	5,652

Table 6: Linear regressions of people's opinions about prohibiting an AI system and regulating it in any way (including by prohibiting it). Models (1) and (3) explores differences based on AI Act risk levels, while models (2) and (4) examines the association between opinions about regulation and perceived risk of AI systems to EU fundamental rights and protected values. We include crossed random intercepts between AI systems and participants to account for multiple measurements. * $p < 0.1$; ** $p < 0.05$; * $p < 0.01$**

Comparison	Estimate	SE	df	<i>t</i>	<i>p</i> -value	Cohen's <i>d</i>
Risk Management			$F(2, 40.353) = 1.472, p = 0.2415$			
Minimal - High	-0.1799	0.1124	42.05	-1.601	0.3509	0.11
Minimal - Unacceptable	-0.1929	0.1390	43.53	-1.388	0.5168	0.12
High - Unacceptable	-0.0130	0.1221	45.83	-0.107	1.0000	0.01
Data Governance			$F(2, 42.721) = 4.6925, p < .05$			
Minimal - High	-0.3044	0.1040	41.55	-2.926	0.0166	0.20
Minimal - Unacceptable	-0.3019	0.1288	43.20	-2.344	0.0712	0.20
High - Unacceptable	0.0024	0.1133	45.88	0.021	1.0000	0.00
Technical Documentation			$F(2, 41.846) = 4.8011, p < .05$			
Minimal - High	-0.3454	0.1118	42.28	-3.089	0.0106	0.23
Minimal - Unacceptable	-0.2583	0.1383	43.67	-1.868	0.2054	0.17
High - Unacceptable	0.0871	0.1214	45.81	0.718	1.0000	0.06
Record-Keeping			$F(2, 41.807) = 5.3359, p < .01$			
Minimal - High	-0.3296	0.1047	41.93	-3.149	0.0090	0.22
Minimal - Unacceptable	-0.3153	0.1295	43.45	-2.435	0.0573	0.21
High - Unacceptable	0.0144	0.1138	45.84	0.126	1.0000	0.01
Compliance Instructions			$F(2, 43.195) = 2.4251, p = 0.1004$			
Minimal - High	-0.1947	0.1050	41.86	-1.854	0.2124	0.13
Minimal - Unacceptable	-0.2604	0.1299	43.40	-2.004	0.1541	0.17
High - Unacceptable	-0.0657	0.1142	45.85	-0.575	1.0000	0.04
Human Oversight			$F(2, 45.31) = 2.6351, p = 0.08268$			
Minimal - High	-0.2973	0.1307	42.73	-2.274	0.0842	0.18
Minimal - Unacceptable	-0.2425	0.1615	43.95	-1.502	0.4210	0.15
High - Unacceptable	0.0548	0.1415	45.75	0.387	1.0000	0.03
Performance			$F(2, 40.266) = 4.2401, p < .05$			
Minimal - High	-0.3138	0.1081	42.02	-2.904	0.0176	0.21
Minimal - Unacceptable	-0.2325	0.1337	43.50	-1.739	0.2673	0.15
High - Unacceptable	0.0813	0.1174	45.84	0.692	1.0000	0.05
Quality Management			$F(2, 42.274) = 2.4252, p = 0.1007$			
Minimal - High	-0.2152	0.1066	41.80	-2.018	0.1501	0.14
Minimal - Unacceptable	-0.2420	0.1320	43.37	-1.834	0.2204	0.16
High - Unacceptable	-0.0269	0.1160	45.86	-0.232	1.0000	0.02
Interaction Disclosure			$F(2, 42.803) = 0.4619, p = 0.6332$			
Minimal - High	-0.0994	0.1170	42.57	-0.849	1.0000	0.06
Minimal - Unacceptable	-0.1215	0.1446	43.85	-0.840	1.0000	0.08
High - Unacceptable	-0.0221	0.1268	45.77	-0.175	1.0000	0.01
Decision Disclosure			$F(2, 43.176) = 1.46, p = 0.2435$			
Minimal - High	-0.1329	0.1126	42.20	-1.180	0.7337	0.09
Minimal - Unacceptable	-0.2336	0.1392	43.62	-1.678	0.3014	0.15
High - Unacceptable	-0.1008	0.1222	45.82	-0.825	1.0000	0.07

Table 7: Pairwise comparisons of the perceived importance of AI Act requirements for an AI system based on their AI Act risk level. Bonferroni corrections were applied to adjust for multiple comparisons.

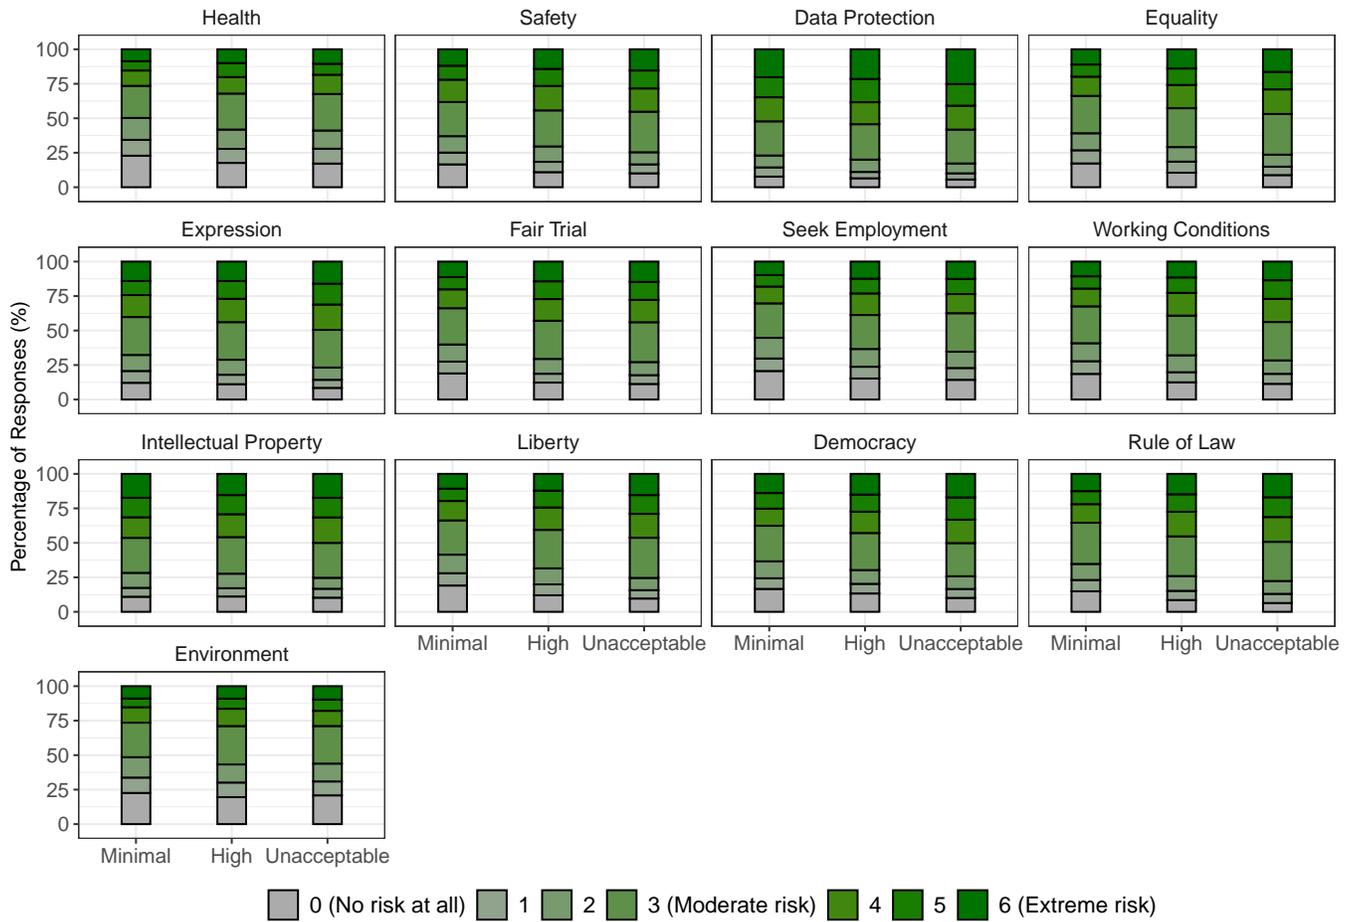


Figure 6: Distribution of judgments of perceived risk posed by AI systems to fundamental rights and protected values. We group judgments based on the AI Act risk level of the AI system being evaluated.

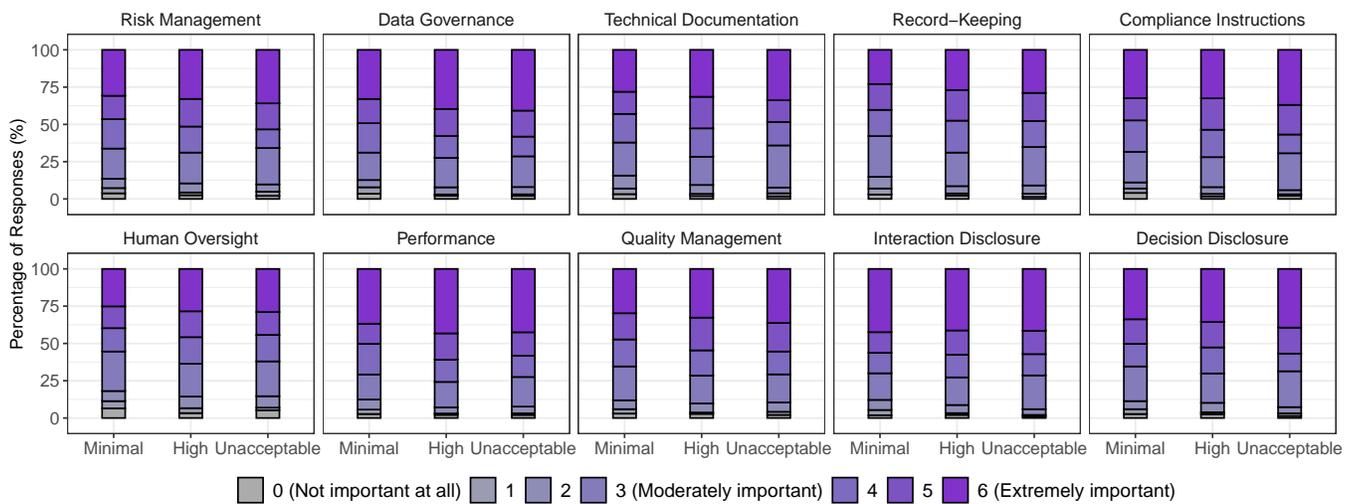


Figure 7: Distribution of judgments of perceived importance of AI Act requirements. We group judgments based on the AI Act risk level of the AI system being evaluated.

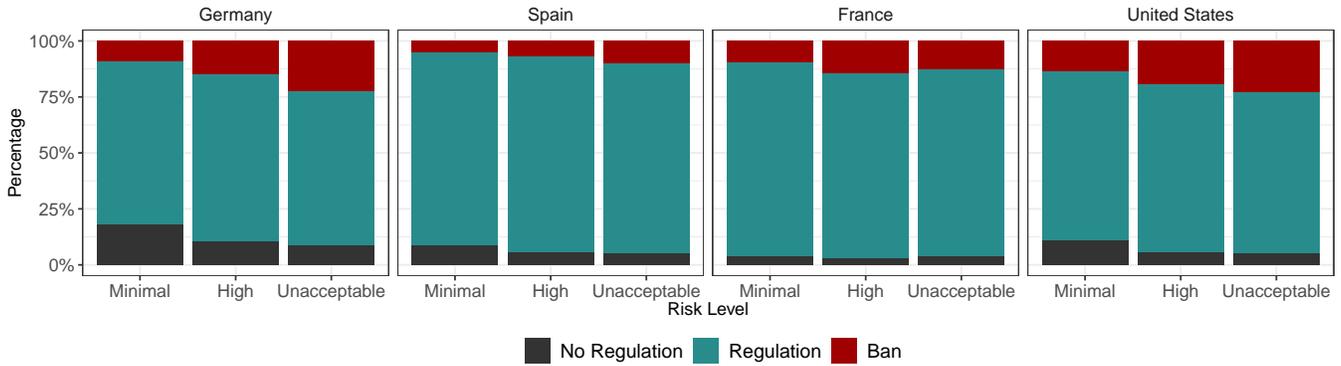


Figure 8: Distribution of participants' views towards regulating or banning AI systems depending on risk level and participants' country.

Comparison	Estimate	SE	df	<i>t</i>	<i>p</i> -value	Cohen's <i>d</i>
Prohibition			$F(3, 1408.6) = 10.559, p < .001$			
Germany - Spain	0.0780	0.0211	1408.94	3.702	0.0013	0.37
Germany - France	0.0230	0.0216	1408.89	1.066	1.0000	0.11
Germany - United States	-0.0354	0.0216	1409.39	-1.642	0.6053	0.17
Spain - France	-0.0550	0.0208	1408.95	-2.646	0.0493	0.26
Spain - United States	-0.1134	0.0208	1408.67	-5.464	<.0001	0.54
France - United States	-0.0584	0.0213	1408.43	-2.746	0.0367	0.28
Regulation			$F(3, 1408.65) = 9.983, p < .001$			
Germany - Spain	-0.0560	0.0154	1408.93	-3.627	0.0018	0.32
Germany - France	-0.0852	0.0158	1408.86	-5.385	<.0001	0.49
Germany - United States	-0.0516	0.0158	1409.49	-3.265	0.0067	0.30
Spain - France	-0.0292	0.0152	1408.92	-1.919	0.3314	0.17
Spain - United States	0.0044	0.0152	1408.58	0.288	1.0000	0.03
France - United States	0.0336	0.0156	1408.29	2.153	0.1891	0.19

Table 8: Pairwise comparisons between countries concerning participants' attitudes towards prohibiting and regulation AI systems. We average the results over the different AI Act risk levels when calculating the contrasts. We apply Bonferroni corrections to account for multiple comparisons.



Figure 9: Perceived importance of AI Act requirements based on risk level and participants' country. The figure includes standard errors for the risk-level group means, though they are too small to be visible.